

What Do Discounted Optima Converge To? A Theory of Discount Rate
Asymptotics in Economic Models

Dutta, Prajit

Working Paper No. 264
January 1991

University of
Rochester

**What Do Discounted Optima Coverage To?
A Theory of Discount Rate Asymptotics in Economic Models**

Prajit Dutta

Rochester Center for Economic Research
Working Paper No. 264

**What Do Discounted Optima Converge To?
A Theory of Discount Rate Asymptotics in Economic Models***

Prajit K. Dutta*
Columbia University

Rochester Center for Economic Research
Working Paper No. 264

January 1991

* I am indebted to Duncan Foley for detailed and invaluable comments on an earlier draft and to Ehud Lehrer for a number of substantive suggestions which greatly improved the quality of the paper. I have benefitted greatly from conversations with M. Ali Khan, Kerry Back, Andrew Caplin, Phil Dybvig, Mukul Majumdar, Andreu Mas-Colell, Aldo Rustichini and Jorgen Weibull, as also the comments of participants at the Columbia University Theory Workshop, the Econometric Society meetings at Munich and Atlanta and the Midwestern Mathematical Economics Meetings at St. Louis. I thank an anonymous referee and an associate editor for insightful comments. I owe a particular debt to Raghu Sundaram for being a patient and instructive sounding board. Unfortunately, I alone am responsible for the remaining errors. Research support from the Social Sciences Research Council at Columbia University is gratefully acknowledged. This version was completed while visiting the University of Rochester.

Abstract : This paper considers a sequence of stochastic dynamic problems in which the discount rate goes to zero. Sufficient conditions are given under which the limit of discounted optimal policies (and normalized values) are a i) long-run average optimal and ii) catching-up optimal policy (and value). These conditions are shown to be satisfied in many economic models. A new decision criterion for the undiscounted problem, the strong long-run average, is introduced. This criterion refines the long-run average, is often equivalent to the catching-up and is easier to use than the latter since it has a recursive representation. A generalization of Fatou's lemma, which has wider applicability, is also proved. *Journal of Economic Literature* Classification Numbers 022, 111.

1. Introduction

The general objective that motivates this research is to develop a theory of parametric variations in dynamic economic models. In particular we would like to examine the continuity of optimal choices and value functions as underlying economic parameters vary. The parameter that we focus on in this paper is the discount factor. A first specific objective of this work will be to examine the question: under what conditions are discounted and undiscounted optimization problems special cases of a unified family of optimization problems? A number of decision criteria have been used for the undiscounted model and a second objective of this paper is to explore the links between them. Specifically we introduce a new criterion, the strong long-run average, which combines attractive features of the two most popular existing criteria, the long-run average and the catching-up, while avoiding their well-known shortcomings.

Variational analyses have been extensively pursued in static economic models and almost not at all in dynamic problems. Part of the reason may be that some of the standard techniques that work for variational problems in static models cannot be used in the dynamic context and one aim of this paper will be to develop alternative techniques. We pick the discount factor as our parameter of choice since the dynamic optimization literature in economics has been largely fragmented, with analyses that exclusively look either at the discounted or at the undiscounted model. Tools and techniques that are useful in one case are thought not to carry over to the other. This paper is an attempt to build a more unified theory. In particular, we attempt to isolate conditions in economic models under which the limit of solutions to the discounted problems are solutions to the undiscounted problem. Under such conditions then, there is a precise sense in which one problem is an approximation to the other. The paradigm within which we carry out this analysis is that of stochastic dynamic programming. The generality of the structure is of course extremely crucial since it accommodates most dynamic economic models.

There are several other reasons why one may be interested in exploring the undiscounted model and the relationship between discounted and undiscounted problems. Firstly, there is no a priori lower bound on discount rates. Time preference, which is often cited as an "explanation" for discounting, clearly admits no logical lower bound. Nor do alternative explanations like the uncertainty of lifetimes. Theoretical consistency requires of us therefore that we admit the possibility of zero discounting. Hence, one needs to analyze all positive discount rates as well as their limit, the undiscounted problem. Secondly, even if the "true" model is a discounted one, it is sometimes more convenient to analyze the

undiscounted problem. In many economic models values and policies under an undiscounted criterion like the long-run average are easy to compute. Whether or not the more tractable problem yields solutions which are good approximations to the true model of low but positive discount rates is a question that needs analysis.¹ Finally, economists like Pigou [26] and philosophers like Rawls [30, pp 293-298] argue against a positive discount rate for social planning problems since it reflects a bias against future generations. Unfortunately, existence of optima under undiscounted criterion like the catching up are often difficult to establish. Our analysis suggests an approach to this problem by treating it as a limit of discounted models, in which existence of optima is assured typically.

There are clearly two senses in which one can discuss continuity in the limit: continuity with respect to optimal policies and continuity with respect to optimal values. The first set of questions addressed in this paper are: what optimality properties, if any, does a limit of discounted optimal policies possess? When is a policy limit optimal under the long-run average criterion, and when is it optimal under the catching-up criterion?² When is there a limit to the maximized returns, i.e. the (normalized) values, and what can be said of this limit? Clearly the following questions are addressed as corollaries. Is an optimization exercise with "low" discount rates a good approximation to an undiscounted one? Conversely, is the dynamical behavior of a system under, e.g. a catching-up optimal policy, "close" to that governed by a discounted optima? Is a long-run average optimal policy e.g., approximately optimal under discounting?

Answers to these convergence questions are not immediate, and may not always be in the affirmative, for several reasons. The primary reason is that the objective, viewed as a function of the discount rate, may be discontinuous at zero. Hence, straightforward arguments like the Berge maximum theorem cannot be used. More subtle arguments will be employed to show that, despite such discontinuities, the Bellman operator is continuous in the discount rate in many instances. Easy counter-examples to the convergence questions may be found in [34] or [15]³. The most damaging implications of these counter-examples may be summarized as: there are examples in which a policy that is

¹ Mortensen [25] shows the validity of such an approximation in the context of a job search model. Inter alia, such an "approximation" was also used in the early literature on inventory (Karlin [21]), money demand (Baumol [2], Tobin [38]) and price adjustment (Barro [1]).

² These criteria are defined in Section 2.

³ For the weakest undiscounted criteria, the long-run average, these counter-examples establish that: in simple problems where discounted optima and their limits exist, a) long-run average optimal policies may not exist, b) even if they do, may be different from the discount limit and c) even discounted values need not converge to the long-run average value.

uniquely optimal under all undiscounted criteria may not be approximately optimal for any positive discount rate. However there is a unique limit to the discounted optima which is of course not optimal for zero discounting.

Those counter-examples turn out to be non-robust for economic problems. And that is one of the main messages of this paper. This task is accomplished in two stages: first we identify two simple sufficient conditions which imply convergence to long-run average and catching-up optima respectively. Then we show that basic structures like convexity and monotonicity, that are inherent in many economic models, make these conditions easy to check and indeed often imply them directly.

A number of different criteria exist for the undiscounted problem. The two most popular ones are clearly the long-run average and the catching-up (and its variants like the overtaking and weak maximality). For usage of the former, the reader is invited to see [16], [25], [27], [28] and the references in footnote one, among others. The latter criteria have been used e.g. in [18], [8], [24] and [35]. The two different criteria reflect two alternative and well-known approaches to the evaluation of infinite streams of returns. In particular, the long-run average considers only the asymptotic behavior of returns whereas the catching-up type criteria weigh finite evaluations as well. To address some well-known disadvantages of these two criteria we propose an intermediate criterion, the strong long-run average. This criterion refines the long-run average, which often discriminates insufficiently between alternative choices. Like the catching-up it gives positive (and equal) weight to each time period and, unlike the catching-up, has a recursive representation. The existence of a recursive representation makes it analytically very attractive since all of the arsenal of dynamic programming can be brought to bear on it. Furthermore, as we show, optimality under this criterion often implies catching-up optimality.

We turn now to a discussion of the results and we do this in an order which is exactly the opposite of that in which we posed the questions. We start with the results on the relationship between the different undiscounted criteria and proceed thereafter to the results on convergence. Finally we discuss a general mathematical result which underlies the rest of the analysis and which we believe is going to be useful in contexts beyond that of the present paper.

The Strong Long-Run Average and the Relationship Between Alternative Undiscounted Criteria: The strong long-run average criterion (formally defined in Section 2) maximizes the sum of expected returns net of long-run average. Theorem 4 says:

Optimality under the strong long-run average criterion is implied by catching-up optimality and in turn implies long-run average optimality. Moreover whenever the strong long-run average value is well-defined⁴ optimality under this criterion also implies catching-up optimality.

Convergence to the Long-Run Average: The only condition for convergence turns out to be boundedness of (normalized) value functions. Value boundedness essentially says that starting from a "higher" state does not imply an infinitely better value than starting from a "lower" state, even if the discount rate goes to zero. Theorem 3 says:

If a dynamic programming problem is value bounded, then the long-run average value is a constant over the state space and is the limit of normalized discounted values, a stationary long-run average optimal policy exists and, if a pointwise limit of stationary discounted optimal policies exist, then it is stationary long-run average optimal.

Further, we show that value boundedness falls out naturally from economic models which are either convex or stochastic (Propositions 1 and 2).

Convergence to the Strong Long-Run Average and Catching-Up: We show that a stronger boundedness condition on (normalized) value functions, called value finiteness, is sufficient for strong long-run average convergence and, by extension, catching-up convergence. Value finiteness is value boundedness plus the requirement that one cannot maintain one-period returns strictly in excess of the long-run average forever. Theorem 5 and (respectively) Corollary 2 say:

Under value finiteness, (normalized) discounted values converge weakly to the strong long-run average (resp. catching-up) value. Further, there is a stationary strong long-run average (resp. catching-up) optimal policy and if a pointwise limit of stationary discounted optimal policies exist, then it is a stationary optimal policy for the strong long-run average (resp. catching-up) problem.

