STRUCTURAL PROPERTIES OF STOCHASTIC DYNAMIC PROGRAMS

JAMES E. SMITH

Fuqua School of Business, Duke University, Box 90120, Durham, North Carolina 27708-0120, jes9@mail.duke.edu

KEVIN F. McCARDLE

The Anderson School at UCLA, 110 Westwood Plaza, Suite D520, Los Angeles, California 90095-1481, kmccardl@anderson.ucla.edu

(Received May 2000; revised January 2001; accepted April 2001)

In Markov models of sequential decision processes, one is often interested in showing that the value function is monotonic, convex, and/or supermodular in the state variables. These kinds of results can be used to develop a qualitative understanding of the model and characterize how the results will change with changes in model parameters. In this paper we present several fundamental results for establishing these kinds of properties. The results are, in essence, "metatheorems" showing that the value functions satisfy property P if the reward functions satisfy property P and the transition probabilities satisfy a stochastic version of this property. We focus our attention on closed convex cone properties, a large class of properties that includes monotonicity, convexity, and supermodularity, as well as combinations of these and many other properties of interest.

1. INTRODUCTION

Stochastic dynamic programming models are pervasive in the operations research, management science, and economics literature. In the study of these models, the researchers often attempt to establish certain structural properties of the value functions—such as monotonicity, convexity, or supermodularity—to develop a qualitative understanding of the model and to derive comparative statics results that describe how the results of the model change with changes in model parameters. For example, in his classic paper on valuing options on stocks, Merton (1973) shows that the value of a call option on a stock is an increasing convex function of the underlying stock price. This result is then paired with Rothschild and Stiglitz's (1970) results on stochastic dominance to conclude that increases in the "riskiness" of the underlying stock increase the value of the call option. The real options literature is full of similar results in nonfinancial contexts (see, e.g., Dixit and Pindyck 1994).

The specific question that motivated this paper was the following: What is required to ensure that the value function for a dynamic program will be increasing and convex? What we found is that the arguments and assumptions we used to ensure that the value function is increasing and convex were perfectly analogous to those used to show that the value function is increasing. To show that the value function is increasing in the underlying state variable, it suffices to show that the reward function is increasing and that the transition probabilities are increasing in the sense of first-order stochastic dominance (or exhibit "positive persistence"). To establish convexity as well, it suffices to show that the reward functions are convex and the transition probabilities exhibit what we will call "stochastic convexity" defined in terms of a stochastic dominance ordering. Pushing further, we found that a variety of other properties can be established using the same argument.

In this paper, we present a set of "metatheorems" that provide conditions describing when value functions will possess certain properties. We begin in §2 by introducing three dynamic programming models that will be used to illustrate the results of the paper. In §3, we define the class of properties that we consider "closed convex cone" or C3 properties. This is a large class of properties that includes monotonicity, convexity, and supermodularity (and combinations of these properties) as well as many other properties of interest. Our main result in this section (Proposition 1) shows that each of these properties can be represented by a system of inequalities with a particular form. In §4, we study properties of conditional expectations and derive necessary and sufficient conditions for conditional expectations to satisfy a particular property. The main result in this section shows that conditional expectations will satisfy a C3 property P if the random variables satisfy a stochastic version of the property, defined by the same system of inequalities with stochastic dominance inequalities replacing the scalar inequalities. In §§5 and 6 we apply these fundamental results to Markov reward and decision processes. The main results there can be summarized as saying that the value functions will satisfy property P if the reward functions satisfy property P and the transition probabilities satisfy a stochastic version of this property.

Subject classifications: Dynamic programming: properties of stochastic models. Decision analysis: properties of sequential models.
Area of review: DECISION ANALYSIS.

Specific properties of dynamic programs have been studied in numerous examples and with varying degrees of generality. Monotonicity properties are particularly well studied. For example, Derman (1963) studies monotonicity properties in a model of optimal replacement rules for decaying systems (see also Ross 1983, p. 36–38). Stokey and Lucas (1989, p. 267–268) present general conditions that guarantee the value function will be monotonic in the underlying state variable; see also Müller (1997). Convexity and concavity properties have been derived in many specific models, including the options examples mentioned earlier. Hinderer (1984) studies the properties "increasing and convex" and "increasing and concave" with some generality. Supermodularity properties of dynamic programs are discussed in general and in specific examples in Topkis (1998). Our metatheorems unify these disparate results into a common framework and provide specific conditions that can be used to establish these and other properties in particular applications.

The analytic framework of this paper is closely related to, and to some extent inspired by, Athey (1998), who provides a general characterization of stochastic dominance relationships and studies properties of stochastic objective functions. Like us, she focuses on closed convex cone properties (her definition of "CCC properties" corresponds to what we call "C5 properties") and considers properties of stochastic objective functions in a general and abstract manner, similar to the treatment in our §§3 and 4. We believe, however, that the inequality representation provided by our Proposition 1—a new, necessary, and sufficient condition for closed convex cone properties—greatly simplifies and clarifies the study of properties of stochastic objective functions. We also differ in our choice of applications. While Athey applies her results to comparative statics problems in nondynamic settings (e.g., principal agent problems, portfolio problems), our interest focuses on stochastic dynamic programs. Müller (1997) also uses stochastic dominance relationships in his study of monotonicity properties of dynamic programs.

2. EXAMPLE APPLICATIONS

In this section, we introduce three dynamic programming models where we apply the results of the paper. The first two examples can be viewed as real options models, where we are interested in valuing options to invest in nonfinancial assets; these were the kind of examples that originally motivated this research. The third example is a stochastic multiproduct inventory model where the goal is to develop a minimum cost ordering policy. These examples are summarized in Table 1.

In each of these models, we index periods counting backwards from the terminal stage ($k = 0$) to the current stage where there are k periods to go. The value functions $v_k^*(x_k)$, describing the value in period k as a function of the then-current state x_k, can be represented recursively as

$$v_k^*(x_k) \equiv \sup_{a_k \in A_k} \{r_k(a_k, x_k) + \delta_k \mathrm{E}[v_{k-1}^*(\tilde{x}_{k-1}(a_k, x_k))]\} \quad \text{for } k > 0, \quad (1)$$
$$v_0^*(x_0) \equiv 0,$$

where a_k denotes the action selected in this period (from set A_k); $r_k(a_k, x_k)$ is the reward earned in state x_k when action a_k is selected, and $\tilde{x}_{k-1}(a_k, x_k)$ denotes the random next-period state, where the probabilities associated with the next-period state are conditioned on starting in state x_k and choosing action a_k.

The goal of the paper is to study properties of these value functions v_k^* and the limiting forms of these value functions $v^*(x) = limit_{k\to\infty} v_k^*(x)$ (if they exist), and present conditions that ensure that v_k^* and v^* satisfy these properties. Given the recursive structure of the models, proofs about properties of value functions usually proceed by induction. To show that v_k^* satisfies a property P for all k, we start by showing that v_0^* satisfies P and then showing that, if v_{k-1}^* satisfies P for period $k-1$, it satisfies P for period k as well. Properties of the limiting form v^* can be established by showing that the set of functions satisfying this property is closed and therefore contains this limiting function.

2.1. Copper Mine Model

Our first example is a discrete-time version of the model of a copper mine developed in Brennan and Schwartz (1985). Suppose a firm owns and operates a copper mine and must decide each period whether to operate the mine, temporarily close the mine, or abandon it altogether. There are costs associated with operating the mine when open and with maintaining the mine when closed, as well as costs associated with opening, closing, or abandoning the mine. The revenue associated with operating the mine depends on the price of copper, and these prices are assumed to follow a discrete-time geometric Brownian motion process. We assume that when the mine operates it produces copper at a fixed rate.

The state variable in this model can be written as $x_k = (s_k, p_k)$, where s_k indicates the status of the mine (*open, closed*, or *abandoned*) and p_k denotes the logarithm of the copper price in period k. Given p_k, the next period's log-price $\tilde{p}_{k-1}(p_k)$ is normally distributed with mean $p_k + \mu$ and standard deviation σ. The value function can then be written as

$$\begin{aligned} v_k^*(s_k, p_k) &= \max\big\{-c(abandoned, s_k) + \delta \mathrm{E}[v_{k-1}^*(abandoned, \tilde{p}_{k-1}(p_k))], \\ &\quad -c(closed, s_k) + \delta \mathrm{E}[v_{k-1}^*(closed, \tilde{p}_{k-1}(p_k))], \\ &\quad -c(open, s_k) + \gamma \exp(p_k) + \delta \mathrm{E}[v_{k-1}^*(open, \tilde{p}_{k-1}(p_k))]\big\}, \end{aligned}$$

Table 1.

Model	Copper Mine	Technology Adoption	Inventory Problem
State (x_k)	$x_k = (s_k, p_k)$ where $s_k =$ state of the mine at beginning of period $\in$ (*open, closed*, or *abandoned*) $p_k =$ log-price of copper $\in \mathbb{R}^1$	$x_k = (m_k, s_k)$ where $m_k =$ estimate of technology value at beginning of period $\in \mathbb{R}^1$ $s_k =$ precision of estimate $\in \mathbb{R}^{1+}$	$x_k = n$-vector of inventory levels for n products $\in \mathbb{R}^n$
Actions (a_k)	$a_k =$ state of mine at the end of period $\in$ (*open, closed* or *abandoned*)	$a_k =$ technology choice in period $\in$ (*adopt, investigate, reject*)	$a_k = n$-vector of order quantities for products; must be non-negative $\in \mathbb{R}^{n+}$
Rewards $(r_k(a_k, x_k))$	$r_k(a_k, s_k, p_k)$ $= \begin{cases} -c(a_k, s_k) + \gamma \exp(p_k) & \text{if } a_k = open \\ -c(a_k, s_k) & \text{otherwise} \end{cases}$	$r_k(a_k, m_k, s_k)$ $= \begin{cases} Am_k - K & \text{if } a_k = adopt \\ -c & \text{if } a_k = investigate \\ 0 & \text{if } a_k = reject \end{cases}$	$r_k(a_k, x_k)$ $= c(a_k) + \mathrm{E}[l(\tilde{x}_{k-1}(a_k, x_k))]$ where $\tilde{x}_{k-1}$ is defined below
Transitions $(\tilde{x}_{k-1}(a_k, x_k))$	$\tilde{x}_{k-1}(a_k, s_k, p_k; \sigma)$ $= (a_k, \tilde{p}_{k-1}(p_k))$ where $\tilde{p}_{k-1}$ is $\mathrm{N}(p_k + \mu, \sigma^2)$.	$\tilde{x}_{k-1}(a_k, m_k, s_k; t)$ $= \begin{cases} (-\infty, \infty) & \text{if } a_k = adopt \text{ or } reject \\ (\tilde{m}_{k-1}(m_k, s_k), s_k + t) & \text{otherwise} \end{cases}$ where $\tilde{m}_{k-1}$ is $\mathrm{N}(m_k, t/s_k(s_k + t))$.	$\tilde{x}_{k-1}(a_k, x_k)$ $= a_k + x_k - \tilde{z}_k$ where $\tilde{z}_k$ is the random demand in period k
Property P $(v(x_k))$	For each s_k, $v(s_k, p_k; \sigma)$ is increasing and convex in p_k and increasing in σ	$v(m_k, s_k; t)$ is increasing and convex in m_k, decreasing in s_k, satisfies the "mixing property" and is increasing in t.	$v(x_k)$ is convex in x_k
Property P^* $(v(a_k, x_k))$	For each a_k and s_k, $v(a_k, x_k, p_k\sigma)$ is increasing and convex in p_k and increasing in σ	For each a_k, $v(a_k, m_k, s_k; t)$ is increasing and convex in m_k, decreasing in s_k, satisfies the "mixing property," and is increasing in t.	$v(a_k, x_k)$ is jointly convex in a_k and x_k

where $c(s_{k-1}, s_k)$ denotes the costs of operating or maintenance costs and any costs of switching from a mine in state s_k to one in s_{k-1}. The decision to abandon the mine can be made irreversible by assuming the costs associated with switching out of an abandoned state are infinitely large. γ is the revenue per unit price of copper and given as the production rate less any proportional taxes or royalties. $\delta > 0$ is the discount factor.

