

EXISTENCE AND DISCOVERY OF AVERAGE OPTIMAL SOLUTIONS IN DETERMINISTIC INFINITE HORIZON OPTIMIZATION

IRWIN E. SCHOCHETMAN AND ROBERT L. SMITH

Department of Mathematical Sciences
Oakland University
Rochester, Michigan 48309
and
Department of Industrial and Operations Engineering
The University of Michigan
Ann Arbor, Michigan 48109

June 24, 1997

ABSTRACT. We consider the problem of making a sequence of decisions, each chosen from a finite action set over an infinite horizon, so as to minimize its associated *average* cost. Both the feasibility and cost of a decision are allowed to depend upon all of the decisions made prior to that decision; moreover, time-varying costs and constraints are allowed. A feasible solution is said to be *efficient* if it reaches each of the states through which it passes at minimum cost. We show that efficient solutions exist and that, under a *state reachability* condition, efficient solutions are also average optimal. Exploiting the characterization of efficiency via a solution's short-run as opposed to long-run behavior, a forward algorithm is constructed which recursively discovers the first, second, and subsequent decisions of an efficient, and hence average optimal, infinite horizon solution.

1. Introduction

Infinite horizon optimization at its most fundamental level is the problem of selecting an infinite sequence of decisions which promises to minimize the associated costs incurred over an unbounded horizon (Bean and Smith [1984], Schochetman and Smith [1989]). One of the key difficulties inherent in this task is the dilemma of evaluating an infinite stream of costs whose cumulative value will typically diverge.

Perhaps the most commonly employed criterion is that of *discounted* costs. This criterion explicitly reflects the time value of money via a discount factor which leads to convergence in cost. Perhaps the second most commonly used criterion is obtained by replacing the original cost stream by its average value per period, the so-called *average-cost* criterion. Advantages of the latter criterion include: a) the value of a discount factor need not be specified; b) it is a numerically stable substitute for discounted problems with discount factor near 1; c) when costs are not measured in dollars, use of a discount factor becomes artificial and

1991 *Mathematics Subject Classification*. Primary 90C20, Secondary 49A99.

Key words and phrases. Average optimality, infinite horizon optimization, efficient solution, state reachability, forward algorithm, equipment replacement.

This work was supported in part by the National Science Foundation under Grant DDM-9214894

potentially meaningless; and d) cumulative costs need not converge. Despite these advantages, the average-cost criterion introduces a number of mathematical pathologies not shared by the discounted-cost criterion. Perhaps the most serious is that, since the average value of an infinite horizon cost stream is insensitive to the costs incurred over any finite horizon, the criterion is extraordinarily underselective. It includes average-cost optimal strategies that are difficult to justify in any other respect (Hopp, Bean, and Smith [1987]). This behavior is correctable in the stationary case by restricting consideration to a set of strategies within which average-cost minimization makes sense. In particular, in stationary problems where neither decision feasibility nor costs are time-dependent, restriction to *stationary* strategies yields an average-cost optimal strategy that avoids the myopic behavior of the larger class of all average-cost optimal strategies (Ross [1983], Puterman [1994]).

In this paper, we extend the latter approach to *nonstationary* problems where either feasibility or costs (possibly both) are time dependent. However, stationary strategies need no longer be optimal here. Instead, we generalize by restricting consideration to *efficient* strategies (see Ryan, Bean, and Smith [1992], Schochetman and Smith [1992] for a similar concept in the discounted case). A strategy is efficient if it is average-cost (or equivalently undiscounted total cost) optimal to each of the states through which it passes. We show that an efficient solution *always* exists and that, under a state-reachability condition, *every* efficient solution is average-cost optimal. This result is perhaps surprising in that the mixing effects due to ergodicity in the more traditional stochastic setting (Ross [1968], Federgruen and Tijms [1978], Hopp, Bean and Smith [1987], Bean, Lasserre, and Smith [1990], Park, Bean, and Smith [1993]) are absent in our purely deterministic model. As a consequence of the above, we conclude existence of an average-cost optimal solution. Moreover, because of its characterization in terms of its optimal behavior *to* states, as opposed to *from* states, we construct a *forward* algorithm that is guaranteed to successively discover the first, second, and subsequent decisions along an efficient, and hence average-cost optimal, strategy for general nonstationary infinite horizon optimization. The algorithm is similar to the forward algorithms proposed in Schochetman and Smith [1992] in the discounted case.

In section 2, we formally introduce the general class of discrete-time infinite horizon problems which we study. We form their finite horizon approximating problems by projection of the feasible region and objective function onto finite dimensional spaces. Efficient solutions are then introduced in section 3 as feasible infinite horizon solutions which are optimal to each of the states through which they pass. The existence of efficient solutions is demonstrated by a topological compactness argument. In section 4, we introduce a state-reachability property which guarantees that every efficient solution is average-cost optimal, and in particular, that an average-cost optimal solution thus exists. In section 5, we turn to the problem of approximating efficient solutions by their computable finite horizon counterparts. We first show that the sequence of sets of *all* efficient solutions of the finite horizon approximating problems converges, in the sense of Kuratowski, to the set of efficient solutions of the infinite horizon problem. We next show *policy* convergence, where the sequence of lexicomin elements of the sets of finite horizon efficient solutions is shown to converge to the lexicomin infinite horizon efficient solution. This is an instance of selection convergence. Average optimal cost convergence is established in section 6. In section 7, a finite forward algorithm, with stopping rule, is given that recursively recovers the first, second, and subsequent decisions in an average-cost optimal efficient solution, whenever the latter is *unique*.

For additional examples and omitted proofs, we refer the reader to Schochetman and Smith [1997].

2. Problem Formulation

We begin with an extremely general deterministic infinite horizon optimization problem, formulated as a dynamic programming problem. The deterministic property removes our model from the stochastic framework required for the majority of average-cost criterion work (Ross [1968], Derman [1966], Federgruen and

Tijms [1978]). These stochastic models generally require a single fixed recurrent class of the Markov chain corresponding to every policy, a property that fails in the purely deterministic case. Essentially, the only restriction in this work, apart from being a deterministic model, is the requirement that at most a finite number of decision alternatives is available at each decision epoch. Included are production planning under nonstationary demand, parallel and serial equipment replacement under technological change, capacity planning under nonlinear demand, and optimal search in a time-varying environment.