In particular, the value finiteness condition can be checked in models where the long-run average value is easily computed. We illustrate with the neo-classical growth model (Proposition 3).

Neither value boundedness nor value finiteness are necessary conditions for convergence to long-run average and catching-up optimality, respectively. However, we

⁴ See Section 2 for a precise interpretation of this statement.

show by way of examples (Example 1 and accompanying discussion), that neither result is true without these assumptions.

General Parametric Variation: The mathematical backbone to the convergence analysis is a generalization of Fatou's lemma (Theorems 1 and 2). This is a generalization in that one studies the convergence of integrals when both the integrand and the probability measures over which they are being integrated change. This problem arises naturally in many parametric variational contexts (and we argue this more fully in Section 3). Consequently the lemma will hopefully be of much wider applicability than the current context.

Let us turn now to a discussion of the available literature. Long-run average convergence and existence results were established for countable state and either a) finite action spaces, by Ross [32], [34] and Taylor [37], or b) compact action spaces by Borkar [7]. Moreover, the discreteness of the state space was crucial to the analytic methods in these papers. Since discrete spaces are not very useful for economic applications, an extension of these results to state continua remained for long an open question. Ross [33] and Bhattacharya and Majumdar [4] did extend these results to general state spaces, but under very strong conditions (see the discussion following Theorem 3).⁵ In this paper we prove results for real state and compact action spaces (see Sections 2 and 4 for details), under a substantially weaker condition.⁶ To the best of our knowledge there is no previous literature on catching-up convergence and the equivalence relations between alternative undiscounted criteria.⁷

Section 2 sets out the dynamic programming problem. Section 3 states and discusses the Generalized Fatou's Lemma. This is used in Sections 4 and 6 where long-run average, strong long-run average and catching-up convergence is discussed. Section 5

⁵ Additionally, all of these results only establish convergence to the lower long-run average (see Section 2 for a definition), whereas our results establish that a) lower and upper values agree and b) that the limiting value is this common long-run average value.

⁶ Since this paper was completed, we have learnt of a result by Lehrer and Monderer [22] in which they show that under a condition similar to value boundedness, discounted average values converge to the upper long-run average value. They also show by way of an example that, under their condition, the lower long-run average value may be strictly less than this limit. It is unclear whether this analysis can be used to examine policy convergence or the existence of a stationary optimal policy in the long-run average problem.

⁷ It should be noted that the strong long-run average criterion has been used before in specific problems. The most celebrated is possibly the Ramsey criterion in optimal growth theory which maximizes the sum of utilities normalized by the utility of consumption at the golden rule. The golden rule consumption is of course the long-run average in the growth model. This concept was also used, in the growth context, by Radner [27].

presents the analysis on the relation between alternative criteria. Section 7 discusses applications to economic models while Section 8 presents a counter-example. Section 9 summarizes the issues and the results obtained in this paper. All proofs are in Section 10.

2. The Stochastic Dynamic Programming Problem

2.1. Some Definitions

A real-valued function f on \mathbb{R}^m is said to be non-decreasing if $z \geq x$ implies $f(z) \geq f(x)$, where \geq implies coordinate-wise weak domination on \mathbb{R}^m . All functions in the sequel will be real-valued. Let $(f_n)_{n \geq 0}$ be a sequence of non-decreasing functions defined on \mathbb{R}^m . By analogy with distribution functions, (f_n) is said to converge weakly⁸ (denoted $f_n \Rightarrow f$) if $\lim_{n \rightarrow \infty} f_n(x) = f(x)$ for all continuity points of f . Let $(\mu_n)_{n \geq 0}$ be a sequence of probability measures on the Borel σ -field of \mathbb{R}^m . Then (μ_n) is said to converge weakly to a probability measure μ , (denoted $\mu_n \Rightarrow \mu$), if $\overline{\lim}_{n \rightarrow \infty} \mu_n(C) \leq \mu(C)$ for all closed sets C (equivalently, $\overline{\lim}_{n \rightarrow \infty} \int f d\mu_n \leq \int f d\mu$ for all upper semi-continuous functions f on \mathbb{R}^m that are bounded above); see the Portmanteau theorem in [5].

2.2. Dynamic Programming Problems

A stochastic dynamic programming problem is specified by a quadruple $\langle S, A, q, r \rangle$. S is the state space, a non-empty closed subset of \mathbb{R}^m . A is the action space, assumed to be a non-empty compact metric space. $\{q(\cdot | s, a) : (s, a) \text{ in } S \times A\}$ is a family of transition probabilities, $q(\cdot | s, a)$ being a probability measure on the Borel σ -field of S conditional on (s, a) . In particular, if the state in the current period is s and the action chosen in the same period is a then the distribution of the state in the next period is generated through $q(\cdot | s, a)$. The family of transition probabilities will be assumed to be weakly continuous, i.e. $(s_n, a_n) \rightarrow (s_0, a_0)$ as $n \rightarrow \infty$ implies that $q(\cdot | s_n, a_n) \Rightarrow q(\cdot | s_0, a_0)$. Finally, r is the one-period reward function and associates with a current state s and action a , a reward $r(s, a)$. $r: S \times A \rightarrow \mathbb{R}$ will be assumed to be a bounded upper semi-continuous function. This structure and these assumptions are standard⁹ in dynamic programming formulations (see, for example,

⁸ The analogy is however incomplete. In particular, f_n does not define a probability measure, and each f_n defines a measure with different "total mass."

⁹ The assumption that is not standard in general versions of this problem is that S is a subset of \mathbb{R}^m . For most economic applications this is not restrictive and of course this buys us a lot of simplification in proofs, notably through the partial order on \mathbb{R}^m . See the comments following the Generalized Fatou's Lemma for possible extensions.

[23] for more details). In particular, this formulation includes continuous deterministic problems and stochastic formulations with continuous transition densities.

Let $h_t = (s_0, a_0, \dots, s_t)$ denote a history of states and actions and let H_t denote the set of all possible histories, $t \geq 0$. A policy $\langle f \rangle = (f_t)_{t \geq 0}$ is a sequence of measurable functions $f_t : H_t \rightarrow A$. A policy f^∞ is stationary if $f_t = f$, for all t , for some Borel-measurable f . A policy induces a conditional distribution on the state and action in every period, conditional on the initial state s_0 .

Let us turn now to the various criteria according to which payoff streams to alternative policies are evaluated.

Discounting: Under discounting we use the standard criterion which sums the expected discounted returns. So, the value function for discount factor δ in $[0,1)$ and initial state s_0 is

$$V(s_0, \delta) = \sup_{\langle f \rangle} E^{\langle f \rangle} \left[\sum_{t=0}^{\infty} \delta^t r(s_t, a_t) \mid s_0 \right] \quad (1)$$

where $E^{\langle f \rangle}$ denotes expectations under a policy $\langle f \rangle$.

For the zero discount rate ($\delta = 1$) problem, the two most popular criteria have been the long-run average and the catching-up. We discuss them first and then introduce a third criterion, the strong long-run average.

Long-Run Average: The long-run average criterion evaluates returns according to the limit of time averages. Since such a limit may not be well-defined, a first approximation looks at the lower limit of these averages. Given an initial state s_0 , the lower long-run average value is

$$v^*(s_0) = \sup_{\langle f \rangle} \liminf_{N \rightarrow \infty} \frac{1}{N} \left[\sum_{t=0}^{N-1} E^{\langle f \rangle} r(s_t, a_t) \mid s_0 \right] \quad (2)$$

The upper long-run average value is defined similarly, with limsup replacing liminf in (2). If the two values coincide, we shall call this the long-run average value function, and denote it simply as $v(s_0)$.

A policy is discounted (resp. long-run average) optimal from s_0 (on S) if it realizes the discounted (resp. the lower long-run average) value at s_0 (on S). If a stationary policy

is optimal, we shall call it a stationary discounted (respectively long-run average) optimal policy.

Catching-Up: The second optimality criterion for the undiscounted problem is the catching-up criterion. A policy $\langle f \rangle$ catches-up to another policy $\langle g \rangle$ if

$$\lim_{N \rightarrow \infty} \sum_{t=0}^{N-1} \left[E\langle f \rangle r(s_t, a_t) - E\langle g \rangle r(s_t, a_t) \mid s_0 \right] \geq 0 \quad (3)$$

$\langle f \rangle$ is catching-up optimal from s_0 (resp. on S) if it catches up with every other policy from s_0 (resp. on S)¹⁰.

Strong Long-Run Average: The two criteria above suffer from some well-known shortcomings. The long-run average is often too weak a criterion in that it does not discriminate among seemingly very different consequences. The catching-up is not a recursive criterion in that it is not the sum of period payoffs. These shortcomings have unpleasant implications for existence and characterization of optima in the undiscounted problem, as we shall argue in greater detail in the sequel. To address these problems we introduce the following criterion. The idea is simple: we want to be able to pick among long-run average optimal policies by looking at those which maximize the sum of expected returns net of the long-run average value. Since this sum may not be well-defined, a first approximation is to look at the lower limit of this sum. The lower strong long-run average value is defined as:¹¹

$$V^*(s_0) = \sup_{\langle f \rangle} \lim_{N \rightarrow \infty} \frac{1}{N} \left[\sum_{t=0}^{N-1} E\langle f \rangle r(s_t, a_t) - v(s_t) \mid s_0 \right] \quad (4)$$

The upper strong long-run average value is defined by replacing the liminf in (4) with a limsup. When the two values coincide we shall call this the strong long-run average value function, and denote it $V(s_0)$. A policy is strong long-run average optimal at s_0 (resp. on S), if it realizes the lower strong long-run average value at s_0 (resp. on S).