We will show that the value function is *increasing and convex* in p_k for each s_k and further, that increases in "risk" induced by increasing the standard deviation σ for the price process lead to *increases* in the value function. As we will see, these conclusions do not depend on the specific assumptions about the price process and will hold for any process with transitions that are stochastically increasing and convex, in a sense to be described later.[1]

2.2. Technology Adoption Model

This model is developed in McCardle (1985) and further analyzed in Lippman and McCardle (1987). Suppose a firm is considering adopting a technology whose value is uncertain and denoted by μ. In period k the firm's uncertainty about μ is described by a normal distribution with mean m_k and precision s_k. (Precision is the reciprocal of variance.) The firm can reject the technology and receive zero, adopt the technology and receive an expected value of $Am_k - K$, or pay a cost $c > 0$ and investigate the technology further. Investigating the technology further yields a noisy observation of the value of the technology; the observation is normally distributed with mean μ and precision t. The firm then updates its estimate of μ according to Bayes' rule: Next period's mean estimate $\tilde{m}_{k-1}$ is normally distributed with mean m_k and precision $s_k(s_k + t)/t$, and next period's estimate will have precision $s_k + t$ (see, e.g., Pratt et al. 1995). The state variable for the problem is thus given by $x_k = (m_k, s_k)$, and the value function is given by

$$v_k^*(m_k, s_k) = \max\{0, Am_k - K, -c + \delta \mathrm{E}[v_{k-1}^*(\tilde{m}_{k-1}(m_k, s_k), s_k + t)]\}.$$

To place this in the format of Equation (1), where the process is assumed to continue for each action, we can assume that if the technology is adopted or rejected, in the next state the estimated value is $-\infty$ and the precision is ∞ (or sufficiently large negative and positive numbers), so the firm will choose to reject and receive a reward of zero in all subsequent periods. (This is done in Table 1.) Alternatively, we could expand the state space to include a variable indicating that the technology has been adopted or rejected in much the same way as the previous example where the state of the mine (*open, closed*, or *abandoned*) is explicitly modeled.

We will show that this value function is *increasing and convex* in m_k, *decreasing* in the prior precision s_k (more precision means there is less uncertainty about the value of the new technology) and *increasing* in the precision t of the observation. We will also show that the value functions satisfy the following *mixing property*: For any $\Delta \geqslant 0$, $s_k \geqslant 0$ and m_k, $0.5v_k^*(m_k+\delta, s_k+\Delta)+0.5v_k^*(m_k-\delta, s_k+\Delta) \geqslant v_k^*(m_k, s_k)$ where $\delta^2 = \Delta/(s_k(s_k+\Delta))$. Intuitively, the property captures the idea that earlier resolution of uncertainty about the value of the technology is preferred to later: The magnitude of δ is chosen so that the variance resolved in the estimate (a 50–50 chance of being revised to $m_k+\delta$ or $m_k-\delta$) is equal to the variance lost $(1/s_k - 1/(s_k+\Delta))$ by increasing the precision of the revised estimate from s_k to $s_k+\Delta$. This property will be used in showing that the value function is increasing in the precision of the observation t. Our assertion that the value function is increasing in t contradicts Lippman and McCardle (1987) who claim the opposite; their error can be traced back to incorrectly assuming the preposterior precision to be $(s_k+t)^2/t$ instead of $s_k(s_k+t)/t$.

2.3. Stochastic Inventory Model

This is a classic model, developed in Karlin (1960) and discussed in many places including Zipkin (2000). In each period, a firm observes its current inventories and decides how much of each of n products to order from its suppliers or produce. Demand for the products in each period is random and assumed to be independent from period to period and independent of the order quantities and inventory levels. Let x_k denote the n-vector of inventories at the beginning of period k; let a_k denote the n-vector of quantities of products ordered and delivered in period k; and let $\tilde{z}_k$ denote the random n-vector of demands for products in period k. The total available for sale in period k is given by $x_k + a_k$ and, assuming that unmet demand in one period carries over to the next, the next period starting inventory is $\tilde{x}_{k-1} = a_k + x_k - \tilde{z}_k$. (Unmet demand is thus treated as negative inventory.) The cost of ordering quantities a_k is given by a convex function $c(a_k)$, with the convexity reflecting potential diseconomies of scale in purchasing or manufacturing the product. There are holding costs associated with carrying unsold inventory as well as penalty costs associated with unmet demand; these combined costs are described by a convex loss function $l(x_{k-1})$. The value function for this model, describing the expected present value of costs under the optimal ordering policies, is

$$v_k^*(x_k) = \min_{a_k \geqslant 0}\{c(a_k) + \mathrm{E}[l(\tilde{x}_{k-1}(a_k, x_k)) + \delta v_{k-1}^*(\tilde{x}_{k-1}(a_k, x_k))]\},$$

where $\delta > 0$ is a discount factor. We will show that $v_k^*(x_k)$ is *jointly convex* in the vector of inventory levels x_k.

3. CLOSED CONVEX CONE PROPERTIES

We begin by defining the class of properties of value functions that we will study. The examples highlighted in the previous section—increasing, decreasing, convex, the mixing property—are all members of this general class of functions. In this section, we consider a number of additional examples and provide a general inequality-based representation of these properties. We conclude this section by considering examples of properties that may be of interest in some applications but are not C3 properties and cannot be represented in this form.

3.1. Definitions

We consider real-valued functions defined on a parameter space Θ with typical element θ. While most of our examples concern properties of functions defined on $\mathbb{R}^1$ or $\mathbb{R}^n$, our results will hold for real-valued functions defined on arbitrary sets Θ. A set of functions $\mathcal{F}$ forms a *convex cone* if for any $f_1, f_2 \in \mathcal{F}$ and scalars $a, b \geqslant 0$, $af_1 + bf_2 \in \mathcal{F}$.

DEFINITION 1. P is a *closed convex cone property* (C3 *property*) if the set of functions satisfying P forms a closed convex cone in the topology of pointwise convergence.

DEFINITION 2. P is a *closed convex cone containing constants property* (C5 *property*) if it is a C3 property and constant functions satisfy P.

To illustrate these definitions consider the following examples of C3 and C5 properties.

1. *Constant.* A function f is constant on Θ if there exists a constant c such that $f(\theta) = c$ for all θ in Θ. This is a C5 property.
2. *Nonnegative.* A function f is nonnegative on Θ if $f(\theta) \geqslant 0$ for all θ in Θ. This is a C3 property but not a C5 property because negative constants functions (e.g., $-\mathbf{1}$) do not satisfy this condition.
3. *Increasing.* A function f is increasing on an ordered space Θ if $f(\theta_1) \leqslant f(\theta_2)$ for all θ_1, θ_2 such that $\theta_1 \leqslant \theta_2$. This is a C5 property.
4. *Convex.* A function f is convex on a convex set Θ if $f(\phi\theta_1 + (1-\phi)\theta_2) \leqslant \phi f(\theta_1) + (1-\phi) f(\theta_2)$ for all θ_1, θ_2, in Θ and ϕ, $0 \leqslant \phi \leqslant 1$. This is a C5 property.
5. *Subadditive.* A function f is subadditive on Θ if $f(\theta_1 + \theta_2) \leqslant f(\theta_1) + f(\theta_2)$ for all θ_1, θ_2 in Θ; here we assume that $\theta_1 + \theta_2$ is in Θ whenever θ_1 and θ_2 are. This is a C3 property but not a C5 property because negative constants functions (e.g., $-\mathbf{1}$) do not satisfy this condition.

6. *Supermodular.* A function f is supermodular on a lattice Θ if $f(\theta_1)+f(\theta_2) \leqslant f(\theta_1 \wedge \theta_2)+f(\theta_1 \vee \theta_2)$ for all θ_1, θ_2 in Θ. Here $\wedge$ and $\vee$ denote componentwise minimization and maximization, respectively; Θ is a lattice if $\theta_1 \wedge \theta_2$ and $\theta_1 \vee \theta_2$ are in Θ whenever θ_1 and θ_2 are. This is a C5 property.

If the functions $\{f_l\}_{l\in L}$ satisfy a C3 property (or a C5 property), the functions $\{-f_l\}_{l\in L}$ satisfy a potentially different C3 (or C5) property. For example, we can generate more C3 and C5 properties by considering the negatives of Examples 2–6: a function f is *nonpositive, decreasing, concave, superadditive,* or *submodular* if $-f$ is *positive, increasing, convex, subadditive*, or *supermodular*, respectively. Intersections of C3 properties are C3 properties and intersections of C5 properties are also C5 properties. Thus, for example, *increasing and convex* is a C5 property because increasing and convex are both C5 properties. Similarly, *linear*, as the intersection of convex and concave, is a C5 property.

The characteristics of C3 properties arise naturally given the recursive structure of dynamic programs. First, the terminal value function $v_0^*(x_0) \equiv 0$ will automatically satisfy any C3 property, because the set of functions satisfying P forms a cone. Second, whenever the reward function (r) and the expected continuation value ($\mathrm{E}[v_{k-1}^*]$) both satisfy a C3 property P, the current value ($r+\mathrm{E}[v_{k-1}^*]$) will also satisfy P because this set of functions is a convex cone. Finally, if the set of functions is closed in the topology of pointwise convergence, the limiting value functions $limit_{k\to\infty} v_k^*$ (when they exist) will be contained in the same set. As discussed in the beginning of §2, these are key steps in proofs establishing properties of value functions.