For convenience, we embed the problem within a discrete-time framework, where each decision is made at the beginning of each of a series of equal time periods, indexed by $j = 1, 2, \dots$. The set of all possible decisions available in period j (irrespective of the period's beginning state) is denoted by Y_j , and is assumed to be finite. For convenience we let $Y_j = \{0, 1, 2, \dots, m_j\}$ with the usual discrete topology, with $m_j \geq 1$, $\forall j = 1, 2, \dots$ (we may interpret decision 0 as signaling no action for that period). We model the problem as a dynamic system governed by the state equation

$$s_j = f_j(s_{j-1}, y_j), \quad \forall j = 1, 2, \dots,$$

where s_0 is the given initial state of the system (beginning period 1); s_j is the state of the system at the end of period j , i.e., beginning period $j + 1$; y_j is the control or decision selected in period j with knowledge of the beginning state s_{j-1} ; and S_j is the given (finite) set of feasible states ending period j , so that $s_j \in S_j$. $\forall j = 1, 2, \dots$. The set $Y_j(s_{j-1})$ is the given (finite) non-empty subset of decisions available in period j when the system is in state $s_{j-1} \in S_{j-1}$, so that $y_j \in Y_j(s_{j-1}) \subseteq Y_j$. Finally, f_j is the given state transition function in period j , where $f_j : F_j \rightarrow S_j$, with

$$F_j = \{(s_{j-1}, y_j) \in S_{j-1} \times Y_j : y_j \in Y_j(s_{j-1})\}, \quad \forall j = 1, 2, \dots$$

We set $S_0 = \{s_0\}$, and require that

$$S_j = \{f_j(s_{j-1}, y_j) : s_{j-1} \in S_{j-1}, y_j \in Y_j(s_{j-1})\}, \quad \forall j = 1, 2, \dots,$$

so that, in particular, $S_1 = \{f_1(s_0, y_1) : y_1 \in Y_1(s_0)\}$. Thus, each S_j is exactly the set of feasible, i.e., attainable, states in period j . Note that our state space formulation is essentially without loss of generality since we can if necessary identify distinct states to every feasible sequence of preceding decisions. We will at times adopt the equivalent view that models the decision problem via a directed acyclic network where the nodes correspond to the states $s_j, j = 1, 2, \dots$ and a directed arc joining state s_j to state s_{j+1} corresponds to an action in $Y_{j+1}(s_j)$ that transforms state s_j into state s_{j+1} via the state transition function f . Every directed path from node s_0 to node s_j then corresponds to a feasible sequence of decisions that results in state s_j at the end of period j .

The product set $Y = \prod_{j=1}^{\infty} Y_j$ which contains all feasible decision sequences or strategies is then a compact topological space relative to the product topology, i.e., the topology of componentwise convergence. (Note that in general not all strategies in Y will be feasible.) Because of the finiteness of the Y_j , componentwise convergence yields eventual agreement in each component of the limiting strategy.

Lemma 2.1. *Let $y \in Y$, $\{y^n\} \subseteq Y$, such that $y^n \rightarrow y$, as $n \rightarrow \infty$. Let $1 \leq K < \infty$. Then there exists n_K sufficiently large such that $y_j^n = y_j$ in $Y_j, \forall j = 1, \dots, K$, for all $n \geq n_K$.*

In section 4, we will find it necessary to assume that the sequence $\{m_j\}$ is bounded by some $1 \leq m < \infty$, i.e., $m_j \leq m$, all j . In this event, the product topology on Y is metrizable with metric d given by

$$d(x, y) = \sum_{j=1}^{\infty} \beta^j |x_j - y_j|, \quad \forall x, y \in Y,$$

where we choose β so that $0 < \beta < 1/(m+1)$.

We will let $\mathcal{K}(Y)$ denote the set of all closed (hence, compact), non-empty subsets of Y . Equipped with the Hausdorff metric D (Berge [1963]) derived from the above metric d , $\mathcal{K}(Y)$ is also a compact metric space (Hausdorff [1962]). It is well-known that the resulting convergence in $\mathcal{K}(Y)$ is set convergence in the sense of Kuratowski (Kuratowski [1966], Hausdorff [1962]). Recall that a sequence of sets T_n in an abstract metric space (Z, ρ) is said to *Kuratowski converge* to T , written $\lim T_n = T$, if $T = \limsup T_n = \liminf T_n$, where:

- (i) $\liminf T_n$ = the set of points x in Z for which there exists $x_n \in T_n$, for n sufficiently large, such that $\rho(x_n, x) \rightarrow 0$, as $n \rightarrow \infty$.
- (ii) $\limsup T_n$ = the set of points x in Z for which there exists a subsequence $\{T_{n_j}\}$ of $\{T_n\}$, and a corresponding sequence $\{x_j\}$, such that $x_j \in T_{n_j}$, $\forall j$, and $\rho(x_j, x) \rightarrow 0$, as $j \rightarrow \infty$.

Now let $x \in Y$. Define x to be *feasible* if $x_j \in Y_j(s_{j-1})$, where $s_j = f_j(s_{j-1}, x_j)$, for all $j = 1, 2, \dots$. We define the *feasible region* X to be the subset of Y consisting of all those x which are feasible.

Lemma 2.2. *The subset X of all feasible solutions is non-empty and closed in Y , i.e., $X \in \mathcal{K}(Y)$.*

Proof. The non-emptiness of X follows immediately from our assumption that $Y_j(s_{j-1}) \neq \emptyset$, for all $s_{j-1} \in S_{j-1}$, and for all $j = 1, 2, \dots$. To show X is closed in the metric space Y , let $\{x^n\}$ be a sequence in X and y an element of Y for which $x^n \rightarrow y$ in Y , i.e., $x_j^n \rightarrow y_j$, as $n \rightarrow \infty$, $\forall j = 1, 2, \dots$. We show that $y \in X$, i.e., $y_j \in Y_j(s_{j-1})$, where $s_j = f_j(s_{j-1}, y_j)$, $\forall j = 1, 2, \dots$.

By Lemma 2.1, we have that for each j , eventually $x_j^n = y_j$. Let $1 \leq K < \infty$. Then there exists m sufficiently large such that $x_j^m = y_j$, $\forall j = 1, 2, \dots, K$. Since $x^m \in X$ by hypothesis, we have that $x_j^m \in Y_j(s_{j-1}^m)$, where $s_j^m = f_j(s_{j-1}^m, x_j^m)$, for all $j = 1, 2, \dots$. Then $y_1 = x_1^m \in Y_1(s_0)$ and $f_1(s_0, y_1) = f_1(s_0, x_1^m) = s_1^m$, so that $s_1 = s_1^m$. Similarly, $y_2 = x_2^m \in Y_2(s_1^m) = Y_2(s_1)$ and $f_2(s_1, y_2) = f_2(s_1, x_2^m) = s_2^m$, so that $s_2 = s_2^m$. Continuing in this manner, we conclude that $y_K = x_K^m \in Y_K(s_{K-1}^m) = Y_K(s_{K-1})$ and $f_K(s_{K-1}, y_K) = f_K(s_{K-1}, x_K^m) = s_K^m$, so that $s_K = s_K^m$. Since K is arbitrary, it follows that $x \in X$. \square

Clearly, the infeasibility of an *infinite* strategy is determinable from observing a *finite* sequence of its initial decisions. Practically speaking, this is hardly a restriction since the feasibility of a potential decision in period j is allowed to depend upon the entire sequence of decisions previously made.

Suppose $y \in Y$ and N is a positive integer. Suppose y is such that $y_j \in Y_j(s_{j-1})$, where $s_j = f_j(s_{j-1}, y_j)$, for $j = 1, 2, \dots, N$, i.e., y is feasible through period N . In this event, for each such $j \leq N$, we define $s_j(y)$ to be the state $s_j = f_j(s_{j-1}, y_j)$, which is an element of S_j . If we do this for each $j = 1, 2, \dots, N$, then notationally we have that $s_j(y) = f_j(s_{j-1}(y), y_j)$. (Note that the state $s_j(y)$ is well-defined.) We will refer to each such $s_j(y)$ as the state through which y passes at the end of period j . If both $z, y \in Y$ satisfy the previous property with respect to N , and $y_j = z_j$, $\forall j = 1, 2, \dots, N$, then $s_j(y) = s_j(z)$ for all $j = 1, 2, \dots, N$. Moreover, if $x \in X$, then the previous property is satisfied for x and *all* positive integers N , so that $s_j(x)$ is defined for *each* period $j = 1, 2, \dots$ in this case. Finally, if $x \in X$, then $(s_{j-1}(x), x_j) \in F_j$, $\forall j = 1, 2, \dots$.