¹⁰ Variants of the catching-up that have been studied in the literature include the overtaking (e.g. Rubinstein [35]), in which the inequality (3) is strict, and the weakly maximal (e.g. Brock [8]), in which a policy is weakly maximal if no other policy catches up to it.

¹¹ We assume in this definition that the long-run average value is well-defined, i.e. that upper and lower values coincide. We show later that this is indeed the case under the value boundedness condition. If the two long-run average evaluations diverge, either could be used to appropriately define the strong long-run average criterion.

3. A Generalization of Fatou's Lemma

The main result proved in this section is a generalization of Fatou's lemma for non-decreasing functions on R^m . We state and discuss the result at this point since it is the critical variational inequality that underlies the comparative dynamic analyses that follow. Moreover, we believe that a result like it is essential to any stochastic variational analysis and hence is likely to be of use in a much wider context than the discount rate asymptotics we study in this paper.

Theorem 1: Let $(f_n)_{n \geq 0}$ be a sequence of non-decreasing upper semi-continuous functions on R^m that are bounded above. Suppose there is a non-decreasing upper semi-continuous function f such that $f_n \Rightarrow f$. Further, let $(\mu_n)_{n \geq 0}$ and μ be probability measures such that $\mu_n \Rightarrow \mu$. Then,

$$\lim_{n \rightarrow \infty} \overline{\int f_n d\mu_n} \leq \int f d\mu \quad (5)$$

The proof of Theorem 1 is to be found in Section 10. A few remarks on the generality of the result are, however, in order. If $\mu_n = \mu$ for all n , Fatou's lemma yields (5) since the functions are bounded above¹². If $f_n = f$ for all n , then $\mu_n \Rightarrow \mu$ implies (5) since f is an upper semi-continuous function that is bounded above. Theorem 1 allows both integrands and measures to vary. The only additional requirement¹³ is that the functions be non-decreasing and, surprisingly, this turns out to be sufficient to drive the general result. Easy examples can be constructed to show that this restriction cannot be dispensed with.

It is immediate that the Generalized Fatou's lemma has wider applicability than the context studied in this paper. Consider any stochastic optimization problem in which an action is chosen to maximize expected returns, where the return function is indexed by some parameter. How do the maximum returns and optimal actions vary in the parameter? The variational problem is precisely that in which the Generalized Fatou's lemma is immediately applicable. It implies that under mild continuity restrictions¹⁴, the maximized returns are upper semi-continuous in the parameter. In the proofs of Theorems 3 and 5 it

¹² As Lemma 3 shows, $f_n \Rightarrow f$ implies that $\lim_{n \rightarrow \infty} \overline{f_n(x)} \leq f(x)$, for all x . Fatou's lemma says, for any sequence $n \rightarrow \infty$ of functions bounded above, $\lim_{n \rightarrow \infty} \overline{\int f_n d\mu} \leq \int \lim_{n \rightarrow \infty} \overline{f_n} d\mu$. This immediately yields (5).

¹³ The domain is also being restricted to R^m . That this generalises to "similar" domains is suggested by the discussion at the end of the proof of Theorem 1.

¹⁴ There are stronger continuity restrictions (see, for example, [5, Theorem 5.5]), under which the limit of the integrals in (5) is actually the integral of the limit.

will be seen that the problem studied in the lemma arises naturally in discount factor analysis. Whenever the discount factor changes, it alters the value function and the optimal action in any state. The latter changes the distribution over the next period's state. The expected continuation payoffs change both because the integrand, the value function, changes and because the future distribution over states, the probability measures, change.

In the sequel the following corollary of the lemma will be useful. Suppose $(f_n)_{n \geq 0}$ are also bounded below pointwise, i.e. there is $M(\cdot)$ s.t. $f_n(x) \geq M(x) > -\infty$, for all n . Given that f_n is a sequence of non-decreasing functions on R^m that are bounded above, Helly's theorem (see [6, Theorem 29.3]) says that there is a non-decreasing upper semi-continuous function f , such that $f_{n'} \Rightarrow f$, along some subsequence. So, one has

Theorem 2: Under the hypothesis of Theorem 1 on (f_n) , (μ_n) and μ , and the additional hypothesis that (f_n) are pointwise bounded below,

a) there is a non-decreasing upper semi-continuous bounded function f such that $f_{n'} \Rightarrow f$ along a subsequence.

b)
$$\overline{\lim}_{n' \rightarrow \infty} \int f_{n'} d\mu_{n'} \leq \int f d\mu \quad (6)$$

Remark: Although the generalized Fatou's lemma is proved here for R^m , the result can be generalized to a separable normed linear space with a partial order, (see [13]).

4. Convergence to the Long-Run Average

The long-run average of returns (see (2) for a definition) is widely used as a decision criterion in economics and game-theory (for example, see [16], [25], [27], [28] among others) as well as in mathematics/operations research ([34] and the references therein). In many simple models the long-run average value is easy to compute and hence is used as a proxy to the discounted problem (for example, [2], [21], [25] and [38]).

In this section we investigate the relationship between the discounting and long-run average criteria, asking in particular: when is it the case that the limit of normalized values

and optimal policies under discounting are the long-run average value and optimal policy?¹⁵ We show that under a condition that we call value boundedness, such a convergence of values and policies does obtain. An implication of the result is that we have general conditions under which long-run average values and optimal policies exist. Since sufficiently general conditions on the existence of long-run average optima are not known (and indeed counter-examples are easy to construct, see Ross [34]), our results (Theorem 3 and Propositions 1 and 2) establish that conditions satisfied by most economic models are sufficient for existence. The convergence results also establish the robustness of using low discounts as an approximation to the undiscounted problem and vice-versa. Further, if the undiscounted optimum is unique then it is "close" to an exact optimum under discounting. The sufficient condition for convergence is easy to check and usually satisfied in economic applications (see Section 7 for a discussion). If the condition fails, the result may not hold and we give an example in Section 8. The results reported here can be adapted to strategic environments as well. An example is provided by an existence result for undiscounted dynamic games, contained in [14].

Fix any z in S . We define a normalized value function as

$$\psi_z(s, \delta) \equiv V(s, \delta) - V(z, \delta), \quad s \text{ in } S \quad (7)$$

For any z in S , the associated normalized value function family is $\{\psi_z(s, \delta) : \delta \text{ in } [0, 1)\}$. The value-boundedness assumption is:

(A1) There is $z \in S$, such that the associated normalized value function family is bounded above uniformly and bounded below pointwise, i.e. there is a function $M(\cdot)$ and $M \in \mathbb{R}$ s.t.

$$-\infty < M(s) < \psi_z(s, \delta) < M < \infty \quad (8)$$

In many economic models, starting from a higher initial state implies a higher value than that for starting from a lower initial state. In such problems, a natural candidate for z is the "highest" state. The uniform upper bound condition is automatically satisfied. The value-boundedness condition says therefore, that by starting from a "higher" initial state one cannot expect to do "infinitely better" than by starting from a "lower" initial state. The amount of "value loss" by starting from a lower state may however depend on how much

¹⁵ The reader should note that the convergence results here are distinct from the folk theorems of repeated games. The problem differs in that folk theorems establish lower hemi-continuity of the (perfect or Nash equilibrium) value correspondence while this analysis investigates the upper hemi-continuity of the policy correspondence and the continuity of the value function, as the discount rate goes to zero.

lower the state is. In Section 7 we show why this assumption is satisfied by many economic models.¹⁶

Theorem 3: Suppose that the dynamic programming problem is value bounded. Then, there is v in R , such that

(i) $v = \lim_{\delta \uparrow 1} (1-\delta)V(s,\delta)$, for all s and v is the long-run average value, which is a

constant over S .

(ii) there is a stationary optimal policy for the long-run average problem.

(iii) let $h(\cdot,\delta)$ be a stationary optimal policy for the discounted dynamic programming problem and suppose $h(\cdot,\delta) \rightarrow h$ pointwise, as $\delta \rightarrow 1$. Then h is a stationary optimal policy for the long-run average problem.

Remark The convergence result above can be straightforwardly extended to a formulation in which the set of feasible actions is state-dependent. Standard assumptions on the feasibility correspondence (e.g. see [17])¹⁷ suffice to establish Theorem 3 for this general case.

A few comments may be in order. (i) asserts that the upper and lower long-run average values coincide and is exactly the (state-independent) limit of discounted average values. So under value boundedness, long-run average value convergence holds in the strongest possible sense. (ii) establishes the undiscounted analog of Blackwell's and Maitra's results in the discounted problem, i.e. that from the point of view of optimality, attention may be restricted to pure stationary policies. Finally (iii) is possibly the most important of the results in that it says that discounted optimal policies have the same approximate qualitative features as undiscounted optima. The converse holds if the undiscounted optimal policy is unique. Example 1 in Section 8 discusses a case in which none of the results (i)-(iii) are valid. Value boundedness is violated in that example.

Theorem 3 was proved for Polish state spaces by Ross [33] and Bhattacharya and Majumdar [4], under the assumptions that the normalized value function family is uniformly bounded and equicontinuous.¹⁸ These conditions are of course strong and often

¹⁶ There are other ways to think of the value boundedness assumption as well. It could be interpreted as an implication of the stability of the state paths generated by the discounted optimal policies. For example, if the states converge to invariant distributions at a uniform rate, value boundedness would obtain.