3.2. Inequality Representation of C3 Properties

Reviewing the examples of C3 and C5 properties given above, we see that many of them are written in the form of a comparison between finite weighted sums of function evaluations at specific values of θ. Our first result shows that all C3 and C5 properties can be represented in this way. We believe that this is a new result. The proof is somewhat involved, but given this result, the rest of our results are quite easy to prove.

PROPOSITION 1. *A property P is a* C3 *property if and only if there exists a collection, indexed by α in $\mathcal{A}$, of finite sets of points $\{\theta_\beta\}_{\beta\in B_\alpha}, \{\theta_\gamma\}_{\gamma\in\Gamma_\alpha}$ and positive weights $\{\lambda_\beta\}_{\beta\in B_\alpha}, \{\lambda_\gamma\}_{\gamma\in\Gamma_\alpha}$ that define a test of satisfaction of the form: f satisfies P if and only if*

$$\sum_{\beta\in B_\alpha} \lambda_\beta f(\theta_\beta) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma f(\theta_\gamma) \quad \text{for all } \alpha \text{ in } \mathcal{A}. \tag{2}$$

Furthermore, if P is a C5 *property, for each α in $\mathcal{A}$, we can normalize the weights to sum to one.*

Thus, all C3 properties can be represented by a set of simple inequalities. The set $\mathcal{A}$ (with generic element α) indexes the inequalities and is potentially infinite. The sets B_α and Γ_α (with generic elements β and γ) index the values and weights used in the summations on each side of each inequality; these sets are finite, and each inequality is a comparison between weighted sums of function evaluations. The examples below will help clarify the interpretation of the proposition.

1. *Increasing*. The index set $\mathcal{A}$ consists of the set of pairs $\alpha=(\theta_1,\theta_2)$ satisfying $\theta_1 \leqslant \theta_2$. For each α, the index sets B_α and Γ_α are both singletons with points $\theta_\beta=\theta_1$ and $\theta_\gamma=\theta_2$ and weights $\lambda_\beta=\lambda_\gamma=1$. The test of satisfaction is thus: f is increasing if and only if $f(\theta_1)\leqslant f(\theta_2)$ for all θ_1, θ_2 such that $\theta_1\leqslant\theta_2$.

2. *Convex*. The index set $\mathcal{A}$ consists of the infinite set of triplets $\alpha=(\phi,\theta_1,\theta_2)$ with θ_1,θ_2 in Θ and scalar ϕ satisfying $0\leqslant\phi\leqslant 1$. B_α is a singleton with point $\theta_\beta=\phi\theta_1+(1-\phi)\theta_2$ being the convex combination of θ_1 and θ_2 and weight $\lambda_\beta=1$. Γ_α contains two elements, γ_1 and γ_2, with $\theta_{\gamma1}=\theta_1$, $\theta_{\gamma2}=\theta_2$ and $\lambda_{\gamma1}=\phi$, and $\lambda_{\gamma2}=(1-\phi)$. The test of satisfaction is thus: f is convex if and only if $f(\phi\theta_1+(1-\phi)\theta_2)\leqslant\phi f(\theta_1)+(1-\phi)f(\theta_2)$ for all θ_1,θ_2, in Θ and $\phi, 0\leqslant\phi\leqslant 1$.

3. *Constant.* Pick some arbitrary $\theta^*\in\Theta$. The test of satisfaction can be written as: f is constant if and only if $f(\theta)\leqslant f(\theta^*)$ for all θ in Θ and $f(\theta^*)\leqslant f(\theta)$ for all θ in Θ.

4. *Linear.* Linearity can be represented by requiring f to be both convex and concave. The test of satisfaction can thus be written as requiring $f(\phi\theta_1+(1-\phi)\theta_2)\leqslant\phi f(\theta_1)+(1-\phi)f(\theta_2)$ for all θ_1,θ_2 in Θ and $0\leqslant\phi\leqslant 1$ and $\phi f(\theta_1)+(1-\phi)f(\theta_2)\leqslant f(\delta\theta_1+(1-\phi)\theta_2)$ for all θ_1,θ_2 in Θ and $0\leqslant\phi\leqslant 1$.

As illustrated in Examples 3 and 4, intersecting two properties—i.e., requiring a function to satisfy two C3 or C5 properties—corresponds to taking the union of the corresponding sets of inequalities. If the set of inequalities $\mathcal{A}_1$ represents the first property and $\mathcal{A}_2$ represents the second, then $\mathcal{A}_1\cup\mathcal{A}_2$ represents the intersection of the two properties.

It is clear that any property that can be represented as in Proposition 1 is a C3 property: If two functions f_1 and f_2 satisfy inequality (2) for some α in $\mathcal{A}$, then for scalars $a, b\geqslant 0$, af_1+bf_2 will also satisfy the inequality. The functions satisfying these inequalities thus form a convex cone. Moreover, if a series of functions $\{f_0, f_1, f_2, \dots\}$ each satisfying (2) converges pointwise to some function f, the limiting function will also satisfy (2). It is harder to prove that all C3 properties can be represented this way. The proof proceeds in two steps. First, we represent the closed convex cone of functions as the intersection of a set (indexed by α in $\mathcal{A}$) of closed half-spaces, each of which can be represented as an inequality for a continuous linear functional. We then show that each of these linear functionals can be represented as a comparison between finite positive sums of function evaluations at specific points. The proof is given in the appendix.

3.3. Other Properties

Though many useful properties are C3 properties, there are other properties that may be of interest in some applications that cannot be represented in this form. Examples include "bounded," "continuous," "differentiable" ($df(x)/dx$ exists), and "integrable" ($\int |f(x)|\,dx$ exists and is finite). In these examples the functions satisfying these properties form a convex cone, but the cone is not closed in the topology of pointwise convergence. The proof of Proposition 1 exploits properties of the topology of pointwise convergence in two ways. First, rather technically, the assumption allows us to establish the existence of the separating hyperplanes without requiring an interior point. Second, we use properties of this topology to represent the hyperplanes in the form of the finite sums of inequality (2). Other properties that may be of interest but do not form a convex cone include the "single-crossing property" and "quasisupermodularity." These properties were introduced in Milgrom and Shannon (1994) and are discussed in Topkis (1998).

4. PROPERTIES OF CONDITIONAL EXPECTATIONS

In this section, we study properties of conditional expectations—what is required to ensure that $\mathrm{E}[u(\tilde{x}(\theta))]$ satisfies property P in θ for functions u in some set U? This is an important step in showing that properties are preserved through the dynamic programming recursion in Equation (1). If we can show that $\mathrm{E}[v^*_{k-1}(\tilde{x}_{k-1}(a_k, x_k))]$ satisfies a C3 property P and the reward function $r_k(a_k, x_k)$ satisfies the same property, we can conclude that the value associated with a particular action, $r_k(a_k, x_k) + \delta\mathrm{E}[v^*_{k-1}(\tilde{x}_{k-1}(a_k, x_k))]$, will also satisfy this property. The main result of this section is a necessary and sufficient condition for $\mathrm{E}[u(\tilde{x}(\theta))]$ to satisfy property P in θ for all functions u in set U.

4.1. Stochastic Dominance Relations

Let $\tilde{x}(\theta)$ denote a randomly selected state chosen from some set X according to a probability measure $\mu(\theta)$ where θ conditions the probabilities associated with the random selection. In the dynamic programming applications, the parameter θ will represent the current state and action, and the random state $\tilde{x}(\theta)$ will correspond to the uncertain next state ($\tilde{x}_{k-1}(a_k, x_k)$). Given a real-valued function u defined on X, we define the expectations as $\mathrm{E}[u(\tilde{x}(\theta))] \equiv \int u\, d\mu(\theta)$.[2]

The following definition of dominance includes the standard stochastic dominance relations as special cases and has been used by several authors, including Athey (1998) and Müller (1997).

Definition 3. $\tilde{x}(\theta_1)$ *dominates* $\tilde{x}(\theta_2)$ on U if $\mathrm{E}[u(\tilde{x}(\theta_1))] \geqslant \mathrm{E}[u(\tilde{x}(\theta_2))]$ for all $u \in U$.

Each set of functions U thus defines a dominance partial ordering. We abbreviate "$\tilde{x}(\theta_1)$ dominates $\tilde{x}(\theta_2)$ on U" by writing $\tilde{x}(\theta_1) \succsim_U \tilde{x}(\theta_2)$ and let $\precsim_U$ denote the opposite ordering where the scalar inequality $\leqslant$ replaces the $\geqslant$ appearing in the definition of $\succsim_U$. Here we talk about dominance holding between the random states. Given how the random states are defined by the probability measures, we can equivalently talk about dominance in terms of the measures and say, for example, $\mu(\theta_1) \succsim_U \mu(\theta_2)$ instead of $\tilde{x}(\theta_1) \succsim_U \tilde{x}(\theta_2)$.

The following are some familiar examples of dominance relations. In the univariate examples (1–3), we let F_1 and F_2 denote the cumulative distribution functions for $\tilde{x}(\theta_1)$ and $\tilde{x}(\theta_2)$, respectively.

1. *Increasing*. $X = \mathbb{R}^1$, $U =$ increasing functions. In this case, $\succsim_U$ corresponds to the usual first-order stochastic dominance ordering and we can check for dominance using the well-known result that $\tilde{x}(\theta_1)$ dominates $\tilde{x}(\theta_2)$ on U if and only if $F_1(x) \leqslant F_2(x)$ for all x.

2. *Concave*. $X = \mathbb{R}^1$, $U =$ concave functions. In this case, $\succsim_U$ corresponds to second-order stochastic dominance, as defined in Rothschild and Stiglitz (1970). We can check for dominance using their result: $\tilde{x}(\theta_1)$ dominates $\tilde{x}(\theta_2)$ on U if and only if $\mathrm{E}[\tilde{x}(\theta_1)] = \mathrm{E}[\tilde{x}(\theta_2)]$ (i.e., the expected values are equal) and $\int_{-\infty}^{a} F_1(x)\,dx \leqslant \int_{-\infty}^{a} F_2(x)\,dx$ for all a.

3. *Increasing and Concave*. $X = \mathbb{R}^1$, $U =$ increasing concave functions. Here $\succsim_U$ corresponds to second-order monotonic stochastic dominance, which can be checked using the well-known result that $\tilde{x}(\theta_1)$ dominates $\tilde{x}(\theta_2)$ on U if and only if $\int_{-\infty}^{a} F_1(x)\,dx \leqslant \int_{-\infty}^{a} F_2(x)\,dx$ for all a.