Turning to the objective function, we also allow the cost of a decision made in period j to depend upon the sequence of previous decisions, or more accurately, upon the state resulting from these decisions. Specifically, we let $c_j(s_{j-1}, y_j)$ be the real-valued cost of decision y_j in period j , if we are in state s_{j-1} beginning period j . We thus obtain cost functions $c_j : F_j \rightarrow \mathbb{R}$, which we assume are uniformly bounded, i.e., there exists $0 < B < \infty$ such that

$$|c_j(s_{j-1}, y_j)| \leq B, \quad \forall s_{j-1} \in S_{j-1}, \quad \forall y_j \in Y_j(s_{j-1}), \quad \forall j = 1, 2, \dots$$

We will adopt the average-cost optimality criterion in this paper. Specifically, for each strategy $x \in X$,

we define the associated average-cost $A(x)$ by

$$A(x) = \limsup_N \frac{1}{N} \sum_{j=1}^N c_j(s_{j-1}(x), x_j),$$

so that $-B \leq A(x) \leq B, \forall x \in X$. One of the pathological aspects of the average cost criterion is its failure to be continuous.

Our average-cost optimization problem (P) can now be formulated as:

$$\min_{x \in X} A(x) \tag{P}$$

We will denote the set of optimal solutions to (P) by X^* . Of course, $\phi \subseteq X^* \subseteq X$, in general. Incidentally, another pathological aspect of the average-cost criterion is that X^* need be not closed in Y , since the optimality of a strategy x is unaffected by costs incurred over early periods.

We remark that the discounted-cost criterion is formally included in the average-cost case. Let $x \in X$. Consider the objective function of the discounted-cost criterion $D(x) = \sum_{j=1}^{\infty} c_j(s_{j-1}(x), x_j)\alpha^j$, where $0 < \alpha < 1$. Set

$$c'_j(s_{j-1}(x), x_j) = j \sum_{l=1}^j c_l(s_{l-1}(x), x_l)\alpha^l - \sum_{l=1}^{j-1} c'_l(s_{l-1}(x), x_l), \quad \forall j = 2, 3, \dots,$$

where $c'_1(s_0, x_1) = c_1(s_0, x_1)$. Then $\frac{1}{N} \sum_{j=1}^N c'_j(s_{j-1}(x), x_j) = \sum_{l=1}^N c_l(s_{l-1}(x), x_l)\alpha^l$, so that

$$A(x) = \limsup_N \frac{1}{N} \sum_{j=1}^N c'_j(s_{j-1}(x), x_j) = \sum_{l=1}^{\infty} c_l(s_{l-1}(x), x_l)\alpha^l = D(x).$$

In this way, we may transform any discounted-cost problem into an equivalent average-cost problem.

Our primary objective in this paper is to approximate optimal solutions of (P) by solutions of finite horizon truncations of (P) . To this end, for each $N = 1, 2, \dots$, let K_N be the (finite) projection of the feasible region X onto $Y_1 \times \dots \times Y_N$, i.e., K_N is equal to the set of $(x_1, \dots, x_N) \in Y_1 \times \dots \times Y_N$ which satisfy: $x_j \in Y_j(s_{j-1})$, where $s_j = f_j(s_{j-1}, x_j), \forall j = 1, 2, \dots, N$.

The set K_N denotes a set of N -horizon feasible strategies in the (finite) set $Y_1 \times \dots \times Y_N$ of all N -horizon strategies. Note that every element of K_N can be *feasibly* completed to an element of X . Moreover, the first N decisions of every element of X lie within K_N . In fact, a subset in \mathbb{R}^N has these two properties if and only if it is the projection of X onto \mathbb{R}^N . It is for this reason that we believe K_N is the natural choice for the feasible region of the N -th approximating subproblem.

Now, for technical reasons, we embed K_N into Y by letting X_N denote the set of all arbitrary extensions of the elements of K_N , i.e., $X_N = K_N \times Y_{N+1} \times Y_{N+2} \times \dots$, so that $X_N \in \mathcal{K}(Y)$, all N . We shall effectively identify X_N and K_N . Note that if $x \in X_N$, then $(s_N(x), x_N) \in F_N$. Moreover, it is clear that $K_{N+1} \subseteq K_N \times Y_{N+1}$, so that

$$X_{N+1} \subseteq X_N, \quad \forall N = 1, 2, \dots,$$

i.e., the sets X_N are nested downward. Moreover, their Kuratowski limit exists and is the infinite horizon feasible region X , as the following lemma demonstrates.

Lemma 2.3. *We have $X = \bigcap_{N=1}^{\infty} X_N$. Thus,*

$$\lim_{N \rightarrow \infty} X_N = X.$$

Proof. Clearly, $X \subseteq X_N$, all N , so that $X \subseteq \bigcap_{N=1}^{\infty} X_N$. Conversely, let $x \in \bigcap_{N=1}^{\infty} X_N$ and suppose $x \notin X$. Then $\delta = d(x, X) > 0$, where as usual, $\delta = \inf\{d(x, y) : y \in X\}$. Let N_0 be sufficiently large such that $\sum_{i=N_0+1}^{\infty} 2^{-i} < \delta$ and $y \in X$. If $x_i = y_i$, $1 \leq i \leq N_0$, then

$$d(x, y) \leq \sum_{i=N_0+1}^{\infty} \frac{1}{2^i} < \delta,$$

which is a contradiction. Thus, for each $y \in X$, there exists $1 \leq i_y \leq N_0$ for which $x_{i_y} \neq y_{i_y}$. This implies that $x \notin X_{N_0}$, which is a contradiction. For the second part, see Aubin [1990]. \square

We define the N -horizon average-cost optimization problem (P_N) as follows:

$$\min_{x \in X_N} A_N(x), \tag{P_N}$$

where

$$A_N(x) = \frac{1}{N} \sum_{j=1}^N c_j(s_{j-1}(x), x_j), \quad \forall x \in X_N,$$

is the average cost of strategy x over horizon N , $\forall N = 1, 2, \dots$. It will be convenient to write $C_N(x) = NA_N(x)$, i.e.,

$$C_N(x) = \sum_{j=1}^N c_j(s_{j-1}(x), x_j), \quad \forall x \in X_N.$$

Clearly, each $A_N : X_N \rightarrow \mathbb{R}$ is a continuous function, since A_N depends only on K_N , which is finite and discrete. Similarly for C_N . Note that if $M > N$, then $C_M(x) - C_N(x) \leq (M - N)B$.

Problem (P_N) is really an N -horizon problem since only the first N entries of x matter, i.e., if $x \in X_N$ is (P_N) -optimal, then so is $(x_1, \dots, x_N, y_{N+1}, y_{N+2}, \dots)$, where $y_j \in Y_j$ can be chosen arbitrarily $\forall j \geq N + 1$. We have defined (P_N) in this way so that solutions to (P_N) are in Y , for all N , i.e., they are comparable to each other and to the solutions of (P) as well. Hence, (P_N) is effectively a finite problem which is finitely solvable (in the worst case) by enumeration of the elements of K_N .