¹⁷ Let $A(s)$ denote the feasible actions at state s . From the proof of Theorem 3 it is easy to see that a sufficient condition for Theorem 3 to be true in this formulation as well is that the feasibility correspondence $A(\cdot)$ be upper hemi-continuous.

¹⁸ The action space in Ross is finite whereas it is a compact metric space in Bhattacharya and Majumdar.

difficult to verify. By restricting attention to a state space which has a lattice structure, we have been able to prove the result under the single, weaker, condition of value boundedness.

A lot of economic applications analyze the single-dimensional state case. In fact in several of them the following policy monotonicity result also holds (by way of examples see the search model of Burdett-Mortensen [10], optimal growth models of Dechert and Nishimura [11] and Dutta [12] etc.)

(A2) Suppose the state space S and the action space A are subsets of R . Then the problem satisfies monotone usage if for all discounted optimal policies $h(\cdot, \delta)$, $s' \geq s$ implies $h(s', \delta) \geq h(s, \delta)$.

By a variant of the Helly selection theorem, we can extract a convergent subsequence of $\{h(\cdot, \delta), \delta \text{ in } [0,1)\}$. That yields:

Corollary 1: Suppose in addition to the hypothesis of Theorem 3, the problem shows monotone usage. Then any sequence of discounted values and optimal policies converge to a long-run average optimal value and policy, as the discount factor goes to one.

Since no single (or finite set of) periods matter for the long-run average value, policies which appear a priori unreasonable on account of their finite period behavior, may sometimes turn out to be long-run average optimal. An easy example of this is Gale's "cake-eating" problem (Gale [18])¹⁹ in which the limit of discounted optimal policies is to consume nothing forever. Since zero is the long-run average consumption of all policies, the pure accumulation policy of consuming nothing forever is in particular long-run average optimal. However, alternative criteria as the catching up or "the maximum total returns among those which have the same long-run average" are more discriminating. When is it the case that the limit of optimal policies under discounting in fact converge to such optima? This question will be addressed in the next sections.

5. The Strong Long-Run Average and the Relation Between Alternative Undiscounted Decision Criteria

¹⁹ A cake of initial size $y > 0$ can be consumed in part every period. The objective is to maximize lifetime returns subject to the constraint that lifetime consumption not exceed y .

This section links the long-run average and the catching-up criteria by introducing a third undiscounted criteria, the strong long-run average. Note that as simple examples like "cake-eating" show, the long-run average criterion may often not be discriminating enough. Since it ignores all finite sets of periods, it may contain too "large" a set of optimal policies. At the same time we would like to find a criterion that is more immediately usable than the catching-up. We show now that the strong long-run average, defined in Section 2.2, meets both of these demands.

Recall that the lower strong long-run average of a policy $\langle f \rangle$ is defined as

$$H^{\langle f \rangle}(s_0) = \liminf_{T \rightarrow \infty} \left[E^{\langle f \rangle} \sum_{t=0}^{T-1} r(s_t, a_t) - vT \mid s_0 \right] \quad (9)$$

where we use the result of the previous section to assume that the long-run average value v exists and is constant over S . The strong long-run average is the aggregate surplus net of a constant, the maximum average return v . Clearly, any long-run average inoptimal policy has a strong long-run average of negative infinity²⁰. We show that it is closely linked to the catching-up criterion and by virtue of having a recursive representation, turns out to be an extremely useful approach to the latter. Recall further that, by analogy with the long-run average criterion, the upper strong long-run average is defined by replacing the liminfs in (9) with limsups. If the maximized returns are independent of this choice, i.e. if the upper and lower strong long-run average values coincide we say of course that the strong long-run average value exists. The theorem that follows characterizes the links between the three undiscounted criteria. In particular, it says that when the strong long-run average value exists, then its ordering is precisely that of the catching-up criterion.

- Theorem 4: (i) Suppose a policy $\langle f \rangle$ is catching-up optimal from s_0 . Then, it is strong long-run average optimal.
(ii) Suppose a policy $\langle f \rangle$ is strong long-run average optimal from s_0 and $V^*(s_0) > -\infty$. Then, it is long-run average optimal.
(iii) Suppose that $-\infty < V_* \leq V^* < \infty$. A strong long-run average optimal policy from s_0 is catching-up optimal if $V_*(s_0) = V^*(s_0)$.

(i) and (ii) establish the increasing sensitivity of undiscounted criteria, from long-run average to strong long-run average to catching-up. (i) and (iii) demonstrate the close

²⁰ This statement was formally proved by Jeanjean [20] in the stochastic growth model where he showed that any policy with a lower strong long-run average greater than negative infinity (called a "good" plan by him, following Gale [18]) must be long-run average optimal.

links between the strong long-run average and catching-up optima. As a corollary this result shows the twin uses of the strong long-run average criterion. In cases in which the long-run average is too non-discriminatory a criterion, (ii) suggests that a characterization of optima through the strong long-run average may yield a more "reasonable" ordering. On the other hand, in problems in which existence of catching-up optima are difficult to establish, (i) and (iii) say that a fruitful approach would be to study the strong long-run average optimality problem instead.

6. Convergence to Strong Long-Run Average and Catching-Up Optima

The catching-up or its variants, the overtaking and weak maximal (see footnote 10), have been widely used in economics and specially in capital theory. For examples see [8], [18], [24] among others, and also [35] for an application in game theory. The attractiveness of the criterion stems from the fact that it weights each period equally and yet gives them positive weights. In this section we ask: under what conditions do the discounted optimal values and policies converge to the strong long-run average optimal values and policies? Since the latter criterion does in fact possess a recursive representation, we can use the structure of dynamic programming to analyze this question. *Pari passu*, we will use the results of the previous section to give conditions under which discounted optima converge to catching-up optima.

How much stronger are the restrictions that need to be placed to guarantee strong long-run average convergence and when do economic models satisfy these conditions? A condition stronger than value boundedness will be required. To see this note that it is easy to show that in the cake-eating example, with bounded marginal utilities, value-boundedness is satisfied. Consequently the discounted limit is long-run average optimal but this limit is the pure accumulation or zero consumption policy, which is trivially not strong long-run average optimal. Further, by an easy adaptation of the counter-example in Flynn [15] we show in Section 8 that even if a problem is known to have a strong long-run average optimal policy (the cake-eating example does not) the discounting limit may not be optimal, under this criterion.

We show that a condition which we call value finiteness is sufficient to establish value and policy convergence. This condition may be thought of as value boundedness plus a stability requirement of candidates for optimal policies. The condition is easy to check in cases where the long-run average is easy to compute and we give an example in Section 7. Of course, value finiteness is also a sufficient condition then, under which

strong long-run average and (if the conditions of Theorem 4 are satisfied) catching-up optima exist. Since general sufficient conditions for existence under these criteria are difficult to come by (see [18] for counter-examples), this result may facilitate analyses under them. The convergence results imply as corollaries the robustness of using the discounted and undiscounted models as approximations for each other.

The analog of the strong long-run average under discounting, we call the normalized discounted value:

$$W(s_0, \delta) = \sup_{\langle f \rangle} E^{\langle f \rangle} \sum_{t=0}^{\infty} \delta^t [r(s_t, a_t) - v \mid s_0] \quad (10)$$

We need this normalization to yield bounded functions. Notice that $W(s, \delta) = \psi_z(s, \delta) + W(z, \delta)$. If there is some state such that it yields asymptotically an exact average of v , i.e. $z \in S$ such that $W(z, \delta) \rightarrow 0$ as $\delta \uparrow 1$, then the two normalizations W and ψ_z are asymptotically equivalent. In order not to introduce further notation we shall maintain this assumption in the sequel. Following Gale [18] define a good policy from initial state s as one with $H^f(s) > -\infty$. Obviously, the search for strong long-run average optima can be restricted to good policies if they exist. The value finiteness assumption is:

(A3 i) Value Boundedness and for some $z \in S$, $W(z, \delta) \rightarrow 0$ as $\delta \uparrow 1$

(A3 ii) For all good policies $\langle f \rangle$, and all initial states, $\lim_{T \rightarrow \infty} \overline{E_T^f \hat{W}} \leq 0$

where $\overline{E_T^f \hat{W}}$ is the expectation of the weak limit of $\{W(., \delta): \delta \text{ in } [0, 1)\}$, under the T -period distribution induced by $\langle f \rangle$ (and initial state s).

Value finiteness may be thought of as an assumption which requires (in addition to value boundedness) stability of good policies in the following sense. The first part of the assumption says that there is some state whose maximum returns are equivalent to the maintenance of the constant long-run average v forever. In other words there is some steady state or invariant distribution which is (approximately) maintained in perpetuity if the initial state is z . Moreover, and this is what the second part of the assumption says, good policies converge to this steady state over time.

Theorem 5: Suppose the dynamic programming problem satisfies value-finiteness. Then,

(i) for any sequence $\delta_n \uparrow 1$, $W(., \delta_n) \Rightarrow V_*(.)$, i.e. the discounted normalized values converge weakly to the undiscounted lower strong long-run average value.

(ii) there is a stationary optimal policy for the strong long-run average optimality problem

(iii) let $h(\cdot, \delta)$ be a stationary optimal policy for the discounted dynamic programming problem. If $h(\cdot, \delta) \rightarrow h(\cdot)$ pointwise, then h is a stationary optimal policy for the strong long-run average problem.

Using Theorem 4, we have conditions under which a discounted limit is in fact catching-up optimal.