4. *Supermodular*. $X = \mathbb{R}^2$, $U =$ supermodular functions. $\tilde{x}(\theta_1)$ dominates $\tilde{x}(\theta_2)$ on U if and only if $F_1(x_1, x_2) \leqslant F_2(x_1, x_2)$ and $F_1(x_1, \text{-}) = F_2(x_1, \text{-})$ and $F_1(\text{-}, x_2) = F_2(\text{-}, x_2)$ for all $(x_1, x_2) \in \mathbb{R}^2$. Here $F_i(x_1, x_2)$ denotes the bivariate cumulative distribution and $F_i(x_1, \text{-})$ and $F_i(\text{-}, x_2)$ denote the univariate marginal distributions for $\tilde{x}(\theta_1)$ and $\tilde{x}(\theta_2)$, respectively (see Athey 1998).

Many other examples of dominance orderings can be found in Shaked and Shantikumar (1994). Athey (1998) provides a general characterization of dominance relations and identifies conditions for one set of functions to serve as a "test set" for establishing dominance on another, larger set of functions. For example, dominance on the set of step functions is sufficient to establish dominance on the set of increasing functions.

4.2. Properties of Conditional Expectations

Using the inequality-based representation of C3 and C5 properties developed in Proposition 1, we can characterize the necessary and sufficient conditions for the conditional expectations $\mathrm{E}[u(\tilde{x}(\theta))]$ to satisfy P in θ for u in some set of functions U. Intuitively, the condition requires the family of measures $\mu(\theta)$ defining $\tilde{x}(\theta)$ to satisfy the system of inequalities defining property P, except the inequalities are now interpreted in the sense of the dominance relation $\succsim_U$ defined for the set of functions U.

Proposition 2. *Let P be a C3 property represented as in Proposition 1. Then $\mathrm{E}[u(\tilde{x}(\theta))]$ satisfies P on Θ for all u in U if and only if the measures $\mu(\theta)$ satisfy*

$$\sum_{\beta\in\beta_\alpha} \lambda_\beta\mu(\theta_\beta) \precsim_U \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma\mu(\theta_\gamma) \quad \textit{for all } \alpha \textit{ in } \mathscr{A}. \tag{3}$$

Proof. Given that the system of inequalities (2) represents $P, \mathrm{E}[u(\tilde{x}(\theta))] \equiv \int u\, d\mu(\theta)$ satisfies P for all u in U if and only if, for each α in $\mathscr{A}$, the corresponding mixed measures $\mu_\alpha^1 \equiv \sum_{\beta\in\beta_\alpha} \lambda_\beta\mu(\theta_\beta)$ and $\mu_\alpha^2 \equiv \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma\mu(\theta_\gamma)$ satisfy $\int u\, d\mu_\alpha^1 \leqslant \int u\, d\mu_\alpha^2$ for all u in U or, equivalently, $\mu_\alpha^1 \precsim_U \mu_\alpha^2$. □

If the condition of the proposition holds, we say the random state $\tilde{x}(\theta)$ *satisfies P in the sense of* some stochastic dominance ordering $\succsim_U$, $\tilde{x}(\theta)$ *satisfies P on U*, or more compactly $\tilde{x}(\theta)$ *satisfies $P(\succsim_U)$*.[3]

Some examples will help clarify this result.

1. *Increasing/Increasing.* $\Theta = X = \mathbb{R}^1$; $P =$ increasing; $U =$ increasing functions. Proposition 2 says $\mathrm{E}[u(\tilde{x}(\theta))]$ is increasing in θ for all increasing u if and only if $\mu(\theta_1) \precsim_U \mu(\theta_2)$ for all $\theta_1 \leqslant \theta_2$, or, in other words, if and only if $\tilde{x}(\theta)$ is *stochastically increasing in θ in the sense of first-order stochastic dominance*.

2. *Concave/Increasing.* $\Theta = X = \mathbb{R}^1$; $P =$ concave; $U =$ increasing functions. Proposition 2 says $\mathrm{E}[u(\tilde{x}(\theta))]$ is concave in θ for all increasing u if and only if $\phi\mu(\theta_1) + (1-\phi)\mu(\theta_2) \precsim_U \mu(\phi\theta_1 + (1-\phi)\theta_2)$ for all θ_1, θ_2, and $\phi, 0 \leqslant \phi \leqslant 1$, or in other words, if and only if $\tilde{x}(\theta)$ *is stochastically concave in θ in the sense of first-order stochastic dominance.*

3. *Concave/Concave.* $\Theta = X = \mathbb{R}^1$; $P =$ concave; $U =$ concave functions. Proposition 2 says $\mathrm{E}[u(\tilde{x}(\theta))]$ is concave in θ for all concave u if and only if $\phi\mu(\theta_1) + (1-\phi)\mu(\theta_2) \precsim_U \mu(\phi\theta_1 + (1-\phi)\theta_2)$ for all θ_1, θ_2, and $\phi, 0 \leqslant \phi \leqslant 1$, or in other words, if and only if $\tilde{x}(\theta)$ *is stochastically concave in θ in the sense of second-order stochastic dominance*.

4. *Increasing and Concave/Increasing and Concave.* $\Theta = X = \mathbb{R}^1$; $P =$ increasing and concave; $U =$ increasing, concave functions. Proposition 2 says $\mathrm{E}[u(\tilde{x}(\theta))]$ is increasing and concave in θ for all increasing and concave u if and only if $\tilde{x}(\theta)$ *is stochastically increasing and concave in θ in the sense of second-order monotonic stochastic dominance.* Note that $\tilde{x}(\theta)$ being stochastically concave in θ in the sense of first-order stochastic dominance is sufficient for this result, but stronger than necessary.

5. *Supermodular/Increasing.* $\Theta = \mathbb{R}^n$; $X = \mathbb{R}^1$; $P =$ supermodular; $U =$ increasing functions. Proposition 2 says $\mathrm{E}[u(\tilde{x}(\theta))]$ is supermodular in θ for all increasing u if and only if $\mu(\theta_1) + \mu(\theta_2) \precsim_U \mu(\theta_1 \wedge \theta_2) + \mu(\theta_1 \vee \theta_2)$ for all θ_1, θ_2 in Θ, or in other words, if and only if μ *is stochastically supermodular in the sense of first-order stochastic dominance*.

Examples 1, 2, and 5 can be found in Topkis (1998; Theorem 3.9.1 and its corollary, p. 160–161). Example 4 is discussed in Shaked and Shantikumar (1994; Theorem 6.A.6, pp 172–173). Note that, as illustrated in Examples 2 and 5, we may choose the property P and set U independently; the functions u appearing inside the integrals need not satisfy the property P. Moreover, the set of functions U need not form a closed convex cone (i.e., correspond to a C3 property) though it does in each of our examples and in most well-known stochastic dominance relationships (Athey 1998).

4.3. Example Applications

Our example dynamic programming applications involve variations and combinations of the abstract examples discussed above, as well as some other properties.

Copper Mine. In this example, we will show that the price transitions are increasing and convex in the sense of second-order, monotonic stochastic dominance: Formally, we take $\Theta = X = \mathbb{R}^1$, with p_k and p_{k-1} being the typical elements of Θ and X, respectively; $P =$ increasing and concave; $U =$ increasing, concave functions. We show that $\tilde{p}_{k-1}(p_k)$ satisfies $P(\succsim_U)$. Recall that in this model the next period's log-price $\tilde{p}_{k-1}$ is normally distributed with mean $p_k + \mu$ and standard deviation σ. An increase in p_k thus leads to a first-order stochastic dominance increase in the distribution for $\tilde{p}_{k-1}$: Thus, $\tilde{p}_{k-1}(p_k)$ is increasing on U. To prove that the transitions are also convex on U, we must show that, for any increasing convex function u, prices p_k^1, p_k^2, and $\phi, 0 \leqslant \phi \leqslant 1$

$$\mathrm{E}[u(\tilde{p}_{k-1}(p_k^\phi))] \leqslant \phi\mathrm{E}[u(\tilde{p}_{k-1}(p_k^1))] + (1-\phi)\mathrm{E}[u(\tilde{p}_{k-1}(p_k^2))],$$

where $p_k^\phi = \phi p_k^1 + (1-\phi)p_k^2$. Given how the change in current price p_k shifts the distribution for $\tilde{p}_{k-1}$, this condition is equivalent to

$$\mathrm{E}[u(p_k^\phi + \tilde{y})] \leqslant \phi\mathrm{E}[u(p_k^1 + \tilde{y})] + (1-\phi)\mathrm{E}[u(p_k^2 + \tilde{y})],$$

where $\tilde{y}$ is normally distributed with mean μ and standard deviation σ. This condition follows from the convexity of u because, for any y,

$$u(p_k^\phi + y) \leqslant \phi u(p_k^1 + y) + (1-\phi)u(p_k^2 + y).$$

Thus, the price transitions $\tilde{p}_{k-1}(p_k)$ are convex in p_k, as well as increasing, on the set of increasing and convex functions. We can also appeal to the results of Rothschild and Stiglitz (1970) and note that, because increasing σ corresponds to an "increase in risk" or "mean-preserving spread" in the sense of second-order stochastic dominance, for any fixed price p_k and increasing convex function u (in fact for any convex function u), $\mathrm{E}[u(\tilde{p}_{k-1}(p_k))]$ is increasing in σ.

Note that these properties of the transitions do not depend on the specific assumption that prices follow a Brownian motion process. One could, for example, assume that prices follow a mean-reverting process with the next period's log-price ($\tilde{p}_{k-1}$) being normally distributed with mean $\kappa p_k + (1-\kappa)\bar{p}$ and standard deviation σ. Here κ controls the rate of mean reversion and $\bar{p}$ is the long-run mean to which prices revert. In this process, changes in the conditioning variables shift the distributions in the same way as before (though multiplied by κ) and the transitions satisfy these same properties. We can use the same arguments with nonnormal transitions, provided changes in the conditioning variables shift the distributions in the same way. The same argument with the inequalities reversed shows that transitions having this "linear translation property" will also be increasing and concave on the set of increasing and concave functions.

Technology Adoption. In this example, we first show that the transitions $(\tilde{m}_{k-1}(m_k, s_k), s_{k-1}(s_k))$ are increasing and convex in the estimated value of the technology (m_k) and decreasing in the precision of this estimate (s_k) for functions $u(m_{k-1}, s_{k-1})$ that are increasing and convex in m_{k-1} and decreasing in s_{k-1}. Here $\Theta = X = \mathbb{R}^1 \times \mathbb{R}^+$ with typical elements (m_k, s_k) for Θ and (m_{k-1}, s_{k-1}) for X; $P =$ increasing and concave in m_k and decreasing in s_k; $U =$ functions that are increasing and concave in m_{k-1} and decreasing in s_{k-1}. Given the current set of state variables m_k and s_k, the next-period estimate of the value $\tilde{m}_{k-1}$ is normally distributed with mean m_k and precision $s_k(s_k+t)/t$ and the next-period precision s_{k-1} is $s_k + t$. Changes in the current estimate of value m_k thus shift the distributions for the next-period estimate $\tilde{m}_{k-1}$ in exactly the same way as the previous example and we can use the same arguments to show that the transitions are stochastically increasing and convex in m_k for functions that are increasing and convex in m_{k-1}. Likewise, just as increases in σ in the previous example led to an increase in expected values for increasing convex functions, here an increase in the precision of the prior estimate s_k increases the precision of $\tilde{m}_{k-1}$ (given by $s_k(s_k+t)/t$) and increases the precision of the next-period estimate $(s_{k-1} = s_k + t)$. Because an increase in precision corresponds to a decrease in variance, both effects lead to a decrease in $\mathrm{E}[u(\tilde{m}_{k-1}(m_k, s_k), s_{k-1}(s_k))]$ for functions $u(m_k, s_k)$ that are increasing and convex in m_k and decreasing in s_k.