The optimal solution set X_N^* to (P_N) is then a non-empty, closed subset of Y , since K_N is finite and non-empty. Thus, $X_N^* \in \mathcal{K}(Y)$, for all N . Although the X_N are nested downward, the X_N^* need not be. We will denote the optimal objective value to (P_N) by A_N^* , $\forall N = 1, 2, \dots$.

Now, for each $N = 1, 2, \dots$ and $s \in S_N$, define

$$X_N(s) = \{x \in X_N : s_N(x) = s\},$$

so that $\{X_N(s) : s \in S_N\}$ is a partition of X_N . If $x \in X_N$, then $s_N(x)$ is the unique element of S_N for which $x \in X_N(s_N(x))$, $\forall N = 1, 2, \dots$.

The following lemma observes that the state at time N of a strategy depends only upon the first N decisions associated with that strategy.

Lemma 2.4. *Let $N = 1, 2, \dots$, and $s \in S_N$. If $x \in X_N(s)$, then*

$$(x_1, \dots, x_N, y_{N+1}, y_{N+2}, \dots) \in X_N(s),$$

where $y_j \in Y_j$ is arbitrary for all $j = N + 1, \dots$.

Proof. This follows from the definition of $X_N(s)$. □

The key property that formally endows states with the intuitive notion that they should incorporate all of the information about past decisions relevant to making future decisions is formalized by the following lemma. Roughly speaking, it says that if two finite horizon feasible strategies with different horizons have the same state at the earlier horizon, followed by identical decisions to the later horizon, then a) the earlier strategy is feasible to the later horizon, with the same ending state as the later strategy, and b) the costs-per-period are the same from the earlier horizon to the later horizon.

Lemma 2.5. *If $0 \leq N < M < \infty$, $x \in X_M$, $y \in X_N$, $s_N(x) = s_N(y)$ and $(x_{N+1}, \dots, x_M) = (y_{N+1}, \dots, y_M)$, then*

- a) $y \in X_M$, $s_M(y) = s_M(x)$ and
- b) $c_j(s_{j-1}(x), x_j) = c_j(s_{j-1}(y), y_j)$, $\forall j = N + 1, \dots, M$.

Proof. Left to the reader. □

We will repeatedly apply the previous lemma in different contexts. Thus, it will be convenient to refer to Lemma 2.5 as being applied to the context (N, M, y, x) . In particular, note that Lemma 2.4 is a consequence of the Lemma 2.5 applied to $(0, M, y, x)$, with $X_0 = Y$ and $s_0(x) = s_0(y) = s_0$.

3. Efficient Solutions

The state-space formulation above associated a unique state at each time period with every infinite horizon feasible solution. Solutions that have the property of *optimally* reaching each of the states through which they pass are called *efficient* solutions (Schochetman and Smith [1992], Ryan, Bean, and Smith [1992]). Such a solution offers little opportunity for retrospective regret in that at every state along its path, there was no better way to reach that state. This efficiency of movement through the state space suggests efficient solutions as candidates for average-cost optimality. This last observation will be explored in the next section, where efficient solutions will indeed be shown to be average-cost optimal under a state reachability property. In this section, we formalize the notion of an efficient solution and prove the existence of such solutions.

For each $1 \leq N < \infty$ and $s \in S_N$, consider the problem $(P_N(s))$ defined as follows:

$$\min_{x \in X_N(s)} A_N(x) \tag{P_N(s)}$$

Optimal solutions to $(P_N(s))$ consist of those N -horizon feasible strategies of least cost which have state s ending period N . As was the case for (P_N) , such optimal solutions exist and form a closed set in Y . Denote them by $X_N^*(s)$, so that $X_N^*(s) \in \mathcal{K}(Y)$, for all N and all $s \in S_N$.

Remark 3.1. Let $N = 1, 2, \dots$. If x is an N -horizon optimal solution, then it is optimal to its own state $s_N(x)$ ending period N , i.e., if $x \in X_N^*$, then $x \in X_N^*(s_N(x))$. This follows because $x \in X_N(s_N(x)) \subseteq X_N$. Hence, if x is least-cost over X_N , the same must be true over the smaller set $X_N(s_N(x))$.

For each $N = 1, 2, \dots$, define \mathcal{X}_N^* to be the set of all N -horizon feasible strategies which are optimal to some feasible state $s \in S_N$, i.e.,

$$\mathcal{X}_N^* = \bigcup_{s \in S_N} X_N^*(s).$$

Since this is a finite (disjoint) union of closed, non-empty subsets of Y , it follows that $\mathcal{X}_N^* \in \mathcal{K}(Y)$, for all N .

Remark 3.2. For each $N = 1, 2, \dots$, we have that $\mathcal{X}_N^* \supseteq X_N^*$ because if $x \in X_N^*$, then $x \in X_N^*(s)$, for the state $s = s_N(x)$ in S_N .

The following lemma is a version of the Principle of Optimality: any optimal solution (to its state) at time N must have been optimal to the state it passed through at each time $M < N$.

Lemma 3.3. *Let $N = 1, 2, \dots$ and $x \in X_N$. If x is optimal to horizon N (necessarily with ending state $s_N(x)$), then x is optimal to every earlier horizon M (with ending state $s_M(x)$), i.e., if $x \in X_N^*(s_N(x))$, then $x \in X_M^*(s_M(x))$, for all $1 \leq M \leq N$.*

Proof. Omitted. □

The following theorem is also a version of the Principle of Optimality: any optimal solution (to its state) at time $N + 1$ must have been optimal to the state it passed through at time N . Equivalently, the efficient solutions are nested downward. (Recall that the optimal solutions X_N^* need not be nested downward.)

Theorem 3.4. *For each $N = 1, 2, \dots$, we have $\mathcal{X}_{N+1}^* \subseteq \mathcal{X}_N^*$.*

Proof. Let $x \in \mathcal{X}_{N+1}^*$. Then $x \in X_N^*(s_{N+1}(x))$, where $s_{N+1}(x) \in S_{N+1}$. By Lemma 3.3, we have that $x \in X_N^*(s_N(x))$, where $s_N(x) \in S_N$, so that $x \in \mathcal{X}_N^*$. □

We can now formally define an efficient solution to be any solution in Y which is a member of \mathcal{X}_N^* , for each $N = 1, 2, \dots$. That is, the set \mathcal{X}^* of (infinite horizon) efficient solutions is given by

$$\mathcal{X}^* = \bigcap_{N=1}^{\infty} \mathcal{X}_N^*.$$

Efficient solutions may thus be characterized as feasible solutions that are optimal to each of the states through which they pass, i.e. $x^* \in X_N^*(s_N(x^*))$ for each $N = 1, 2, \dots$. Note that \mathcal{X}^* is a closed subset of Y , while X^* is *not* closed in general.