Corollary 2: Suppose that $V_* = V^*$ on S . Then, under the hypothesis of Theorem 5, the results of the theorem hold for the catching-up optimality criterion.

7. Applications and Discussion

A crucial test of the usefulness of the general convergence results, Theorems 3 and 5, is of course how conveniently the value boundedness and value finiteness conditions can be checked. A related question is: are such conditions satisfied in a large class of economic models? In particular, are there general features of economic problems which imply such boundedness?

7.1 Value Boundedness in Resource Rich Problems

We now show that value boundedness is satisfied in dynamic economic models in which some "communication" is possible, or expressed differently, it is possible to (eventually) grow from "low" (in payoff terms) states. This by itself should come as no surprise since the value boundedness condition says that payoffs from "low" states should not lag infinitely behind payoffs from "high" states. What is extremely useful is that by exploiting underlying structures like convexity and monotonicity, in economic models such communication conditions can be expressed in very parsimonious terms.

For expositional ease we shall restrict ourselves to a class of monotone problems that we call resource rich problems. In many economic applications, the value function is in fact non-decreasing in the initial state, i.e. $s' \geq s$ implies $V(s', \delta) \geq V(s, \delta)$, for all δ . Such problems we may call resource rich.²¹ We now show that in resource rich problems, value boundedness is easy to check and often satisfied. To facilitate the presentation we shall restrict attention to the state space as a compact interval of R , say $[0, 1]$. It should be emphasized that both restrictions are being placed to allow an elementary presentation of the

²¹ Typically, this follows from the fact that one period returns are non-decreasing in the state and "higher states are more likely from higher states", i.e. $s' \geq s$ implies $q(\cdot | s', a)$ stochastically dominates $q(\cdot | s, a)$. Alternatively, free disposal by itself suffices.

results that follow. In each case the further arguments, when these restrictions are relaxed, are straightforward. Further, we present two of a number of possible variants and make no attempt to state the strongest possible results.

The first result focusses on convex problems, i.e. those in which the value function $V(\cdot, \delta)$ is concave, for δ in $[0,1)$. Many economic applications are resource rich and convex. Examples may be found in aggregative growth models [9], search models [25], macro models (see [36]), natural resource models [31], etc.

Proposition 1: Suppose the dynamic programming problem is resource rich and convex. Suppose further that there is some action a' in A such that $q((0,1] | 0, a') > 0$. Then, value boundedness holds.

To paraphrase Proposition 1: in convex, resource rich problems, it suffices to check the lowest state. If there is some action which induces a positive probability of being in higher states next period, such "communication" possibilities suffice to establish value boundedness. Such a condition is not satisfied if the lowest state is absorbing as e.g. in the neo-classical growth model. However if growth is possible from any positive state, value boundedness holds on $(0,1]$.²² The following result can then be proved:

Corollary 3: Suppose the dynamic programming problem is resource rich and convex. Suppose further that there is $s^* > 0$ with the property that for any s in $(0, s^*]$, there is an action $a(s)$ for which i) $q([0, s] | s, a(s)) = 0$ and ii) $q((s, 1] | s, a(s)) > 0$. Then, value boundedness holds on $S=(0,1]$.

In the neoclassical growth problem, the familiar Inada conditions on the production possibilities are precisely i) and ii). Clearly, growth and renewable resources satisfy the assumptions of Corollary 3. The search and macro models referred to earlier satisfy the hypothesis of Proposition 1 directly. In fact the only example that comes to mind where growth is impossible is the pure exhaustible resource or cake-eating case. By independent argument it should be easy to see that this case exhibits value boundedness if and only if the marginal utilities of consumption are bounded.

²² The usual Inada conditions on utility imply that it is inoptimal to get to state 0 from any positive state. The entire exercise of Theorem 3 can then be carried out on $S=(0,1]$

An alternative structure in economic models in which value boundedness follows directly is the pure stochastic case. Let us define a dynamic programming problem to be stochastic if there is a positive density, $g: S^2 \times A \rightarrow R_{++}$ s.t.

$$q([c,d] | s,a) = \int_c^d g(s'; s,a) ds'$$

Proposition 2: Suppose the dynamic programming problem is resource rich and stochastic. Then, it is value bounded.²³

7.2. Value Finiteness

The additional hypothesis in value finiteness requires comparison of the discounted values to the long-run average value. Unlike value boundedness therefore, this is a condition whose verification requires a qualitative analysis of both the discounted and long-run average problems. In models where the long-run average value is easy to compute such verification may not, however, be difficult. We illustrate with the deterministic aggregative neo-classical growth model. This model is extremely well-known (see for instance [36]) and hence will not be repeated here.

Proposition 3: In the aggregative growth model, discounted optimal policies and (normalized) values converge to the unique catching-up optimal policy and value.²⁴

8. An Example

Flynn [15] showed in an example that discounted limits may not be long-run average optimal even when such optima exist. An easy argument shows that even strong long-run average optima exist in his example. So even in problems that are "well-behaved" and have optima under a strong ordering like the strong long-run average, discounted limits may violate the weakest optimality criteria. Of course this example violates value boundedness.

Example: The discounted optimal policy and value limits are not long-run average optimal even though strong long-run average optima exist.

²³ As with Corollary 3, if $s=0$ is a distinguished absorbing state, then value boundedness holds on $S=(0,1]$.

²⁴ By catching-up value I mean the strong long-run average value, which will be shown to be an equivalent criterion here.

The state space is the positive integers and zero and $A = \{0,1\}$. Whenever the state is any positive integer it moves up one independently of the action chosen. At state zero it moves up one if action 1 is chosen but remains at zero if action 0 is chosen. Let $r(n,i) = r_n$, $n \geq 1$ and $r(0,i) = \lambda r^* + (1-\lambda)r_*$, λ in $(0,1)$ where $(r_n)_{n \geq 1}$ is a bounded sequence such that

$$r^* \equiv \liminf_{\delta \uparrow 1} (1-\delta) \sum_{n=1}^{\infty} \delta^{n-1} r_n > \liminf_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N r_n \equiv r_* \quad (11)$$

Denote $r(0,i) \equiv k$ and let $(\delta_m)_{m \geq 0}$ be a sequence such that

$$(1-\delta_m) \sum_{n=1}^{\infty} \delta_m^{n-1} r_n \geq k$$

Given (11) it is not difficult to show that taking the action 1 always is δ optimal for all $\delta < 1$, but from state 0 one should take action 0 under the long-run average criterion. In fact,

Lemma $a(n)=0$ for all n, is a stationary strong long-run average optimal policy.

On the other hand, as we argued earlier, the cake-eating example with bounded marginal utilities provides an example where the discounting limit is not catching up optimal, even though value boundedness is satisfied. Value finiteness fails in that example.

9. Conclusions

We examined the general issue of parametric variation in dynamic models in the context of variations in the discount factor. We looked for conditions under which discounted optimal policies and values converge to undiscounted optima. It is well known that the discounted average of an infinite sequence of numbers may not converge to the long-run average as the discount factor goes to one. This discontinuity of the objective function, with respect to the discount factor, is at the heart of counter-examples that show that a unique limit of discounted optima may be inoptimal in the long-run average sense whereas the (unique) undiscounted optimum may not even be ϵ -optimal in the discounted problem. For economic models such examples were shown to be pathological. Long-run average convergence always holds under the condition of value boundedness and this condition was shown to follow immediately in convex or stochastic structures which many economic models satisfy. In all such models one can then view the discounted and

undiscounted problems as special cases of the same family of optimization problems and use optimal policies in one to approximate optimality in the other.

Since the long-run average is insufficiently discriminating as a decision criterion we next introduced a refinement of it, the strong long-run average. This was shown to have two other attractive features: it has a recursive representation and whenever the strong long-run average value function is well-defined, i.e. the upper and lower values coincide, it is equivalent to the catching-up criterion.

Discounted optimal values and policies were shown to converge to the strong long-run average optimal values and policy if value finiteness is satisfied. This condition, although not as easy to check as value boundedness, was shown to be verifiable in economic models where the long-run average value is easy to establish as in the neoclassical growth model. A final contribution of this paper is the Generalized Fatou's lemma which provides a necessary tool for other problems in stochastic dynamic variational analysis.

10. Proofs

10.1 Generalized Fatou's Lemma

In this sub-section we prove the generalized Fatou's lemma, Theorem 1, by way of proving the following preliminary results (Lemmas 1-4). Since all the functions are uniformly bounded above, without loss of generality we take them to be non-positive.

Lemma 1: a) There is a sequence of continuous functions $(v_k)_{k \geq 0}$ s.t. $0 \geq v_k \geq f$, $v_k \downarrow f$ pointwise. Further, b) $v_k \geq -k$.

Lemma 2:
$$\lim_{n \rightarrow \infty} \overline{\int v_k d\mu_n} \leq \int v_k d\mu, k \geq 0 \quad (12)$$

Lemma 3: For all $s_n \rightarrow s$,
$$\lim_{n \rightarrow \infty} \overline{f_n(s_n)} \leq f(s)$$

Lemma 4:
$$\lim_{n \rightarrow \infty} \overline{\int f_n d\mu_n} \leq \lim_{n \rightarrow \infty} \overline{\int v_k d\mu_n}, \text{ for all } k. \quad (13)$$

Clearly (12) and (13) will suffice to prove

$$\lim_{n \rightarrow \infty} \overline{\int f_n d\mu_n} \leq \int v_k d\mu, k \geq 0$$

By the construction of v_k , it is clear that $\int v_k d\mu \rightarrow \int f d\mu$ as $k \rightarrow \infty$, from a straightforward application of the monotone convergence theorem. Hence the proof of Theorem 1 will be complete.