In the appendix we show that the transitions are increasing in the precision of the observation in each period (t) and satisfy the mixing property discussed in §1 for functions that also satisfy the mixing property.

Stochastic Inventory Model. In this example, we show that the transitions $\tilde{x}_{k-1}(a_k, x_k)$ are jointly convex in the n-vector of current inventory levels x_k and the n-vector of order quantities a_k for functions that are similarly jointly convex. Here $\Theta = \mathbb{R}^{n+} \times \mathbb{R}^n$ with typical element (a_k, x_k) and $X = \mathbb{R}^n$ with typical element x_{k-1}, and we show that the transitions $\tilde{x}_{k-1}(a_k, x_k)$ satisfy $P(\precsim_U)$ where P is joint convexity in a_k and x_k and U denotes the set of functions that are convex in x_{k-1}. We must show that for any convex function u, n-vectors a_k^1, a_k^2, x_k^1, x_k^2, and scalar ϕ, $0 \leqslant \phi \leqslant 1$,

$$\mathrm{E}[u(\tilde{x}_{k-1}(a_k^\phi, x_k^\phi))] \leqslant \phi\,\mathrm{E}[u(\tilde{x}_{k-1}(a_k^1, x_k^1))] + (1-\phi)\,\mathrm{E}[u(\tilde{x}_{k-1}(a_k^2, x_k^2))],$$

where $a_k^\phi = \phi a_k^1 + (1-\phi)a_k^2$ and $x_k^\phi = \phi x_k^1 + (1-\phi)x_k^2$. Because the next-period inventory $\tilde{x}_{k-1}$ is given as $a_k + x_k - \tilde{z}_k$ and the demand $\tilde{z}_k$ is independent of a_k and x_k, this is equivalent to

$$\mathrm{E}[u(a_k^\phi + x_k^\phi - \tilde{z}_k)] \leqslant \phi\mathrm{E}[u(a_k^1 + x_k^1 - \tilde{z}_k)] + (1-\phi)\mathrm{E}[u(a_k^2 + x_k^2 - \tilde{z}_k)].$$

The inequality then follows from the assumption that u is convex because for any z_k,

$$u(a_k^\phi + x_k^\phi - z_k) \leqslant \phi u(a_k^1 + x_k^1 - z_k) + (1-\phi)u(a_k^2 + x_k^2 - z_k).$$

Thus, the transitions are jointly stochastically convex in a_k and x_k on the set of convex functions. We can reverse the inequalities in this argument and show that transitions are jointly stochastically concave in a_k and x_k on the set of concave functions.

5. CHARACTERIZING VALUE FUNCTIONS FOR MARKOV REWARD PROCESSES

Before considering the full stochastic dynamic programming model, we first consider Markov reward processes of the form

$$v_k(x_k) \equiv r_k(x_k) + \delta_k \mathrm{E}[v_{k-1}(\tilde{x}_{k-1}(x_k))] \quad \text{for } k > 0,$$

$$v_0(x_0) \equiv 0,$$

where v_k is the value function with k periods remaining, x_k is the period-k state variable selected from some state space X, r_k is the period reward given that you start in state x_k, $\delta_k > 0$ is a discount factor, and $\tilde{x}_{k-1}(x_k)$ denotes the random next-period state (also selected from X) with transition probabilities conditioned on starting in the current state x_k.[4]

Suppose P is a C3 property. We can show that a value function v_k satisfies P using an induction argument of the form discussed in §2. First, because the set of functions satisfying any C3 property is a cone, $v_0(x_0) \equiv 0$ satisfies P. The fact that the set of functions satisfying P is a convex cone implies that $v_k(x_k) = r_k(x_k) + \delta_k\,\mathrm{E}[v_{k-1}(\tilde{x}_{k-1}(x_k))]$ satisfies P whenever both of the summands satisfy P. Thus, if $r_k(x_k)$ satisfies P for each k and we can somehow ensure that $\mathrm{E}[v_{k-1}(\tilde{x}_{k-1}(x_k))]$ satisfies P whenever $v_{k-1}(x_{k-1})$ satisfies P, then by induction we can conclude that v_k satisfies P for all k. Moreover, the fact that the set of functions satisfying property P is closed under pointwise limits implies that the limiting function, $\text{limit}_{k\to\infty} v_k$, if it exists, also satisfies P. We can appeal to Proposition 2 to provide conditions to ensure that the conditional expectations $\mathrm{E}[v_{k-1}(\tilde{x}_{k-1}(x_k))]$ will satisfy P.

PROPOSITION 3. *Let U be a set of functions on X satisfying a C3 property P. If for all k,*

(a) the reward functions $r_k(x_k)$ satisfy P and

(b) the transitions $\tilde{x}_{k-1}(x_k)$ satisfy $P(\precsim_U)$,

then each v_k satisfies P and $\text{limit}_{k\to\infty} v_k$, if it exists, also satisfies P.

Proof. This follows from the inductive argument given before the proposition and noting that by Proposition 2, $\mathrm{E}[f(\tilde{x}_{k-1}(x_k))]$ satisfies P for all f satisfying P if and only if the transitions $\tilde{x}_{k-1}(x_k)$ satisfy $P(\succsim_U)$. □

Some simple and abstract examples of this result follow.

1. *Increasing*. v_k is increasing if the reward functions are increasing and the transitions are stochastically increasing in the sense of first-order stochastic dominance.

2. *Increasing and Concave.* v_k is increasing and concave if the reward functions are increasing and concave and the transitions are stochastically increasing and concave in the sense of second-order monotonic stochastic dominance.

3. *Increasing and Convex.* v_k is increasing and convex if the reward functions are increasing and convex and the transitions are stochastically increasing and convex in the sense of first-order stochastic dominance.

While in the previous section we could distinguish between the set of functions U and the property P (for instance, in Example 2 in §4.2, U consisted of increasing functions and P was concave), here, because of the recursive structure of the problem, we must equate these two classes of functions.

Proposition 3 shows how C3 properties are preserved through the recursive calculations of the Markov reward process. We can't quite say that *only* C3 properties are preserved through these recursive calculations. Given the additive structure of the Markov reward process, if a property is to be preserved for arbitrary discount factors ($\delta_k \geqslant 0$), then the set of functions satisfying the property must form a convex cone. The closure of this set under pointwise limits is sufficient to ensure that the limiting functions, if they exist, will also be in that set. One could, however, use other arguments to prove that value functions satisfy properties that are not C3 properties. For example, if one assumes that the reward functions are bounded and continuous and the transitions satisfy the "Feller property," we can guarantee that the value functions will also be bounded and continuous.[5] In this case, the set of functions satisfying this property forms a convex cone but it is not closed in the topology of pointwise convergence.

6. CHARACTERIZING VALUE FUNCTIONS FOR MARKOV DECISION PROCESSES

Now we consider a Markov decision process of the form

$$v_k^*(x_k) \equiv \sup_{ak\in Ak} \{r_k(a_k, x_k) + \delta_k \mathrm{E}[v_{k-1}^*(\tilde{x}_{k-1}(a_k, x_k))]\} \quad \text{for } k > 0,$$

$$v_0^*(x_0) \equiv 0,$$

where the definitions are as in the previous section, except here we choose the action a_k from a set of feasible actions A_k in period k to maximize the value functions v_k^*. To simplify the analysis, we will assume that the set of feasible actions A_k does not vary by state. The analysis in this section is complicated by the fact that the rewards and transitions depend on the actions as well as the current state and we must consider what is required for the properties to be preserved through the maximization in the selection of the action. We consider the maximization problem in isolation before proceeding to the full stochastic dynamic programming problem.

6.1. Properties Preserved by Maximization

What properties must a function $f(a,\theta)$ need to satisfy to ensure that $g(\theta) \equiv \sup_{a\in A}\{f(a,\theta)\}$ satisfies P in θ? One might speculate that a sufficient condition for $g(\theta)$ to satisfy a property P in θ is for $f(a,\theta)$ to satisfy P in θ for each a. The intuition is that, if the property holds for every action a, we might expect it to hold for the optimal action a^*. Unfortunately, this intuition is faulty: The optimal action a^* depends on θ and it is possible for $f(a,\theta)$ to satisfy P in θ for each a and yet have $g(\theta)$ not satisfy P. For example, with $A = \Theta = \mathbb{R}^1$, $f(a,\theta) = -\exp(-a^2)((a-\theta)^2+1)$ is concave in θ for each a. The optimal action is to take $a=\theta$ and this gives $g(\theta) = -\exp(-\theta^2)$, which is not concave in θ.

This intuition does, however, hold for a special class of properties that we will call *single-point properties*. We consider this special case because it is easy to check in applications and because it illustrates what is required for a property to be preserved under maximization; we will consider a more general class of properties in a moment. We say a C3 property P is a *single-point property* if the inequalities in the representation of P in Proposition 1 involve only a single function evaluation on the left side; that is, g satisfies P if and only if

$$g(\theta_\alpha) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma g(\theta_\gamma) \quad \text{for all } \alpha \text{ in } \mathscr{A}.$$

The C3 properties increasing, decreasing, convex, and subadditive are all single-point properties. Intersections of single-point properties are also single-point properties. Concavity is an example of a C3 property that is not a single-point property: It can be represented as an inequality involving a single function evaluation on one side of the inequality, but the single function evaluation is on the wrong side of the inequality.

It is easy to show that if P is a single-point property on Θ and $f(a,\theta)$ satisfies P for each $a \in A$, then $g(\theta) \equiv \sup_{a\in A}\{f(a,\theta)\}$ also satisfies P. If $f(a,\theta)$ satisfies P for each $a \in A$, then for any $\alpha \in \mathscr{A}$ and action $a \in A$,

$$f(a,\theta_\alpha) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma f(a,\theta_\gamma) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma g(\theta_\gamma).$$

The first inequality holds because f satisfies P for each $a \in A$; the second because $f(a,\theta) \leqslant g(\theta)$ for all a and θ. Taking the supremum over actions $a \in A$ on the left side of this inequality (i.e., choosing the action a to be optimal for θ_α), we find $g(\theta_\alpha) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma g(\theta_\gamma)$. Because this holds for each α in $\mathscr{A}$, g satisfies P.