Lemma 3.5. *The set \mathcal{X}^* is non-empty, i.e., $\mathcal{X}^* \in \mathcal{K}(Y)$.*

Proof. Let $x_N^* \in \mathcal{X}_N^*$, all N . Then $\{x_N^*\}$ is a sequence in the compact metric space Y . Thus, there exists y in Y and a subsequence $\{x_{N_k}^*\}$ of $\{x_N^*\}$ such that $x_{N_k}^* \rightarrow y$ in Y , as $k \rightarrow \infty$.

Fix N . By Lemma 3.3, there exists k_N sufficiently large such that $x_{N_k}^* \in \mathcal{X}_N^*$, for all $k \geq k_N$. Since \mathcal{X}_N^* is closed, it follows that $y \in \mathcal{X}_N^*$. But N is arbitrary. Hence, $y \in \mathcal{X}^*$. □

By the previous lemma, efficient solutions always exist in our setting. We demonstrate in the next section that, under an additional hypothesis, they are necessarily average-cost optimal.

4. State-Reachability and Average-Cost Optimality

In this section, we show that, under a state-reachability property, efficient solutions are average-cost optimal. We then conclude the existence of an average-cost optimal solution under state-reachability, as a consequence of the existence of efficient solutions.

Recall that for each $N = 1, 2, \dots$, we have $X \subseteq X_N$, so that $\{s_N(x) : x \in X\} \subseteq S_N$. Conversely, since every finite horizon feasible strategy is extendable to a feasible infinite horizon strategy, we have that

$\{s_N(x) : x \in X\} \supseteq S_N$, so that $S_N = \{s_N(x) : x \in X\}$, $\forall N = 1, 2, \dots$, i.e., every N -horizon feasible state is realizable as the N -horizon state of some infinite horizon feasible solution. Thus, *there is no distinction between finite and infinite horizon feasible states.*

We now present a notion of *bounded state-reachability* which requires roughly that it be possible to feasibly reach any state within a bounded number of periods from any other feasible state.

Definition 4.1 (Bounded Reachability Property). *There exists a positive integer R such that for each $1 \leq K < \infty$, each $s \in S_K$ and each finite sequence of states (t_K, \dots, t_{K+R}) in $S_K \times \dots \times S_{K+R}$, there exists $K \leq L \leq K + R$ and $y \in X_L$ for which $s_K(y) = s$ and $s_L(y) = t(L)$. We will refer to this as the Bounded Reachability Property for $(K, s; (t_K, \dots, t_{K+R}))$. If there exists $R > 0$ for which this property holds for all $K, s; (t_K, \dots, t_{K+R})$, then problem (P) is said to satisfy Bounded Reachability.*

We now can prove the main result of this paper.

Theorem 4.2. *Suppose (P) satisfies Bounded Reachability. Then every efficient solution is average-cost optimal, i.e.,*

$$\mathcal{X}^* \subseteq X^*.$$

Proof. Let $v \in \mathcal{X}^*$. Since $\mathcal{X}^* = \liminf_N \mathcal{X}_N^*$, there exists a sequence $\{s_N\}$ and a corresponding sequence $\{x_N^*(s_N)\}$ such that $s_N \in S_N$, $x_N^*(s_N) \in X_N^*(s_N)$, all N , and $x_N^*(s_N) \rightarrow v$ in Y , as $N \rightarrow \infty$. Since $X_N^*(s_N) \subseteq X_N$, all N , it follows that $v \in \liminf_N X_N = X$, i.e., v is feasible for (P) . We next show that v is optimal for (P) , i.e., $A(v) \leq A(x)$, $\forall x \in X$, so that $v \in X^*$.

Let x be an arbitrary element of X . Fix $1 \leq K < \infty$. Let n_k be as in Lemma 2.1, B as in section 2 and R as in the Bounded Reachability Property. Fix T sufficiently large so that $T \geq n_k$ and $T \geq K + R$. Without loss of generality, we may assume that $n_k \geq K + R$. Set $N = T$ and $y = x^*(T, s_T)$ for convenience, so that $s_T(y) = s_T$ and y is optimal for $(P_T(s_T))$. Since $x^*(N, s_N) \rightarrow v$, as $N \rightarrow \infty$, and $T \geq n_k$, it follows from Lemma 2.1 that $y_j = v_j$, $j = 1, \dots, K$. Furthermore, since s_0 is the initial state for *all* strategies, we have that $s_K(v) = s_K(y)$ and

$$c_j(s_{j-1}(v), v_j) = c_j(s_{j-1}(y), y_j), \quad \forall j = 1, \dots, K,$$

by Lemma 2.5 applied to $(0, K, y, v)$ (or $(0, K, v, y)$).

Applying the Bounded Reachability Property to $(K, s_K(x); (s_K(y), \dots, s_{K+R}(y)))$, there exists $K \leq L \leq K + R$ and $z \in X_L$ such that $s_K(z) = s_K(x)$ and $s_L(z) = s_L(y)$. Observe that $z \in X_L(s_L(z)) = X_L(s_L(y))$ and $y \in X_T \subseteq X_L$.




Figure 1 - Every efficient solution is average-cost optimal

Now let

$$w = (x_1, \dots, x_K, z_{K+1}, \dots, z_L, y_{L+1}, y_{L+2}, \dots).$$

Then $w \in X_K(s_K(x))$ (Lemma 2.4), so that

$$s_K(w) = s_K(x) = s_K(z),$$

$$w \in X_L(s_L(z)) = X_L(s_L(y)),$$

$$s_{j-1}(w) = s_{j-1}(z),$$

$$s_L(w) = s_L(z) = s_L(y),$$

and

$$c_j(s_{j-1}(w), w_j) = c_j(s_{j-1}(z), z_j), \quad \forall j = K+1, \dots, L,$$

by Lemma 2.5 for (K, L, w, z) . Similarly, by Lemma 2.5 for (L, T, w, y) , we have that $w \in X_T(s_T) = X_T(s_T(y))$,

$$s_T(w) = s_T = s_T(y),$$

$$s_{j-1}(w) = s_{j-1}(y),$$

and

$$c_j(s_{j-1}(w), w_j) = c_j(s_{j-1}(y), y_j), \quad \forall j = L+1, \dots, T.$$

Since y is optimal for $(P_T(s_T))$, it follows that $C_T(w) \geq C_T(y)$. But

$$\begin{aligned} C_T(w) &= C_K(w) + [C_L(w) - C_K(w)] + [C_T(w) - C_L(w)] \\ &= C_K(x) + [C_L(z) - C_K(z)] + [C_T(y) - C_L(y)], \end{aligned}$$

and

$$\begin{aligned} C_T(y) &= C_K(y) + [C_L(y) - C_K(y)] + [C_T(y) - C_L(y)] \\ &= C_K(v) + [C_L(y) - C_K(y)] + [C_T(w) - C_L(w)], \end{aligned}$$

so that

$$C_K(v) \leq C_K(x) + [C_L(z) - C_K(z)] + [C_K(y) - C_L(y)],$$

which implies that

$$\begin{aligned} C_K(v) &\leq C_K(x) + 2(L - K)B \\ &\leq C_K(x) + 2RB. \end{aligned}$$

Hence, $A_K(v) \leq A_K(x) + 2RB/K$, where K is arbitrary. Consequently,

$$A(v) = \limsup_K A_K(v) \leq \limsup_K A_K(x) + \limsup_K \frac{2RB}{K} = A(x),$$

which implies that $v \in X^*$. Thus, $\mathcal{X}^* \subseteq X^*$. □

Corollary 4.3. *Suppose (P) satisfies Bounded Reachability. Then there exists an average optimal solution for (P) , i.e., $X^* \neq \emptyset$.*