Proof of Lemma 1²⁵: Part a) of the lemma is a well-known result for real-valued functions (defined not just on \mathbb{R}^m but on any metric space as a matter of fact); see, for example, Goffman [19]²⁶. If the sequence of continuous functions so generated, say v_k , does not satisfy b), modify it as follows: pick the function $\max(v_k, -k)$. To avoid further notation we shall call this modified sequence v_k as well.

Proof of Lemma 2: The result follows from the fact that $\mu_n \Rightarrow \mu$, given that v_k is a continuous function (see [5, Theorem 2.1]).

Proof of Lemma 3: Suppose not. Then, there is $s_n \rightarrow s$, $\theta > 0$, $\epsilon > 0$, and $N < \infty$ such that $f_n(s_n) > f(s) + 2\epsilon$, $n \geq N$ and $\|s_n - s\| < \theta$, where $s + \theta e \equiv (s_1 + \theta, \dots, s_m + \theta)$ is a continuity point of f and $f(s + \theta e) < f(s) + \epsilon$. All this is possible to do since f is right continuous and non-decreasing (and hence has a countable set of discontinuity points on the ray from the origin passing through s). Combining the inequalities we get

$$f(s + \theta e) < f_n(s + \theta e) - \epsilon, \quad n \geq N \quad (14)$$

Since $f_n \rightarrow f$ at all continuity points, and $s + \theta e$ is one such, (14) yields a contradiction.

Proof of Lemma 4: Define, for a fixed k ,

$$E_p = \{s : f_n(s) < v_k(s)(1+\epsilon), n \geq p\}$$

Let us first show that E_p is open and $E_p \uparrow \mathbb{R}^m$. Consider any sequence in E_p^c . For any such sequence $(s_t)_{t \geq 0}$, $s_t \rightarrow s$, we have (renumbering wlog) either $f_t(s_t) \geq v_k(s_t)(1+\epsilon)$, or $f_n(s_t) \geq v_k(s_t)(1+\epsilon)$ for some fixed n . In the former case, Lemma 3 and the continuity of v_k implies $f(s) \geq v_k(s)(1+\epsilon)$, a contradiction. So, it must be the latter case and hence from the upper semi-continuity of f_n , it follows that s is in E_p^c . Similarly, if there is s in E_p^c for all p , then along a subsequence $f_n(s) \geq v_k(s)(1+\epsilon)$. Lemma 3 again yields a contradiction.

²⁵ This proof of Lemma 1 was suggested by Mukul Majumdar.

²⁶ A candidate sequence is $v_k(s) = \sup_y \{ f(y) - k \|s - y\| \}$.

Since $E_p \uparrow R^m$, for $\epsilon > 0$, there is p such that $\mu(E_p^c) < \epsilon$. From hereon, the index p , in addition to the index k , will also remain fixed. Since $\mu_n \Rightarrow \mu$ and since E_p^c is a closed set, there is $N < \infty$ such that $\mu_n(E_p^c) < \epsilon$, $n \geq N$. Notice that

$$\int v_k d\mu_n \geq \int_{E_p} v_k d\mu_n - k\epsilon \quad n \geq N \quad (15)$$

Hence,

$$\begin{aligned} \int f_n d\mu_n &= \int_{E_p} f_n d\mu_n + \int_{E_p^c} f_n d\mu_n \\ &\leq (1+\epsilon) \int_{E_p} v_k d\mu_n \\ &\leq (\int v_k d\mu_n + k\epsilon)(1+\epsilon) \end{aligned} \quad (16)$$

From (16) the lemma follows immediately. Hence, the proof of Theorem 1 is complete. •

Suppose $R^m \supset S$ has the "northeastern cone" property, i.e. for all x in S , $x' \geq x$ implies $x' \in S$. The reader is invited to check that all the above arguments go through if the functions f_n happen to be defined on such a domain.²⁷

Corollary A: Suppose S is a closed set with the "northeastern cone" property. Under the hypothesis of Theorem 1, the Generalized Fatou's lemma holds.

10.2. Long-Run Average Convergence

In this sub-section we prove Theorem 3 by way of lemmas 5-9.

Lemma 5: A closed set S , $R^m \supset S$, can be embedded in a closed, "northeastern cone" C , $R^{2m} \supset C$.

Proof: For any s in R^m , let $s^* = (s, -s)$. Define $S^* = \{s^* \in R^{2m}: s \in S\}$. Consider the northeastern cone of S^* ,

$$C = \{z \in R^{2m}: \exists s^* \text{ in } S^*, z \geq s^*\} \quad (17)$$

It is immediate that $z \in C$, $z' \geq z$ implies $z' \in C$. Let $(z_n)_{n \geq 0}$ be a sequence in C , $z_n \rightarrow \bar{z}$. In particular, $z_n \geq s_n^*$, for some s_n^* in S^* . Writing z_n^1 for the first m coordinates of z_n , and z_n^2 for the second, we have: $z_n^1 \geq s_n \geq -z_n^2$. Clearly, (s_n) is a bounded

²⁷ Naturally, the definition of weak convergence has to be adapted to a domain, the set S , which may not now be open. We say $f_n \Rightarrow f$, if $f_n(s) \rightarrow f(s)$, for all continuity points of f which are in the interior of S .

sequence, and hence on a subsequence has a limit, $\bar{z} \in S$. But then, $\bar{z} \geq \bar{s}^*$, i.e. $\bar{z} \in C$. That establishes the closedness of C .•

Lemma 6: Any upper semi-continuous function $f: S \rightarrow \mathbb{R}$, can be extended to the domain C , through a function $F: C \rightarrow \mathbb{R}$ with the property that

- i) F is non-decreasing, upper semi-continuous
- ii) $F(s^*) = f(s)$, s in S

Proof: Consider the following feasible correspondence and maximized returns

$$\begin{aligned}\phi(z) &= \{s \in S: z \geq s^*\} \quad z \in C \\ F(z) &= \max f(s), \quad s \in \phi(z)\end{aligned}$$

Since ϕ is a non-decreasing correspondence, it immediately follows that F is non-decreasing. By construction, $\phi(z) \neq \emptyset$, z in C . Standard arguments, paralleling the proof in Lemma 5, establish that ϕ is an upper hemi-continuous correspondence. Since f is usc, Theorem 2 of Berge [3, p. 116] establishes that F is usc. Finally, note that $\phi(s^*) = \{s\}$, s in S . From this, it follows that $F(s^*) = f(s)$.•

From [23] it is known that, under the hypotheses of dynamic programming used in this paper, the value function in any discounted problem, $V(\cdot, \delta)$, is usc. Consequently, the construction of Lemma 6 yields that the normalized value functions $\psi(\cdot, \delta)$ (which are usc) can be extended to a non-decreasing usc function $\Psi(\cdot, \delta)$ defined on the northeastern cone of S . Consider any sequence $\delta_n \uparrow 1$. Given value boundedness, by Helly's theorem there is a subsequential weak limit $\Psi(\cdot, \delta_n) \Rightarrow \hat{\Psi}(\cdot)$. Maintain the numbering of the original sequence.

Lemma 7: For all $s^* \in S^*$, $\epsilon > 0$, there is $s_n^* \in S^*$, $s_n \rightarrow s^*$ such that

$$\liminf_{n \rightarrow \infty} \Psi(s_n^*, \delta_n) \geq \hat{\Psi}(s^*) - \epsilon \quad (18)$$

Proof: We first show that there is $z_n \in C$, $z_n \downarrow s^*$ s.t. $\liminf \Psi(z_n, \delta_n) \geq \hat{\Psi}(s^*) - \epsilon$. Consider any sequence of continuity points of $\hat{\Psi}$, $z_n \downarrow s^*$. For each z_n , there is k_n , s.t.

$$\begin{aligned}\Psi(z_n, \delta_{k_n}) &> \hat{\Psi}(z_n) - \epsilon \\ &\geq \hat{\Psi}(s^*) - \epsilon\end{aligned}$$

Re-numbering the sequence z_n if need be, we have the desired inequality. But $\Psi(z_n, \delta_n) = \Psi(s_n^*, \delta_n)$ for some $s_n \in \phi(z_n)$. The lemma follows.•

Pick any $s^* \in S^*$, $\varepsilon > 0$ and consider the sequence (s_n^*) established in Lemma 7. The standard Bellman optimality equation yields by easy manipulation

$$\Psi(s_n, \delta_n) + (1 - \delta_n)V(z, \delta_n) = r(s_n, a_n) + \delta_n \int \Psi(\cdot, \delta_n) dq(\cdot | s_n, a_n) \quad (19)$$

for some optimal choice a_n at s_n . In the obvious way, extend r and V to domains on C , i.e. $r^*(s^*, a) = r(s, a)$, and $V^*(s^*, \delta) = V(s, \delta)$ ²⁸. Similarly, q can be straightforwardly extended to a probability measure q^* on C . Re-writing (19),

$$\Psi(s_n^*, \delta_n) + (1 - \delta_n)V^*(z^*, \delta_n) = r^*(s_n^*, a_n) + \delta_n \int \Psi(\cdot, \delta_n) dq^*(\cdot | s_n^*, a_n)$$