To develop a more general condition that ensures that $g(\theta) \equiv \sup_{a\in A}\{f(a,\theta)\}$ satisfies a C3 property P, consider the inequality representation of this property: g satisfies P if and only if

$$\sum_{\beta\in B_\alpha} \lambda_\beta\, g(\theta_\beta) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma\, g(\theta_\gamma) \quad \text{for all } \alpha \text{ in } \mathscr{A}.$$

Now fix α in $\mathscr{A}$ and consider the following series of equalities and inequalities:

$$\sum_{\beta\in B_\alpha} \lambda_\beta\, g(\theta_\beta) = \sum_{\beta\in B_\alpha} \lambda_\beta\, f(a^*(\theta_\beta), \theta_\beta)$$
$$\leqslant \sum_{\beta\in B_\alpha} \lambda_\gamma\, f(a_\gamma, \theta_\gamma) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma\, g(\theta_\gamma). \tag{4}$$

The outer elements correspond to g satisfying this equality (α) of property P. In the second term on the left, we let $a^*(\theta_\beta)$ denote the maximizing action a for a given value θ_β, assuming for expository purposes that this maximum is obtained by some action. (This assumption is not required and will be dropped later.) The equality on the left of (4) then follows from this definition. The inequality on the right follows the definition of g, since $f(a, \theta) \leqslant g(\theta)$ for all a and θ. The key to proving that g satisfies P, then, is showing that f satisfies the middle inequality for some set of actions $\{a_\gamma\}$.

We can ensure that (4) holds by assuming that f satisfies a *joint extension* of P defined as follows. Given a C3 property P on Θ represented as in Equation (3), we say that a C3 property P^* on $A \times \Theta$ is a *joint extension* of P if for any α in $\mathscr{A}$ in Equation (3) and any set of actions $\{a_\beta\}_{\beta\in B_\alpha}$, there exists a set of actions $\{a_\gamma\}_{\gamma\in\Gamma_\alpha}$ such that

$$\sum_{\beta\in B_\alpha} \lambda_\beta\, f(a_\beta, \theta_\beta) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma\, f(a_\gamma, \theta_\gamma). \tag{5}$$

This condition ensures that the middle inequality in Equation (4) is satisfied for some set of actions $\{a_\gamma\}_{\gamma\in\Gamma_\alpha}$ and is sufficient to ensure that $g(\theta) \equiv \sup_{a\in A}\{f(a, \theta)\}$ satisfies P.

PROPOSITION 4. *Let P be a C3 property on Θ and let P^* be a joint extension of P on $A \times \Theta$. If $f(a, \theta)$ satisfies P^*, then $g(\theta) \equiv \sup_{a\in A}\{f(a, \theta)\}$ satisfies P.*

PROOF. By definition of a joint extension, for any constraint α in the inequality representation of P, given a set of actions $\{a_\beta\}_{\beta\in B_\alpha}$ for the left side of Equation (5) and the corresponding set of actions $\{a_\gamma\}_{\gamma\in\Gamma_\alpha}$ for the right, we have

$$\sum_{\beta\in B_\alpha} \lambda_\beta\, f(a_\beta, \theta_\beta) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma\, f(a_\gamma, \theta_\gamma) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma\, g(\theta_\gamma).$$

The first inequality follows from the assumption that f satisfies an extended version of P and the second follows from the definition of g. Taking the supremum inside the summation on the left side of the inequality (i.e., choosing the a_β to be optimal for the corresponding θ_β), we have

$$\sum_{\beta\in B_\alpha} \lambda_\beta\, g(\theta_\beta) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma\, g(\theta_\gamma).$$

Thus, g satisfies this inequality and, more generally, satisfies property P. □

Some examples may help clarify the nature of this condition. The first three are positive examples; the fourth is a negative example in which the condition does not hold.

1. *Single-Point Properties.* When we considered single-point properties, we assumed that the property was satisfied for all choices of actions, i.e.,

$$f(a, \theta_\alpha) \leqslant \sum_{\gamma\in\Gamma_\alpha} \lambda_\gamma\, f(a, \theta_\gamma) \quad \text{for all } a \text{ in } A \text{ and } \alpha \text{ in } \mathscr{A}.$$

In this case, f satisfies a joint extension of P on $A \times \Theta$ because the inequalities are satisfied whenever the actions on the right side of the inequality match the action chosen on the left.

2. *Joint Concavity.* Given convex parameter and action spaces (Θ and A), joint concavity requires that for any choice of θ_1 and θ_2 in Θ and a_1 and a_2 in A and ϕ such that $0 \leqslant \phi \leqslant 1$,

$$\phi f(a_1, \theta_1) + (1-\phi) f(a_2, \theta_2) \leqslant f(a_\phi, \theta_\phi)$$

for $a_\phi \equiv \phi a_1 + (1-\phi)a_2$ and $\theta_\phi \equiv \phi\theta_1 + (1-\phi)\theta_2$. In this case, for any given θ_1, θ_2, and ϕ, we are free to choose actions a_1 and a_2 on the left side of the inequality, and there exists an action on the right side of the inequality (a_ϕ) such that this inequality is satisfied. Joint concavity on $A \times \Theta$ is thus a joint extension of concavity on Θ. (Note that it is not enough for f to be concave in θ for each a, as illustrated by the example at the beginning of this section.)

3. *Joint Supermodularity.* Given parameter and action spaces (Θ and A) that are lattices, joint supermodularity requires that, for all θ_1 and θ_2 in Θ and a_1 and a_2 in A,

$$f(a_1, \theta_1) + f(a_2, \theta_2) \leqslant f(a_1 \wedge a_2, \theta_1 \wedge \theta_2) + f(a_1 \vee a_2, \theta_1 \vee \theta_2).$$

Like the previous example, for any given θ_1 and θ_2 for any choice of actions a_1 and a_2 on the left side of the inequality, there will be actions for the right side of the inequality ($a_1 \wedge a_2$ and $a_1 \vee a_2$) such that the inequality is satisfied. Joint supermodularity on $A \times \Theta$ is thus a joint extension of supermodularity on Θ.

4. *Joint Submodularity.* Given parameter and action spaces (Θ and A) that are lattices, joint submodularity requires that, for all θ_1 and θ_2 in Θ and a_1 and a_2 in A,

$$f(a_1 \wedge a_2, \theta_1 \wedge \theta_2) + f(a_1 \vee a_2, \theta_1 \vee \theta_2) \leqslant f(a_1, \theta_1) + f(a_2, \theta_2).$$

This does not satisfy the joint extension condition because we cannot independently select actions $(a_1 \wedge a_2)$ and $(a_1 \vee a_2)$ on the left side of this inequality and be sure that there will be actions on the right such that this inequality holds, because $(a_1 \wedge a_2) \leqslant (a_1 \vee a_2)$ for any a_1 and a_2.

The joint extension condition of Proposition 4 is sufficient but not necessary for g to inherit the C3 property P.

The middle inequality in Equation (4) need only hold for the actions $\{a_\beta\}_{\beta\in B_\alpha}$ that are optimal for the corresponding points $\{\theta_\beta\}_{\beta\in B_\alpha}$. We could therefore relax this condition by requiring Equation (5) to hold only for the actions that are optimal for the left side.

6.2. Markov Decision Processes

To apply these results to Markov decision processes, we take the function to be maximized to be $f(a_k, x_k) = r_k(a_k, x_k) + \delta_k \,\mathrm{E}[v^*_{k-1}(\tilde{x}_{k-1}(a_k, x_k))]$ and look for conditions that ensure that this function satisfy a joint extension of the property P. We can use essentially the same argument as in the last section to achieve this.

Proposition 5. *Let U be a set of functions on X satisfying a C3 property P and let P^* be a joint extension of P on $A\times\Theta$. If, for all k,*

(a) the reward functions $r_k(a_k, x_k)$ satisfy P^* and

(b) the transitions $\tilde{x}_{k-1}(a_k, x_k)$ satisfy $P^*(\succsim_U)$,

then each v^*_k satisfies P and $\mathrm{limit}_{k\to\infty} v^*_k$, if it exists, also satisfies P.

Proof. We proceed by induction. First, because the set of functions satisfying any C3 property is a cone, $v^*_0(x_0) = 0$ satisfies P. Now assume that $v^*_{k-1}(x_{k-1})$ satisfies P. The fact that the set of functions satisfying P^* is a convex cone implies that $r_k(a_k, x_k) + \delta_k \,\mathrm{E}[v^*_{k-1}(\tilde{x}_{k-1}(a_k, x_k))]$ satisfies P^* whenever both of the summands satisfy P^*; if this satisfies P^*, then $v^*_k(x_k)$ satisfies P by Proposition 4. The first summand $r_k(a_k, x_k)$ satisfies P^* by the first assumption and the second $\mathrm{E}[v^*_{k-1}(\tilde{x}_{k-1}(a_k, x_k))]$ satisfies P^* for v^*_{k-1} satisfying P by the second assumption and Proposition 2. Thus, for each k, v^*_k satisfies P. As before, the fact the set of functions satisfying property P is closed under pointwise limits implies that the limiting function, $\mathrm{limit}_{k\to\infty} v^*_k$, if it exists, also satisfies P. □

Thus, to show that a value function satisfies some property P, we need to identify a joint extension P^* of P and check that the rewards and transitions satisfy P^*. We first consider a few simple and abstract examples of this result and then consider the example applications introduced in §2.

1. *Increasing.* Because increasing is a single-point property, $v^*_k(x_k)$ is increasing in x_k if the reward functions $r_k(a_k, x_k)$ are increasing in x_k for each a_k and the transitions are stochastically increasing in x_k in the sense of first-order stochastic dominance for each a_k.

2. *Increasing and Convex.* Because increasing and convex are both single-point properties, $v^*_k(x_k)$ is increasing and convex in x_k if the reward functions $r_k(a_k, x_k)$ are increasing and convex in x_k for all a_k and the transitions are stochastically increasing and convex in x_k in the sense of first-order stochastic dominance for each a_k.

3. *Increasing and Concave.* $v^*_k(x_k)$ is increasing and concave in x_k if the reward functions $r_k(a_k, x_k)$ are increasing in x_k for all a_k and jointly concave in a_k and x_k (joint concavity on $A\times X$ being a joint extension of concavity on X) and the transitions are increasing in x_k and concave in a_k and x_k in the sense of second-order monotonic stochastic dominance.

6.3. Example Applications

The example applications involve variations of the previous abstract examples and other properties. The examples and the corresponding properties are summarized in Table 1.