Proof. Follows from Lemma 3.5 and Theorem 4.2. □

5. Optimal Solution Convergence

We showed in the previous section that efficient solutions, in the presence of Bounded Reachability, are average-cost optimal, i.e., $\emptyset \neq \mathcal{X}^* \subseteq X^*$. However, in general, the inclusion is strict, $\mathcal{X}^* \subset X^*$, i.e., there are average-cost optimal solutions which may fail to be efficient. This is clear since average-cost optimality is a long-run property while efficiency is a short-run property. It is in this sense that efficient solutions strengthen the under-selective criterion of average-cost optimality (Hopp, Bean, and Smith [1987]). Moreover, efficient solutions inherit the long-run properties of average-cost optimality through their behavior in the short-run, i.e., optimality to every state through which they pass. It is this latter characterizing property that we exploit in this section to approximate an efficient, and hence average-cost optimal, solution by solving for short-run optimal solutions to (P) .

The following result says that the N -horizon efficient solutions \mathcal{X}_N^* arbitrarily well approximate the infinite horizon efficient solutions \mathcal{X}^* , for sufficiently long horizon N .

Lemma 5.1. *(Average-Cost Optimal Solution Set Convergence) The sequence $\{\mathcal{X}_N^*\}$ of N -horizon efficient solution sets converges in $\mathcal{K}(Y)$ to the set \mathcal{X}^* of all infinite horizon efficient solutions.*

Proof. This follows from the definition of \mathcal{X}^* and Theorem 3.4. □

The previous lemma assures us that for every efficient solution $x^* \in \mathcal{X}^*$, there is a sequence of N -horizon efficient solutions $x_N^* \in \mathcal{X}_N^*$ which converges to x^* . We turn next to the problem of finding such a selection

$\{x_N^*\}$ from the $\{\mathcal{X}_N^*\}$. Following Ryan, Bean, and Smith [1992], we recall the notion of *lexicomin* of a subset of Y .

Define the *lexico order* \prec on Y as follows. For $x, y \in Y$, $x \prec y$ if there exists a positive integer k for which $x_j = y_j$, for $j = 1, \dots, k-1$, and $x_k < y_k$. We will write $x \preceq y$ if $x = y$ or $x \prec y$. Clearly, \preceq is a partial order on Y . Since it also satisfies the Law of Trichotomy, i.e., for each $x, y \in Y$, either $x \prec y$, $x = y$ or $y \prec x$, the relation \prec is a linear order on Y . It will be convenient to denote the element $(0, 0, \dots)$ in the compact metric space Y by θ . For our choice of β , i.e., $0 < \beta < 1/(m+1)$, the following lemma shows that the distances to elements of Y from θ determine the lexico order \prec .

Lemma 5.2. *Let $x, y \in Y$. Then $x \prec y$ if and only if $d(\theta, x) < d(\theta, y)$. Hence, $x = y$ if and only if $d(\theta, x) = d(\theta, y)$.*

Proof. Omitted. □

Now let V be an arbitrary non-empty subset of Y and $\lambda \in Y$. We define λ to be a *lexicomin* of V if $\lambda \in V$ and $\lambda \prec x$, for all $x \in V$ such that $x \neq \lambda$. Note that lexicomins are unique, when they exist. (Suppose λ and μ are distinct lexicomins of V . Then $\lambda \prec \mu$ and $\mu \prec \lambda$ necessarily, and without loss of generality, it follows that there exists a positive integer k such that $\lambda_k \leq \mu_k$ and $\mu_k < \lambda_k$. Contradiction.) We will denote the lexicomin of V by $\lambda(V)$, whenever it exists.

The following lemma emphasizes that a lexicomin is a nearest point selection.

Lemma 5.3. *Let V be a subset of Y and $x \in V$. Then x is the lexicomin of V if and only if x is the point in V nearest to θ .*

Proof. Follows immediately from Lemma 5.2. □

The following lemma assures us that such a nearest-point selection from a set exists whenever the set is closed.

Lemma 5.4. *Let V be a closed non-empty subset of Y . Then there exists a (unique) lexicomin of V . In particular, θ is the lexicomin of Y .*

Proof. The function $x \rightarrow d(\theta, x)$ is continuous on the compact set V . Hence, it attains its minimum at some point of V , which is necessarily the lexicomin of V by the previous lemma. □

The converse of this lemma is not true in general.

The following proposition assures us that the lexicomin selection from a sequence of Kuratowski-converging sets is convergent.

Proposition 5.5. *Let $\{F_n\}$ be a sequence in $\mathcal{K}(Y)$ which converges to F in $\mathcal{K}(Y)$. Then the sequence $\{\lambda(F_n)\}$ converges to $\lambda(F)$ in Y .*

Proof. In the terminology of Schochetman and Smith [1989], θ is a *uniqueness point* for F and $\{\lambda(F_n)\}$ is a *nearest-point selection* from the F_n defined by θ . The result then follows from Theorem 3.4 of Schochetman and Smith [1989]. □

We are now in position to obtain our average-cost optimal solution convergence result in terms of solution strategies instead of solution strategy sets.

Theorem 5.6. *(Average-Cost Optimal Policy Convergence) The sequence $\{\lambda(\mathcal{X}_N^*)\}$ of lexicomin N -horizon efficient solutions converges in Y to the lexicomin infinite horizon efficient solution $\lambda(\mathcal{X}^*)$, i.e.,*

$$\lim_{N \rightarrow \infty} \lambda(\mathcal{X}_N^*) = \lambda(\mathcal{X}^*).$$

Proof. Recall that $\mathcal{X}_N^* \rightarrow \mathcal{X}^*$ in $\mathcal{K}(Y)$, as $N \rightarrow \infty$. Since \mathcal{X}^* and the \mathcal{X}_N^* are closed and non-empty, each contains a lexicomin denoted by $\lambda(\mathcal{X}^*)$ and $\lambda(\mathcal{X}_N^*)$, respectively, $\forall N = 1, 2, \dots$. By Proposition 5.5, we have that $\lim_{N \rightarrow \infty} \lambda(\mathcal{X}_N^*) = \lambda(\mathcal{X}^*)$ in Y , where $\lambda(\mathcal{X}^*) \in X^*$, i.e., it is an optimal solution for (P) . Thus, we have average-cost optimal solution convergence in the familiar point sense. \square

6. Value Convergence

In the previous section, we demonstrated policy convergence of solutions of the approximating problems (P_N) to an infinite horizon average-cost optimal solution of (P) . However, since the average-cost operators A_N fail to converge uniformly to A (indeed, A is not a continuous function in general), we cannot conclude value convergence in the average-cost case. In particular, the optimal average cost A_N^* of problem (P_N) need not converge to the optimal average cost A^* of (P) ; in fact, it need not converge at all. Despite this, there is a weaker sense in which the optimal average-cost values $\{A_N^*\}$ of the subproblems (P_N) approximate the optimal average-cost value A^* of the problem (P) .