By the Generalized Fatou's lemma (Corollary A to Theorems 1 and 2),

$$\overline{\lim}_{n \rightarrow \infty} \int \Psi(\cdot, \delta_n) dq^*(\cdot | s_n^*, a_n) \leq \int \hat{\Psi}(\cdot) dq^*(\cdot | s, a') \quad (20)$$

where a' is a subsequential limit of a_n and this limit exists since A is compact. Further $\overline{\lim}_{n \rightarrow \infty} r(s_n, a_n) \leq r(s, a')$. Letting $v \equiv \lim (1 - \delta_n)V^*(z^*, \delta_n)$, which limit exists possibly on a further subsequence, the above inequalities, together with Lemma 7 yields

$$\begin{aligned} \int \hat{\Psi} dq(\cdot | s^*, a') + r^*(s^*, a') &\geq \overline{\lim} r^*(s_n^*, a_n) + \overline{\lim} \int \Psi(\cdot, \delta_n) dq^*(\cdot | s_n^*, a_n) \\ &\geq \overline{\lim} [r^*(s_n^*, a_n) + \int \Psi(\cdot, \delta_n) dq^*(\cdot | s_n^*, a_n)] \\ &\geq \underline{\lim} \Psi(s_n^*, \delta_n) + \lim (1 - \delta_n) V^*(z^*, \delta_n) \\ &\geq \hat{\Psi}(s^*) + v - \varepsilon \end{aligned}$$

Hence,

$$\max_a \left\{ r^*(s^*, a) + \int \hat{\Psi}(\cdot) dq^*(\cdot | s^*, a) \right\} \geq \hat{\Psi}(s^*) + v$$

or equivalently, defining $\hat{\Psi}(s) = \hat{\Psi}(s^*)$, s in S ,

$$\max_a \left\{ r(s, a) + \hat{\Psi}(\cdot) dq(\cdot | s, a) \right\} \geq \hat{\Psi}(s) + v \quad (21)$$

Lemma 8: $v \leq v^*(s)$, i.e. v is no greater than the long-run average value, defined through lower averages, from any s in S .

²⁸ The V^* here should not be confused with the same notation used for the strong long-run average value function defined elsewhere!

Proof: By the measurable selection theorem of [23]²⁹, there is a measurable function h which attains the maximum in (21).

$$r(s, h(s)) + \int \hat{\psi}(\cdot) dq(\cdot | s, h(s)) \geq \hat{\psi}(s) + v, s \text{ in } S \quad (22)$$

Iterating the inequality (27) yields

$$\begin{aligned} \hat{\psi}(s) + Tv &\leq E^h \sum_{t=0}^{T-1} r_t + E_T^h \hat{\psi} \\ &\leq E^h \sum_{t=0}^{T-1} r_t \end{aligned} \quad (23)$$

Dividing by T , letting $T \rightarrow \infty$ and taking limits-inferior yields the fact that the long-run average returns defined along lower averages, from the stationary policy h^∞ is at least as large as v . The lemma follows. •

Lemma 9: $v \geq v^*(s)$

Proof³⁰: Pick any policy $\langle f \rangle$ and define $c = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} r_t^f$, where r_t^f is the expected returns in period t from using the policy $\langle f \rangle$, conditional on initial state s . We show, $v \geq c$.

Claim 1: $\forall \epsilon > 0$ and $\lambda > 0$, there is T s.t.

$$\frac{1}{k} \sum_{t=T}^{T+k-1} r_t^f > c - \epsilon \quad 0 < k \leq \lambda \quad (24)$$

Pf. Clearly a contradiction to (24) would be the existence of some $\epsilon > 0$ and $\lambda > 0$ such that no matter which T we pick, there is an associated $0 < k(T) \leq \lambda$ with average returns between T and $T + k(T) - 1$ no larger than $c - \epsilon$. Clearly this implies the existence of a sequence $\{T_m\}$ satisfying: i) $T_{m+1} - T_m \leq \lambda$ and ii) $\frac{1}{T_{m+1} - T_m} \sum_{t=T_m}^{T_{m+1}-1} r_t^f \leq c - \epsilon$. For

notational convenience let us denote the average to returns over T periods as $A(T)$. From the preceding argument it is immediate that $A(T_m) \leq c - \epsilon$, for all m . It is also clear that

²⁹ Although this selection theorem was originally proved for bounded usc functions, the same proof can be shown to work for usc functions bounded above.

³⁰ This proof was suggested by Ehud Lehrer, who also pointed out a mistake in an earlier proof.

there is T and T_m satisfying i) $A(T) > c - \frac{\varepsilon}{2}$ and ii) $|A(T) - A(T_m)| < \frac{\varepsilon}{2}$. It is then obvious that the last two sets of properties yield a contradiction.

Claim 2: $\forall \varepsilon > 0$ and $0 < \delta < 1$, there is λ (and an associated T from claim 1) with the property

$$(1-\delta) \sum_{t=T}^{\infty} \delta^{t-T} r_t^f \geq c - 2\varepsilon \quad (25)$$

Proof: Pick any sequence $\lambda(\delta)$ with the property that $\delta^{\lambda(\delta)} \rightarrow 0$ as $\delta \rightarrow 1$. A candidate for such a sequence is $\lambda(\delta) = \frac{1}{(1-\delta)^2}$. Since returns are bounded we can, wlog, take them to be non-negative. Then,

$$(1-\delta) \sum_{t=T}^{\infty} \delta^{t-T} r_t^f \geq (1-\delta) \sum_{t=T}^{T+\lambda-1} \delta^{t-T} r_t^f \quad (26)$$

Let us define a sequence r_t^f with the property that its average over the cycle T through $T+\lambda-1$ is exactly $c - \varepsilon$. In particular, we can do this by defining $r_t^f = r_t^f$, $t = T, \dots$

$T+\lambda-2$ and $r_{T+\lambda-1}^f = r_{T+\lambda-1}^f + \lambda(c-\varepsilon) - \sum_{t=T}^{T+\lambda-1} r_t^f$. This new sequence satisfies (24), strictly

for $T, \dots, T+\lambda-2$ and with equality at $T+\lambda-1$, and $(1-\delta) \sum_{t=T}^{T+\lambda-1} \delta^{t-T} r_t^f \geq (1-\delta) \sum_{t=T}^{T+\lambda-1} \delta^{t-T} r_t^f$.

Consider finally the constant sequence $(c-\varepsilon)$ over $t = T, \dots, T+\lambda-1$. Given (24) it follows that

$$(1-\delta) \sum_{t=T}^{T+\lambda-1} \delta^{t-T} r_t^f \geq \left[(1-\delta) \sum_{t=T}^{T+\lambda-1} \delta^{t-T} \right] (c-\varepsilon). \text{ By (24), } r_T^f > (c-\varepsilon). \text{ Hence taking that first}$$

return to be exactly $(c-\varepsilon)$ and "shifting" returns to $T+1$ can only decrease the discounted sum. Since $\frac{1}{2} \sum_{t=T}^{T+1} r_t^f > c - \varepsilon$ the second return is now greater than $(c-\varepsilon)$. So discounted

sums are only lowered by taking the returns at $T+2$ to be $(c-\varepsilon)$ and "shifting" some returns to $T+2$. The argument iterates to yield the claimed inequality. Then,

$$(1-\delta) \sum_{t=T}^{T+\lambda-1} \delta^{t-T} r_t^f \geq (c - \epsilon) (1 - \delta^\lambda) \quad (27)$$

From (26) and (27) the claim is seen to follow easily for sufficiently high δ .

Inequality (25) implies that there is at least one state in the distribution induced by policy $\langle f \rangle$ at T , say s' , s.t.

$$(1 - \delta) V(s', \delta) \geq c - 2\epsilon \quad (28)$$

From value boundedness it then follows that

$$(1 - \delta) M + (1 - \delta) V(z, \delta) \geq (1 - \delta) V(s', \delta) \quad (29)$$

(29) immediately yields $v \geq c - 2\epsilon$. Since this holds for all $\epsilon > 0$ and all policies $\langle f \rangle$, we get $v \geq v^*$. Lemma 9 has been therefore proved. •

We have so far shown that v is the long-run average value for the dynamic programming problem. Given value boundedness it is easy to see that $(1 - \delta)V(z, \delta)$ has the same limit as $(1 - \delta)V(s, \delta)$ for any state s . So Theorem 3i) has been completely proved. Further, the stationary policy h^∞ in Lemma 7 is clearly optimal. Finally, the argument of Lemmas 7 and 8 yield policy convergence, i.e. (iii) of the theorem. •

10.3 Alternative Undiscounted Criteria

Proof of Theorem 4: (i) To save on notation let us write F_T for $\sum_{t=0}^T E^{\langle f \rangle} [r(s_t, a_t) | s_0]$ etc.

From the catching-up optimality of $\langle f \rangle$, for all $\epsilon > 0$ there is $T(\epsilon)$ s.t

$$F_T - G_T \geq -\epsilon, \quad T \geq T(\epsilon)$$

$$\text{So, } \lim_{T \rightarrow \infty} [F_T - vT] \geq \lim_{T \rightarrow \infty} [G_T - vT]$$

ii) Suppose that $\langle f \rangle$ is strongly long-run average optimal. Then, $\lim [F_T - vT] > -\infty$. Hence, $\lim \frac{F_T}{T} = v$.

iii) Suppose $\langle f \rangle$ is strongly long-run average optimal and that $V_* = V^*$. Then, $\lim F_T - vT = V_* = V^* \geq \lim \overline{G_T - vT}$. But that implies $\langle f \rangle$ catches up with $\langle g \rangle$. •

10.4 Convergence to Catching-Up and Strong Long-Run Average

Proof of Theorem 5: An alternative way of writing (23) is

$$\hat{\psi}(s) \leq E^h \sum_{t=0}^{T-1} [r_t - v] + E_T^h \hat{\psi} \quad (30)$$

It is clear that h^∞ is a good policy. By (A3), (30) immediately implies that $\hat{\psi} \leq V_*$.

Suppose in fact that $\hat{\psi}(s) < V_*(s)$, for some s . This implies in particular that $\lim_{\delta \uparrow 1} \bar{W}(s, \delta) < V_*(s)$. But then, there is a policy whose discounted returns are strictly greater than $W(s, \delta)$, for high δ . That is clearly a contradiction and hence we have established $\hat{\psi} = V_*$. (ii) and (iii) of the Theorem follow for the same reason as in Theorem 3. •

10.5 Applications

Proof of Proposition 1: Wlog, let $r \geq 0$. Let $\bar{r} = \max r$. Consider $s = 0$. Let $s' > 0$ and a in A be such that $q([s', 1] | 0, a) > \epsilon$, for some action a and some $\epsilon > 0$. Then,

$$V(0, \delta) \geq \delta \{ (1 - \epsilon)V(0, \delta) + \epsilon V(s', \delta) \}$$

Re-arranging,

$$[V(0, \delta) - V(s', \delta)](1 - \delta(1 - \epsilon)) \geq (\delta - 1)V(s', \delta) \geq -\bar{r}$$

This implies

$$V(0, \delta) - V(s', \delta) \geq \frac{-\bar{r}}{\epsilon} \quad (31)$$

Combining (31) with concavity of the value functions, we clearly have the existence of $M(s) > -\infty$, such that $V(s, \delta) - V(1, \delta) \geq M(s)$. The uniform upper bound for this difference is clearly zero. •

Proof of Corollary 3: The proof is identical to that of Proposition 1 except for the replacement of $s=0$ by s in $(0, s^*]$. •

Proof of Proposition 2: Consider $s=1$. Using the fact that the transition probabilities have no mass points, together with their weak continuity, it is not difficult to show that: for all

small $\epsilon > 0$ there is $s(\epsilon)$ in $(0,1)$ s.t. $q([0,s(\epsilon)] 1,a) > \epsilon$ for all a in A . Fix such an ϵ . Pick an arbitrary a and denote $q([0,s(\epsilon)] 0,a) = p > 0$.

$$\begin{aligned} V(0,\delta) &\geq \delta\{pV(0,\delta) + (1-p)V(s(\epsilon),\delta)\} \\ V(1,\delta) &\leq \bar{r} + \delta\{(1-\epsilon)V(1,\delta) + \epsilon V(s(\epsilon),\delta)\} \end{aligned}$$

Collecting these inequalities, some manipulation yields,

$$V(1,\delta) - V(0,\delta) \leq \delta(1-\delta)V(s(\epsilon),\delta)k$$

where k is a constant that depends on p and ϵ . Value boundedness immediately follows. •

Proof of Proposition 3: Let the production and utility functions be denoted f and u . By Proposition 1, the aggregative growth model satisfies value boundedness. Let z be the golden-rule of the model, i.e. z is generated by an investment x such that $f'(x)=1$. The maximum sustainable consumption is $f(x)-x$, i.e. the golden rule is the long-run average of this model. Clearly, $W(z,\delta) \rightarrow 0$, as $\delta \rightarrow 1$. Further, all good policies (Gale [18]) converge to the golden rule. Hence, the asymptotic zero integral property is also satisfied. For the same reason, $V^* = V_*$. So, any discounted policy and value limit, if they exist, are strong long-run average (and catching-up) optimal. But in the growth model optimal policies show monotone usage. Hence, by an identical argument to Corollary 1, policy limits exist (and so obviously do value limits). Uniqueness follows from the concavity of the production function and strict concavity of the utility function. •

10.6 Example

Proof of Lemma : A contradiction implies that for all $\epsilon > 0$ there is $N < \infty$ s.t.

$$\begin{aligned} \sum_{t=1}^T [k-r_t] &< -\epsilon, \quad T > N \\ \text{or } k &< \frac{1}{T} \sum_{t=1}^T r_t - \frac{\epsilon}{T} \\ \Rightarrow k &\leq \liminf \frac{1}{T} \sum_{t=1}^T r_t \end{aligned} \tag{32}$$

(32) contradicts (11). ♦

References:

1. R. Barro, A Theory of Monopolistic Price Adjustment, *Rev. Econ. Stud.* 34 (1972), 17-26.
2. W.J. Baumol, The Transactions Demand for Cash: An Inventory Theoretic Approach, *Quart. J. Econ.* 66 (1952), 545-556.
3. C. Berge, "Topological Spaces," Macmillan, New York, 1963.
4. R.N. Bhattacharya and M. Majumdar, Controlled Semi-Markov Models Under Long-Run Average Rewards, Cornell University Working Paper # 403, 1988.
5. P. Billingsley, "Convergence of Probability Measures," John Wiley, New York, 1968.
6. P. Billingsley, "Probability and Measure," John Wiley, New York, 1986.
7. V. Borkar, Control of Markov Chains with Long-Run Average Cost Criterion, in W. Fleming and P.L. Lions eds., "Stochastic Differential Systems, Stochastic Control Theory and Applications," Springer-Verlag, New York, 1988.
8. W. Brock, On Existence of Weakly Maximal Programmes in a Multi-Sector Economy, *Rev. Econ. Stud.* 37 (1970), 275-280.
9. W. Brock and L. Mirman, Optimal Economic Growth and Uncertainty: The Discounted Case, *J. Econ. Theory* 4 (1972), 479-513.
10. K. Burdett and D. Mortensen, Search, Layoffs and Labor Market Equilibrium, *J. Polit. Econ.* 41 (1980), 652-672.
11. W. Dechert and K. Nishimura, A Complete Characterisation of Optimal Growth in An Aggregated Model with a Non-Concave Production Function, *J. Econ. Theory* 31 (1983), 332-352.
12. P.K. Dutta, Capital Deepening and Impatience Equivalence in a Stochastic Aggregative Growth Model, *J.Econ. Dynam. Control* 11 (1987), 519-530.
13. P.K. Dutta and M. Majumdar, Limit Integration Theorems for Monotone Functions with Applications to Dynamic Programming, University of Rochester Working Paper forthcoming, 1990.
14. P.K. Dutta and R. Sundaram, Stochastic Games of Resource Allocation: Existence Theorems for Discounted and Undiscounted Models, Columbia University Working Paper # 483, 1990.
15. J. Flynn, Averaging vs. Discounting in Dynamic Programming: A Counter-Example , *Ann. Math. Stat.* 2 (1974), 411-413.
16. D. Fudenberg and E. Maskin, The Folk Theorem in Repeated Games With Discounting and with Incomplete Information, *Econometrica* 54 (1986), 533-554.
17. N. Furukawa, Markovian Decision Processes with Compact Action Spaces, *Ann. Math. Stat.* 5 (1972), 1612-1623.

18. D. Gale, On Optimal Development in a Multi-Sector Economy, *Rev. Econ. Stud.* 34 (1967), 1-18.
19. G. Goffman, "Real Functions," Prindle, Weber and Schmidt, Boston, Mass., 1953.
20. P. Jeanjean, Optimal Development Programmes under Uncertainty: The Undiscounted Case, *J. Econ. Theory* 7 (1974), 66-92.
21. S. Karlin, Steady State Solutions, in K.J. Arrow, S.Karlin and H.Scarf eds., *Studies in the Mathematical Theory of Inventory and Production*, Stanford University Press, Stanford, California, 1958.
22. E. Lehrer and D. Monderer, Discounting versus Averaging in Dynamic Programming, mimeo, Northwestern University, 1989 .
23. A. Maitra, Discounted Dynamic Programming on Compact Metric Spaces, *Sankhya Ser. A* 30 (1968), 211-216.
24. T. Mitra, On Optimal Depletion of Exhaustible Resources: Existence and Characterization Results, *Econometrica* 48 (1980),1431-1450.
25. D. Mortensen, Job Search and Labor Market Analysis, Northwestern University Working Paper # 594, 1984.
26. A.C. Pigou, "The Economics of Welfare," Macmillan, London, 1920.
27. R. Radner, Optimal Growth in a Linear Logarithmic Economy, *Int. Econ. Rev.* 7 (1966), 1-35.
28. R. Radner, Monitoring Cooperative Agreements in a Repeated Principal Agent Relationship, *Econometrica* 49 (1981), 1127-1148.
29. F.P. Ramsey, A Mathematical Theory of Savings, *Econ. J.* 38 (1928), 543-559.
30. J. Rawls, "A Theory of Justice", Belknap Press, Harvard, Cambridge, Massachusetts, 1971.
31. W. Reed, A Stochastic Model for the Economic Management of a Renewable Animal Resource, *Math. Biosci.* 22 (1974), 313-337.
32. S.M. Ross, Non-discounted Denumerable Markovian Decision Models, *Ann. Math. Stat.* 39 (1968a), 412-423.
33. S.M. Ross, Arbitrary State Markovian Decision Processes, *Ann. Math. Stat.* 39 (1968b), 2118-2122.
34. S.M. Ross, "Introduction to Stochastic Dynamic Programming," Academic Press, New York, 1983.
35. A. Rubinstein, Equilibrium in Supergames with the Overtaking Criterion, *J. Econ. Theory* 21 (1979), 1-9.

36. N. Stokey and R. Lucas, "Recursive Methods in Economic Dynamics," Harvard University Press, Cambridge, Massachusetts, 1989.
37. H. Taylor, Markovian Sequential Replacement Processes, *Ann. Math. Stat.* 36 (1965), 1677-1694.
38. J. Tobin, The Interest Elasticity of the Transactions Demand for Cash, *Rev. Econ. Stat.* 38 (1956), 241-247.