Copper Mine. Here we show that the value functions are increasing and convex functions of the log price p_k for each mine state s_k (open, closed, or abandoned); call this property P. The joint extension of this property (P^*) requires functions to be increasing and convex in p_k for each mine state s_k (*open, closed*, or *abandoned*) and each action a_k (making the mine *open, closed*, or *abandoned*); because increasing and convex are both single-point properties, P^* is indeed a joint extension of P. To check Condition (a) of Proposition 5, we note that the reward functions are increasing and convex in p_k for each mine state s_k and action a_k, and thus satisfy P^*. We saw in §4.3 that the price transitions $\tilde{p}_{k-1}(p_k)$ are stochastically increasing and convex in p_k on the set of functions that are increasing and convex functions in p_k. Because this holds for each mine state s_k and action a_k, the full transitions $\tilde{x}_{k-1}(a_k, s_k, p_k)$ satisfy P^* as well. By Proposition 5, we can conclude that value functions $v^*_k(p_k, s_k)$ are increasing and convex in p_k as well.

We can also show that the value functions are increasing in σ. Again this is a single-point property. The reward functions are constant (and therefore increasing) in the volatility (σ) of prices and for each a_k, s_k, and p_k and, as shown in §4, the price transitions $\tilde{p}_{k-1}(p_k)$ are stochastically increasing in σ on the set of functions that are increasing and convex in p_k. The result follows from Proposition 5. One can similarly also show that the value functions are increasing and convex in the revenue margin (γ) as well as other properties.

Technology Adoption. Here we show that the value functions $v^*_k(x_k, s_k, t_k)$ are (i) increasing and convex functions of the current estimate of the value of the technology x_k, (ii) decreasing in the precision s_k associated with this estimate, (iii) increasing in the precision t_k associated the observation in each period, and (iv) satisfy the mixing property described in §2; this combination of properties is the property P that we want to show the value function satisfies. Each of these properties is a single-point property and therefore P is as well. The joint extension P^* of P requires P to hold for each action a_k. The reward functions satisfy P^* for each choice of action (this is easy to check) and thus satisfy P^*. We showed in §4 and the appendix that the transitions also satisfy P for each action a_k and thus satisfy P^*. Proposition 5 then implies that the value functions satisfy property P.

Stochastic Inventory Model. We show that the cost functions $v_k^*(x_k)$ are convex in the n-vector of product inventory levels x_k; this is the desired property P. Because joint concavity on a_k and x_k is a joint extension of concavity on x_k for a maximization by problem (as in the example in §6.1), joint convexity of a function $f(a_k, x_k)$ in a_k and x_k will ensure that $g(x_k) = \inf_{a\in A} f(a_k, x_k)$ is convex in x_k; joint convexity in a_k and x_k is thus the property P^*.

To verify the first condition of Proposition 5, we show that the reward function for the model $r_k(a_k, x_k) = c(a_k) + \mathrm{E}[l(\tilde{x}_{k-1}(a_k, x_k))]$ is jointly convex in a_k and x_k for each k. The costs of production $c(a_k)$ were assumed to be convex in the amount ordered (a_k). Because the penalty function $l(x_{k-1})$, reflecting shortage and inventory costs, is assumed to be convex and the transitions $\tilde{x}_{k-1}(a_k, x_k)$ are stochastically convex in (a_k, x_k) for convex functions (as shown in §4.3), by Proposition 2, the expected penalty $\mathrm{E}[l(\tilde{x}_{k-1}(a_k, x_k))]$ is jointly convex in a_k and x_k. The reward function, as the sum of two convex functions, is thus jointly convex in a_k and x_k. The second condition of Proposition 5, that the transitions $\tilde{x}_{k-1}(a_k, x_k)$ be stochastically convex for convex functions, was established in §4.3. Proposition 5 then implies that the cost functions are convex in the n-vector of inventory level x_k, as sought.

7. SUMMARY

We have developed a set of metatheorems that describe how properties of value functions are preserved and propagated through Markov reward and decision processes. The metatheorems provide specific conditions that can be checked in applications to establish specific properties of the value functions. Perhaps more importantly, the theorems clarify the structure of these kinds of results. The main results can be summarized intuitively as saying that the value functions satisfy property P if the reward functions satisfy property P and the transition probabilities satisfy a stochastic version of this property. With Markov reward processes, this result holds quite generally for closed convex cone (C3) properties. With Markov decision processes, we need to consider C3 properties that are closed under maximization (or minimization) in the choice of actions; we provide general conditions that ensure this is the case.

APPENDIX

PROOF OF PROPOSITION 1. Let E denote the set of real-valued functions defined on Θ, endowed with the topology of pointwise convergence, also called the product topology for $\mathbb{R}^\Theta$. As the product of locally convex topologies, the product topology $\mathbb{R}^\Theta$ is also locally convex. E is thus a locally convex topological vector space. Let $\Pi \subseteq E$ denote the closed convex cone of functions satisfying a C3 property P. As a nonempty closed convex set in E, Π can be represented as the intersection of all the closed half-spaces that contain Π (see, e.g., Berberian 1974, Theorem 34.3). Each of these closed half-spaces can be represented as an inequality involving a continuous linear functional F and scalar a such that $F(f) \geqslant a$ for all $f \in \Pi$. Thus, if we let $\mathcal{A}$ be an index set describing the collection of closed half-spaces containing Π, the set Π can be represented as the set of solutions to a set of linear inequalities of the form $F_\alpha(f) \geqslant a_\alpha$ for all $\alpha \in \mathcal{A}$.

We now show that because Π is a cone, we can restrict ourselves to inequalities of the form $F_\alpha(f) \geqslant a_\alpha = 0$ for each $\alpha \in \mathcal{A}$. Consider any α. If $a_\alpha > 0$, then for any $f \in \Pi$ we have $F_\alpha(f) \geqslant a_\alpha$. Yet by linearity of F_α, for a sufficiently small positive β we would have $F_\alpha(\beta f) = \beta F_\alpha(f) < a_\alpha$. This is a contradiction: Because Π is a cone, $\beta f \in \Pi$ and thus $F_\alpha(\beta f) \geqslant a_\alpha$ whenever $\beta \geqslant 0$ and $f \in \Pi$. Thus, $a_\alpha \leqslant 0$. Similarly, if $a_\alpha < 0$ and there exists an $f \in \Pi$ such that $0 > F_\alpha(f) \geqslant a_\alpha$, then for a sufficiently large $\beta > 0$, we would have $F_\alpha(\beta f) = \beta F_\alpha(f) < a_\alpha$. Again, this is a contradiction because $\beta f \in \Pi$ whenever $f \in \Pi$. Finally, if $a_\alpha < 0$ and there is no $f \in \Pi$ such that $0 > F_\alpha(f) \geqslant a_\alpha$, then we can replace a_α with 0 and still have $F_\alpha(f) \geqslant 0$ for all $f \in \Pi$. Thus the set Π can be represented as the set of solutions to a set of linear inequalities of the form $F_\alpha(f) \geqslant 0$ for all α in $\mathcal{A}$.

We now show that each linear functional $F_\alpha(f)$ depends on the value of f at a finite set of points $\theta \in \Theta$. Pick an $\varepsilon > 0$ and consider the inverse image of the set $(-\varepsilon, \varepsilon)$ under the functional F_α, $F_\alpha^{-1}[(-\varepsilon, \varepsilon)]$; this is the set of functions whose evaluations under F_α are within ε of 0. Because F_α is continuous, $F_\alpha^{-1}[(-\varepsilon, \varepsilon)]$ is an open set. Moreover, $F_\alpha^{-1}[(-\varepsilon, \varepsilon)]$ is not empty because for the zero function $f_o(\theta) \equiv 0$ for all θ in Θ, we have $F(f_o) = 0$. Considering the base for the product topology $\mathbb{R}^\Theta$ (see, e.g., Royden 1968, pg. 152), there exists a nonempty, open "basic set" B such that $f_o \in B \subseteq F_\alpha^{-1}[(-\varepsilon, \varepsilon)]$ and B is of the form $\{f : f(\theta_1) \in O_1, f(\theta_2) \in O_2, \ldots, f(\theta_n) \in O_n\}$ where $\{\theta_1, \theta_2, \ldots, \theta_n\}$ is some finite subset of Θ and the sets $\{O_1, O_2, \ldots, O_n\}$ are open sets in $\mathbb{R}$ containing 0. Thus, membership in B depends only on the value of the function at the finite set of points $\{\theta_1, \theta_2, \ldots, \theta_n\}$. This then implies that the functional $F_\alpha(f)$ depends only on the values of the function f at these points. To prove this, suppose that $F_\alpha(f)$ depends on the values of f outside of these points. Then there must be two functions f_1 and f_2 such that $f_1(\theta_i)$ and $f_2(\theta_i)$ for $i = 1, \ldots, n$ but such that $F_\alpha(f_1) \neq F_\alpha(f_2)$. By linearity, this implies that $F_\alpha(f_\nabla) \neq 0$ for $f_\nabla \equiv f_1 - f_2$. Because $F_\alpha(f_\nabla) \neq 0$, for sufficiently large β, $F_\alpha(\beta f_\nabla) = \beta F_\alpha(f_\nabla)$ will be outside of $(-\varepsilon, \varepsilon)$. Because $f_\nabla(\theta_i) = 0$ for $i = 1, \ldots, n$, $f_\nabla \in B$ and, moreover, for any β, $\beta f_\nabla \in B$. Because $\beta f_\nabla \in B$ and $B \subseteq F_\alpha^{-1}[(-\varepsilon, \varepsilon)]$, this implies $F_\alpha(\beta f_\nabla) \in (-\varepsilon, \varepsilon)$, contradicting our earlier finding that $F_\alpha(\beta f_\nabla) \notin (-\varepsilon, \varepsilon)$.

Given that the linear functional $F_\alpha(f)$ depends on values of f at only these points, it can be represented as a weighted sum $F_\alpha(f) = \sum_{i=1}^n \omega_{\alpha i} f(\theta_{\alpha i})$ for some set of weights $\{\omega_{\alpha 1}, \omega_{\alpha 2}, \ldots, \omega_{\alpha n}\}$. By separating the sets of points and weights into groups according to whether the weights are positive or negative, we arrive at the representation of the proposition. (We can drop any zero weights;

the index sets B_α and Γ_α correspond to the negative and positive weights, respectively.)

If the property P is a C5 property as well as a C3 property, because **1** and **−1** both satisfy P, then for each α, the sum of the weights on both sides of the inequality in the proposition must be equal, and hence can be normalized to sum to one. □

PROOF OF PROPERTIES FOR THE TECHNOLOGY ADOPTION MODEL. Here we show that the transitions satisfy a stochastic version of the mixing property and are increasing in t for functions u that satisfy the mixing property. To show that the transitions satisfy this mixing property, we must show that $\mathrm{E}[u(\tilde{m}_{k-1}(m_k, s_k), s_{k-1}(s_k))]$ satisfies:

$$0.5\mathrm{E}[u(\tilde{m}_{k-1}(m_k+\delta, s_k+\Delta), s_{k-1}(s_k+\Delta))] + 0.5\mathrm{E}[u(\tilde{m}_{k-1}(m-\delta, s_k+\Delta), s_{k-1}(s_k+\Delta))] \geqslant \mathrm{E}[u(\tilde{m}_{k-1}(m_k, s_k), s_{k-1}(s_k))]$$

for any $\Delta, s_k \geqslant 0$ and $\delta^2 = \Delta/(s(s+\Delta))$, for any function u satisfying the mixing property. Using the definition of the transitions, this condition is equivalent to

$$0.5\mathrm{E}[u(m_k+\tilde{y}+\delta, s_k+t+\Delta)] + 0.5\mathrm{E}[u(m_k+\tilde{y}-\delta, s_k+t+\Delta)] \geqslant \mathrm{E}[u(m_k+\tilde{y}, s_k+t)],$$

where $\tilde{y}$ is normally distributed with mean 0 and precision $s_k(s_k+t)/t$. If the function u satisfies the mixing property, for any y we have

$$0.5u(m_k+y+\delta, s_k+t+\Delta) + 0.5u(m_k+y-\delta, s_k+t+\Delta) \geqslant u(m_k+y, s_k+t).$$

Because this holds for all y, this implies that the previous inequality holds and the transitions satisfy the mixing property for functions u satisfying the mixing property.

To show that the transitions are increasing in t for functions satisfying the mixing property, we first show that the mixing property can be extended from binary to normal mixtures. That is, for any m, and s, $\Delta \geqslant 0$, and with $\tilde{z}$ a normally distributed random variable with mean 0 and variance equal to $\Delta/(s(s+\Delta))$, we have $\mathrm{E}[u(m+\tilde{z}, s+\Delta)] \geqslant u(m, s)$, for any u that satisfies the mixing property. The interpretation is the same as before: The expected value added by resolving the uncertainty $\tilde{z}$ increases the expected value more than the value lost due to increasing the precision of the revised estimate by Δ. To prove this, we will construct a series of binary random variables whose sums approximate the normally distributed random variable $\tilde{z}$ (similar to a "binomial tree" approximation used to value options on stocks; see Cox et al. 1979) where we can use the binary mixing property at each stage in the summation. The binary random variables $\{\tilde{z}_j\}$, $j=1$ to n, are assumed to independent random variables and equally likely to take on values $\pm\delta_j$ where $\delta_j^2 = 1/(s+\Delta(j-1)/n) - 1/(s+\Delta j/n)$. Let $\tilde{w}_i$ denote the partial sum $\sum_{j=1}^{i} \tilde{z}_j$. The variance of $\tilde{w}_i$ is then $\sum_{j=1}^{i} \delta_j^2 = (i/n)\Delta/(s(s+(i/n)\Delta))$. The variance of $\tilde{w}_n$ is thus $\Delta/(s(s+\Delta))$, the variance of $\tilde{z}$. In the limit as $n \to \infty$, $\tilde{w}_n$ converges in distribution to $\tilde{z}$; convergence follows from the Lindeberg Theorem (Billingsley 1986, p. 369), a variation of the Central Limit Theorem. The binary mixing property then implies that, for $i = 1, \ldots, n$, $\mathrm{E}[u(m+\tilde{w}_i, s+\Delta i/n)] \geqslant \mathrm{E}[u(m+\tilde{w}_{i-1}, s+\Delta(i-1)/n)]$. Here each stage i corresponds to the addition of the binary random variable $\tilde{z}_i$ with a corresponding increase in precision of Δ/n. Chaining these inequalities together we have $\mathrm{E}[u(m+\tilde{w}_n, s+\Delta)] \geqslant u(m, s)$ for each n. In the limit as $n \to \infty$, w_n converges in distribution to $\tilde{z}$, so $\mathrm{E}[u(m+\tilde{w}_n, s+\Delta)]$ approaches $\mathrm{E}[u(m+\tilde{z}, s+\Delta)]$ and therefore $\mathrm{E}[u(m+\tilde{z}, s+\Delta)] \geqslant u(y, s)$.

We now show that the transitions are increasing in t for functions that satisfy the mixing property. Consider the process starting in state (m_k, s_k). If the observation is made with precision t, $\tilde{m}_{k-1}$ will be normally distributed with mean m_k and variance $t/s_k(s_k+t)$; let $\tilde{y}$ be a random variable with this distribution. If the observation is made with precision $t+\Delta$, then $\tilde{m}_{k-1}$ will be normally distributed with mean m_k and variance $(t+\Delta)/(s_k(s_k+t+\Delta))$. Let $\tilde{z}$ be a random variable that is normally distributed with mean 0 and variance $(t+\Delta)/(s_k(s_k+t+\Delta)) - t/(s_k(s_k+t))$ and independent of $\tilde{y}$. The sum $\tilde{y}+\tilde{z}$ will have the same distribution as $\tilde{m}_{k-1}$ if the observations are made with precision $t+\Delta$. The transitions will be increasing in t for functions u satisfying the mixing property if

$$\mathrm{E}[u(\tilde{y}+\tilde{z}, s_k+t+\Delta)] \geqslant \mathrm{E}[u(\tilde{y}, s_k+t)] \tag{A1}$$

holds for these functions. For any value y, the mixing property for normal distributions implies that $\mathrm{E}[u(y+\tilde{z}, s_k+t+\Delta)] \geqslant u(y, s_k+t)$. The fact that this holds for any y implies that it holds when taking expectations over $\tilde{y}$ as in (A1). Thus the transitions are increasing in t for functions satisfying the mixing property. □

ENDNOTES

1. In Table 1, we write the reward functions and transitions as function of the state variables (x_k) using notation specific to the examples. Where we want to study sensitivity to a parameter that is not a state variable, we include these exogenous parameters as arguments to the relevant functions following the state variables and separated by a semicolon. For example, in the copper mine model, we are interested in the sensitivity of the value functions to the volatility of prices (σ), which is not a state variable in the model.

2. Formally, let $(X, \Im)$ be a measurable space and let $\mu(\theta)$ be a family of measures defined on this space; the functions u are assumed to be measurable with respect to $(X, \Im)$. The use of the random state notation simplifies some of our later discussions because we can consider properties

of the random function $\tilde{x}(\theta)$ in the same way we consider properties of nonrandom functions.

3. If P is a C5 property, the weights in (3) sum to one and consequently, if each of the measures $\mu(\theta)$ is a probability measure, the weighted measures involved in the dominance comparison in (3)—referred to as μ_α^1 and μ_α^2 in the proof—are also probability measures. If P is a C3 property but not a C5 property, μ_α^1 and μ_α^2 are not necessarily probability measures. The same definition of dominance applies however—$\mu_\alpha^1 \precsim_U \mu_\alpha^2$ if and only if $\int u\, d\mu_\alpha^1 \leqslant \int u\, d\mu_\alpha^2$ for all u in U—and the result of Proposition 2 holds.

4. We assume that the state variables are all random variables defined on some measurable space $(X, \Im)$ and that the reward functions are all measurable with respect to this space.

5. A transition function $\tilde{x}_{k-1}(x_k)$ satisfies the "Feller property" if $\mathrm{E}[u(\tilde{x}_{k-1}(x_k))]$ is a bounded and continuous function of x_k for functions u that are bounded and continuous in x_{k-1}; see Stokey and Lucas (1989). Stokey and Lucas also provide equivalent probabilistic conditions (p. 376).

ACKNOWLEDGMENTS

The authors are grateful for the research support provided by Chevron and the National Science Foundation under the Joint NSF/Private Sector Research Opportunities Initiative (grant number SBR-9809176). They are also grateful for the helpful comments provided by an anonymous associate editor.

REFERENCES

Athey, S. 1998. Characterizing properties of stochastic objective functions. Working paper, MIT, Cambridge, MA.

Berberian, S. K. 1974. *Lectures in Functional Analysis and Operator Theory*. Springer-Verlag, New York.

Billingsley, P. 1986. *Probability and Measure*, 2nd ed. John Wiley and Sons, New York.

Brennan, M. J., E. S. Schwartz. 1985. Evaluating natural resource investments. *J. Bus.* **58** 135–157.

Cox, J. C., S. A. Ross, M. Rubinstein. 1979. Option pricing: A simplified approach. *J. Financial Econom.* **7** 229–263.

Derman, C. 1963. On optimal replacement rules when changes of state are Markovian. R. Bellman, ed. *Mathematical Optimization Techniques*. University of California Press, Berkeley, CA.

Dixit, A. K., R. S. Pindyck. 1994. *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ.

Hinderer, K. 1985. On the structure of solutions of stochastic dynamic programs. *Proc 7th Conference on Probability Theory, Brasov 1984*. M. Iosifescu, ed. VNU Science Press, Utrecht, The Netherlands, 173–182.

Karlin S. 1960. Dynamic inventory policy with stochastic demands. *Management Sci.* **6** 231–258.

Lippman, S. A., K. F. McCardle. 1987. Does cheaper, faster, or better imply sooner in the timing of innovation decisions? *Management Sci.* **33** 1058–1064.

McCardle, K. F. 1985. Information acquisition and the adoption of new technology. *Management Sci.* **31** 1372–1389.

Merton, R. C. 1973. The theory of rational options pricing. *Bell J. Econom. Management Sci.* **4** 141–183.

Milgrom, P., C. Shannon. 1994. Monotone comparative statics. *Econometrica* **62** 157–180.

Müller, A. 1997. How does the value function of a Markov decision process depend on the transition probabilities? *Math. Oper. Res.* **22** 872–885.

Pratt, J. W., H. Raiffa, R. Schlaifer. 1995. *Introduction to Statistical Decision Theory*. MIT Press, Cambridge, MA.

Ross, S. M. 1983. *Introduction to Stochastic Dynamic Programs*. Academic Press, San Diego, CA.

Rothschild, M., J. E. Stiglitz. 1970. Increasing risk I: A definition. *J. Econom. Theory* **2** 225–243.

Royden, H. L. 1968. *Real Analysis*, 2nd ed. Macmillan Publishing Co., Inc., New York.

Shaked, M., J. G. Shantikumar. 1994. *Stochastic Orders and Their Applications.* Academic Press, San Diego, CA.

Stokey, N., R. E. Lucas Jr., E. C. Prescott. 1989. *Recursive Methods in Economic Dynamics.* Harvard University Press, Cambridge, MA.

Topkis, D. M. 1998. *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.

Zipkin, P. H. 2000. *Foundations of Inventory Management*. McGraw-Hill, Boston, MA.