Theorem 6.1. *Suppose (P) satisfies Bounded Reachability. Then*

$$A^* = \limsup_N A_N^*.$$

Proof. Let $x^* \in \mathcal{X}^*$ (Lemma 3.5) and $x_N^* \in X_N^*$, so that $A^* = A(x^*)$ (Theorem 4.2) and $A_N^* = A_N(x_N^*)$, $\forall N = 1, 2, \dots$. Since $x^* \in X$ and $X \subseteq X_N$, it follows that $x^* \in X_N$, so that $A_N(x_N^*) \leq A_N(x^*)$, $\forall N = 1, 2, \dots$. Hence, since both sequences are bounded, by the results of Goldberg [1964], we have that $\limsup_N A_N(x_N^*) \leq \limsup_N A_N(x^*)$, i.e., $\limsup_N A_N^* \leq A^*$.

Conversely, let R be as in the Bounded Reachability Property. Fix $N > 2R$, so that $M = N - R > R$, and set $v = x_N^*$ for convenience. Applying the Bounded Reachability Property to the data

$$(M, s_M(v); (s_M(x^*), \dots, s_N(x^*))),$$

we conclude that there exists $M \leq L \leq N$ and $w \in X_L$ such that $s_M(w) = s_M(v)$ and $s_L(w) = s_L(x^*)$. Define

$$z = (v_1, \dots, v_M, w_{M+1}, \dots, w_L, x_{L+1}^*, x_{L+2}^*, \dots).$$




Figure 2 - Optimal average-cost convergence

Clearly, $z \in X_L$ by Lemma 2.5. Also by Lemma 2.5 for (L, N, z, x^*) , we have that $z \in X_N$ and $s_N(z) = s_N(x^*)$, so that $z \in X_N(s_N(x^*))$, and

$$c_j(s_{j-1}(z), z_j) = c_j(s_{j-1}(x^*), x_j^*), \quad \forall j = 1, 2, \dots, N.$$

Also recall that since $x^* \in \mathcal{X}^*$, we have that $x^* \in X_N^*(s_N(x^*))$.

Since x^* is an optimal solution to $(P_N(s_N(x^*)))$ and $z \in X_N(s_N(x^*))$, it follows that

$$\begin{aligned} C_N(x^*) &\leq C_N(z) = C_M(z) + [C_N(z) - C_M(z)] \leq C_M(z) + (N - M)B \\ &= C_M(z) + RB \leq C_M(v) + RB = C_N(v) + [C_M(v) - C_N(v)] + RB \\ &\leq C_N(v) + |C_N(v) - C_M(v)| + RB \leq C_N(v) + (N - M)B + RB = C_N(v) + 2RB, \end{aligned}$$

so that $A_N(x^*) \leq A_N(v) + 2RB/N$, for $N > 2R$. Hence, by the results of Goldberg [1964], we have

$$\begin{aligned} A^* &= \limsup_N [A_N(x^*)] \leq \limsup_N \left[A_N(v) + \frac{2RB}{N} \right] \\ &\leq \limsup_N [A_N(v)] + 2 \lim_{N \rightarrow \infty} \frac{RB}{N} = \limsup_N A_N(v), \end{aligned}$$

i.e., $A^* \leq \limsup_N A_N^*$, which completes the proof. \square

Remark 6.2. Theorem 6.1 says that $A(x) = \limsup_N A_N(x^N)$, for $x = x^* \in \mathcal{X}^*$ and $x^N = x_N^* \in X_N^*$, $\forall N = 1, 2, \dots$. For a general sequence $\{x^N\}$ such that $x^N \in X_N$, all N , even if $x^N \rightarrow x$ in Y , as $N \rightarrow \infty$, it is not necessarily the case that $A(x) = \limsup_N A_N(x^N)$. Thus, the convergence claimed in Theorem 6.1 is restricted to optimal average-cost values, unlike in the discounted case.

7. A Next-Decision Forward Algorithm

We know from average-cost optimal policy convergence (Theorem 5.6) that the first decision of the lexicomin efficient solution to problem (P_N) , for large N , is the same as the first decision of the lexicomin infinite horizon average-cost optimal solution. (Note the trivial nature of this claim when made with respect to an arbitrary average-cost optimal solution, since finite leading decisions are irrelevant to the property of being average cost optimal!) A finite algorithm for discovery of this first decision is then lacking only a procedure for numerically discovering how large N must be for the claimed invariance to occur. We incorporate the determination of such a horizon, called a *solution horizon* (Bes and Sethi [1988]), within the forward algorithm below. We find the efficient solution set and the associated lexicomin solution for increasingly longer horizons N , until a stopping rule is met. At this point, a solution horizon has been discovered; hence, the first decision of the lexicomin infinite horizon average-cost optimal efficient solution has also been discovered. (The algorithm is similar in spirit to that proposed in Lasserre [1986] in the discounted case except that we use a lexicomin selection thus avoiding the assumption there that the optimal first decision is unique.)

Forward Algorithm

- (1) Set horizon $N = 1$.
- (2) Find the set $X_N^*(s_N)$ of all solutions optimal to state s_N at horizon N , for each state s_N in the set S_N of all N -horizon feasible states.
- (3) Determine the lexicomin efficient solutions $\lambda_N^*(s_N)$ to state s_N at horizon N , for all feasible states $s_N \in S_N$. Set the lexicomin efficient solution λ_N^* for horizon N equal to the lexicomin of $\{\lambda_N^*(s_N) : s_N \in S_N\}$.
- (4) (**Stopping Rule**) If $(\lambda_N^*)_1 = (\lambda_N^*(s_N))_1$, for all $s_N \in S_N$, stop. The first decision of the lexicomin infinite horizon average optimal solution is $\lambda_1^* = (\lambda^*(N))_1$. Otherwise, increment the horizon N to $N + 1$ and go to step 2.

It is clear from the Principle of Optimality that if the Stopping Rule is met, then the first decision of the lexicomin infinite horizon efficient solution has been found. The next theorem provides a sufficient condition that guarantees the Stopping Rule will eventually be met.

Theorem 7.1. *Suppose that for every choice (s_1, s_2, \dots) in $\prod_{N=1}^{\infty} S_N$, we have that*

$$\lim_{N \rightarrow \infty} X_N^*(s_N) = \mathcal{X}^*$$

in $\mathcal{K}(Y)$. Then the Stopping Rule holds at some $1 \leq N < \infty$.

Proof. For each $N = 1, 2, \dots$, and each $s_N \in S_N$, let $\lambda_N^*(s_N) = \text{lexicomin of } X_N^*(s_N)$, $\lambda_N^* = \text{lexicomin of } \mathcal{X}_N^*$ and $\lambda^* = \text{lexicomin of } \mathcal{X}^*$. Note that λ_N^* must be in $X_N^*(r_N)$, for some unique $r_N \in S_N$, by definition of \mathcal{X}_N^* . Necessarily, λ_N^* is the lexicomin of $X_N^*(r_N)$, i.e., $\lambda_N^* = \lambda_N^*(r_N)$, $\forall N = 1, 2, \dots$.

Now suppose that the hypothesis is satisfied, but not the Stopping Rule. Then, for each $N = 1, 2, \dots$, there exists $t_N \in S_N$ such that $r_N \neq t_N$ and

$$(\lambda_N^*(t_N))_1 \neq (\lambda_N^*(r_N))_1 = (\lambda_N^*)_1.$$

If not, then the Stopping Rule would be satisfied at any N where this was not the case. But we have seen that $\lambda_N^* \rightarrow \lambda^*$ in Y , so that $(\lambda_N^*)_1 = \lambda_1^*$, for large N (Lemma 2.1). Consequently, $(\lambda_N^*(t_N))_1 \neq \lambda_1^*$, for large N . But, by hypothesis, $\lim_{N \rightarrow \infty} X_N^*(t_N) = \mathcal{X}^*$, so that $\lim_{N \rightarrow \infty} \lambda_N^*(t_N) = \lambda^*$, which implies (Lemma 2.1) that $\lambda_N^*(t_N)_1 = \lambda_1^*$, for large N . This is a contradiction. \square

The hypothesis of Theorem 7.1 is unfortunately difficult to establish directly in specific cases. We can however provide a sufficient condition for this hypothesis which, although quite strong, is not vacuous (see section 8 for an instance where it is met).

Theorem 7.2. *If \mathcal{X}^* is a singleton $\{x^*\}$, then the hypothesis of Theorem 7.1 is satisfied.*

Proof. Let $s_N \in S_N$, for each $N = 1, 2, \dots$. It suffices to show that $\limsup_N X_N^*(s_N) \subseteq \{x^*\}$ and $x^* \in \liminf_N X_N^*(s_N)$.

If $x \in \limsup_N X_N^*(s_N)$, then $x \in \limsup_N \mathcal{X}_N^*$, by definition of \mathcal{X}_N^* . But $\limsup_N \mathcal{X}_N^* = \mathcal{X}^*$, so that $x = x^*$ by hypothesis.

Conversely, since $\lim_{N \rightarrow \infty} \mathcal{X}_N^* = \mathcal{X}^*$, x^* is the limit of any selection from the \mathcal{X}_N^* . In particular, x^* is the limit of any selection from the $X_N^*(s_N)$, i.e., $x^* \in \liminf_N X_N^*(s_N)$. \square

From a practical point of view, it is important to note that even in the absence of conditions that guarantee that the Stopping Rule will eventually be met, if for whatever reason the Stopping Rule *is* eventually met in an application of the forward algorithm, then we are assured that the first decision of an efficient infinite horizon average-cost optimal solution has been found.

8. Concluding Remarks

Virtually every discrete infinite horizon decision making problem is included within the framework of the general problem (P) posed in section 2. See Schochetman and Smith [1997] for an extensive application of the above theory to the problem of acquiring and retiring equipment over an infinite horizon in the presence of technological change (Denardo [1982], Bean, Lohmann, and Smith [1985], [1994]). We also establish there that the key property of Bounded Reachability is satisfied for this case.

Consequently, all of the previous results hold for this model. In particular, the Forward Algorithm of section 7 may be employed, whenever the Stopping Rule is eventually met, to find the next replacement decision in a rolling horizon procedure, to recover an optimal replacement schedule over the infinite horizon. Similar stopping rules have been empirically observed to be usually met in practice, at least in the discounted-cost case (Bean, Lohmann, and Smith [1985]). A sufficient theoretical condition for it to be met in the average-cost case is (by Theorem 7.2) that $\mathcal{X}^* = \{x^*\}$. For example, in the classic problem of replacing equipment under stationary technology, since the problem is mathematically equivalent to a knapsack problem, a singleton efficient solution obtains whenever the turnpike policy equipment type that minimizes the infinite horizon average-cost under a stationary replacement strategy is unique (Gillmore and Gomory [1966]). We note that it is not difficult to extend this observation to the nonstationary case of technological change.

Acknowledgment: We are grateful to an anonymous referee for several suggestions that significantly improved the clarity of the exposition.

REFERENCES

1. Aliprantis, C., K. C. Border (1994). *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer-Verlag, NY.
2. Aubin, J.-P. (1990). *Set-Valued Analysis*, Birkhauser, Boston.
3. Bean, J. C., R. L. Smith, J. B. Lasserre (1990). Denumerable State Nonhomogeneous Markov Decision Processes. *Journal of Mathematical Analysis and Applications*. **153** 64–77.
4. Bean, J. C., J. R. Lohmann, R. L. Smith (1994). Equipment Replacement Under Technological Change. *Naval Research Logistics*. **41** 117–128.
5. ——— (1985). A Dynamic Infinite Horizon Replacement Economy Decision Model *The Engineering Economist*. **30** 99–120.

6. Bean, J. C., R. L. Smith (1984). Conditions for the Existence of Planning Horizons. *Mathematics of Operations Research* **9** 391–401.
7. Berge, C. (1963). *Topological Spaces*, Oliver and Boyd, London.
8. Bes, C., S. Sethi (1988). Concepts of Forecast and Decision Horizons: Applications to Dynamic Stochastic Optimization Problems. *Mathematics of Operations Research*. **13** 295-310.
9. Denardo, E. V. (1982). *Dynamic Programming: Models and Applications*, Prentice-Hall, Englewood Cliffs.
10. Derman, C. (1966). Denumerable state Markovian decision processes - average cost case. *Annals of Mathematical Statistics*. **37** 1545–1554.
11. Federgruen, A., H. C. Tijms (1978). The Optimality equation in average cost denumerable state semi-Markov decision problems, recurrency conditions and algorithms. *Journal of Applied Probability*. **15** 356–373.
12. Gillmore, P. C., R. E Gomory (1966). The Theory and Computation of Knapsack Functions. *Operations Research*. **14** 1045-1074.
13. Goldberg, R. R. (1964). *Methods of Real Analysis*, Blaisdell Pub. Co., Waltham.
14. Hausdorff, F. (1962). *Set Theory*, 2nd ed., Chelsea, New York.
15. **35** 875-883.
16. Kuratowski, K. (1966). *Topologie I, II*, Academic Press, New York.
17. Lasserre, J. (1986). Detecting Planning Horizons in Deterministic Infinite Horizon Optimal Control. *IEEE Transactions on Automatic Control* **31** 70-72.
18. Park, Y., J. C. Bean, R. L Smith (1993). Optimal Average Value Convergence in Nonhomogeneous Markov Decision Processes. *Journal of Mathematical Analysis and Applications*. **179** 525-536.
19. Puterman, M. L. (1994). *Markov decision processes : discrete stochastic dynamic programming*, John Wiley & Sons, New York.
20. Ross, S. (1968). Non-discounted Denumerable Markov Decision Models. *Annals of Mathematical Statistics*. **39** 412–423.
21. _____ (1983). *Introduction to Dynamic Programming*, Academic Press, NY.
22. Ryan, S. M., J. C. Bean, R. L Smith (1992). A Tie-breaking Algorithm for Discrete Infinite Horizon Optimization. *Operations Research*. **40** S117–S126.
23. Schochetman, I. E., R. L Smith (1989). Infinite Horizon Optimization. *Mathematics of Operations Research*. **14** 559–574.
24. _____ (1992). Finite Dimensional Approximation in Infinite Dimensional Mathematical Programming. *Mathematical Programming*. **54** 307–333.
25. _____ (1997). *Existence and Discovery of Average-Cost Optimal Solutions in Deterministic Infinite Horizon Optimization*, Technical Report 97-07, Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor.