This is a manuscript published as Ryan, S. M., J. Bean and R. L. Smith, "A Tie-Breaking Algorithm for Discrete Infinite Horizon Optimization," Operations Research, 40, S117-S126 (1992). Available at http://dx.doi.org/10.1287/opre.40.1.S117. Posted with permission.

A TIE-BREAKING RULE FOR DISCRETE INFINITE HORIZON OPTIMIZATION *

Sarah M. Ryan

Department of Industrial Engineering
University of Pittsburgh
Pittsburgh, Pennsylvania 15261

James C. Bean
Robert L. Smith

Department of Industrial and Operations Engineering
The University of Michigan
Ann Arbor, Michigan 48109

April, 1989

Abstract

We study discrete infinite horizon optimization problems without the common assumption of a unique optimum. A method based on solution set convergence is employed for finding optimal initial decisions by solving finite horizon problems. This method is applicable to general discrete decision models which satisfy a weak reachability condition. The algorithm, together with a stopping rule, is applied to solve capacity expansion problems, and computational results are reported.

*This work was supported in part by the National Science Foundation under Grant ECS-8700836 to The University of Michigan.

In sequential decision problems with indefinite horizons, the length of an appropriate planning horizon is an important issue. Data far in the future is difficult to forecast accurately, and yet myopia must be avoided. Recently, progress has been made in proving the existence of, and then discovering, forecast and solution horizons. These are horizons long enough to guarantee that the optimal initial decision or decision sequence found is optimal over any longer horizon (including the infinite horizon). The decision maker may implement this decision or sequence with confidence, and then uncover subsequent decisions in a rolling horizon fashion.

Solution horizon procedures typically consist of solving finite horizon problems, increasing the horizon until some stopping criterion is satisfied. Two questions then arise:

1. (Solution Horizon Existence) Is there a horizon long enough to guarantee optimality?
2. (Solution Horizon Discovery) How long a horizon is needed to *know* that optimality is guaranteed?

Nearly all solution horizon existence and discovery results have required a unique optimal initial decision (see Bean and Smith 1984; Hopp, Bean, and Smith 1987; Bès and Sethi 1988). However, Ryan and Bean [1987] show that in discrete decision problems, such a requirement may be difficult to meet. Moreover, discrete decisions arise in many problems, including production planning, capacity expansion and equipment replacement.

Recently, Schochetman and Smith [1987] began to attack problems with multiple optima by studying the convergence of sets of finite horizon optimal solutions. They showed that if these sets converge as the horizon is lengthened, then a solution horizon may be found by making appropriate selections from the sets. In this paper we force set convergence by a suitable construction of the finite horizon sets. Rather than merely finding optimal solutions, our algorithm finds the sets of solutions optimal to their own state. Similar ideas appear in Bean and Smith [1986] and Chand and Morton [1986]. We provide a simple structural condition under which these sets converge. Further, we derive an easily implemented selection, or tie-breaking rule, for selecting one solution from each finite horizon set. The sequence of selected solutions converges to an infinite horizon optimal strategy.

These constructions lead to a straightforward tie-breaking algorithm for choosing one of potentially many optimal strategies. Our stopping rule, while not requiring uniqueness, may still fail when there is more than one optimal initial decision. However, as illustrated by an example, even when the algorithm fails to stop, the optimal strategies can be recognized readily by examining the finite horizon sets. When applied to regeneration point problems, the algorithm generalizes the forward algorithms of Shapiro and Wagner [1967] and Bean and Smith [1984].

Section 1 is a mathematical statement of the problem and assumptions. In Section 2 we review concepts of set convergence and apply them to discrete infinite horizon optimization. The construction of finite horizon solution sets, conditions under which they converge, and the algorithm and stopping rule appear in Section 3. In Section 4 we apply the algorithm to capacity expansion problems and discuss the results of computational tests. Finally, Section 5 contains conclusions.

1 Problem Definition and Assumptions

We model the infinite horizon sequential decision problem as in Bean and Smith [1986] with an infinite directed decision network $(\mathcal{N}, \mathcal{A}, C)$ where $\mathcal{N}$ is the set of nodes or *decision points*, $\mathcal{A}$ is the set of arcs or *decisions*, and $C : \mathcal{A} \mapsto \Re$ is a cost function.

We impose three structural assumptions on $(\mathcal{N}, \mathcal{A})$. First, we require that there be a unique root node with in-degree one. Second, we assume that all node out-degrees are non-zero and uniformly bounded. Finally, we assume that the cumulative in-degrees of all nodes are finite, where the cumulative in-degree of a node is the sum of its in-degree and all in-degrees of nodes from which there is a directed path to that node. From these assumptions we can number the nodes $\mathcal{N} = \{0, 1, 2, \ldots\}$ such that $(i, j) \in \mathcal{A}$ only if $i < j$ (Skilton 1985, p. 230). Hence the node numbers can serve as a surrogate for time. Moreover, it follows that there is a directed path from the root node 0 to each node i which can be continued over the infinite horizon.

A path through the network, $(i_0, i_1, \ldots)$, where $i_0 = 0$, represents a feasible strategy $\pi = (\pi_1, \pi_2, \ldots)$ in which the n^{th} decision $\pi_n = (i_n, i_{n+1})$. Let Π_n be the set of decisions available after $n - 1$ decisions have been made. We assume the decisions in Π_n are indexed by a subset of $\{0, 1, \ldots, M\}$, where M is uniform over n. Associated with each node i is a time T_i, called a *decision epoch*, such that $i \leq j$ if and only if $T_i \leq T_j$. We assume $T_i \to \infty$ as $i \to \infty$.

Denote the set of *feasible strategies* (or paths) by $\Pi \subseteq \times_{n=1}^{\infty} \Pi_n$. We define a metric on $\times_{n=1}^{\infty} \{0, 1, \ldots, M\}$ as follows: for $\pi, \hat{\pi} \in \times_{n=1}^{\infty} \{0, 1, \ldots, M\}$,

$$\rho(\pi, \hat{\pi}) \equiv \sum_{n=1}^{\infty} \beta^n |\pi_n - \hat{\pi}_n|,$$

where $\beta < \frac{1}{M+1}$. Under this metric, the closeness of strategies is measured by agreement in early decisions. This metric induces a topology on Π which is identical to the product topology; therefore Π is a compact metric space (see Bean and Smith 1986).

Let $i_k(\pi)$ denote the k^{th} node visited by strategy π, where $i_0(\pi) = 0$. The cost of π is given by

$$f_\pi \equiv \sum_{n=0}^{\infty} C(i_n(\pi), i_{n+1}(\pi)).$$

We assume that f_π is uniformly convergent over $\pi \in \Pi$. The problem we wish to solve is:

$$f \equiv \min_{\pi \in \Pi} f_\pi.$$

The minimum exists since Π is compact and f_π is a uniformly convergent sequence of continuous functions and therefore continuous over $\pi \in \Pi$. A strategy $\hat{\pi}$ is termed *infinite horizon optimal* if it minimizes f_π. We will also refer to these as *optimal* strategies. An *optimal initial decision* is an initial decision for some optimal strategy. Let Π^* and Π_1^* denote sets of optimal strategies and optimal initial decisions, respectively.

The solution horizon approach involves solving finite horizon problems. We define the T–horizon cost of π as

$$f_\pi(T) \equiv \sum_{\{n | T_{i_n(\pi)} \leq T\}} C(i_n(\pi), i_{n+1}(\pi)).$$

We also define the following finite horizon sets of strategies:

$$\Pi^*(T) \equiv \{\pi \in \Pi | \pi \in \arg\min f_\pi(T)\}$$

$$\hat{\Pi}(i) \equiv \left\{ \pi | \pi \in \underset{\{\pi | i = i_k(\pi) \text{ for some } k\}}{\arg\min} \sum_{n<k} C(i_n(\pi), i_{n+1}(\pi)) \right\}$$

$$\bar{\Pi}(T) \equiv \{\pi | \pi \in \hat{\Pi}(i_n(\pi)),\ n \text{ such that } T_{i_{n-1}(\pi)} < T \leq T_{i_n(\pi)}\}.$$

The set $\Pi^*(T)$ is the set of T–horizon *optimal* strategies, while $\hat{\Pi}(i)$ is the set of strategies optimal to node i. The set $\bar{\Pi}(T)$ is the set of strategies optimal to their own node at or just beyond time T. We will refer to this as the set of T–horizon *efficient* strategies. In a discrete time capacity expansion or production planning problem, these are the strategies optimal to their own capacity or inventory level at time T, respectively. Note that $\Pi^*(T) \subseteq \bar{\Pi}(T)$. A subscript of k on any set of strategies will denote the corresponding set of k^{th} decisions.

A *solution horizon* is a time, $\hat{T}$, such that for $T \geq \hat{T}$, $\Pi_1^*(T) = \{\pi_1^*\}$, for some fixed $\pi_1^* \in \Pi_1^*$. A *general solution horizon* is a time, $\hat{T}$, such that for $T \geq \hat{T}$, $\Pi_1^*(T) \subseteq \Pi_1^*$. Most algorithms in the literature have required the existence of solution horizons. In this paper, we provide an algorithm that selects an optimal initial decision from Π_1^* in the presence of a general solution horizon.

2 Set Convergence in Discrete Infinite Horizon Optimization

Previous authors (Bean and Smith 1984, Schochetman and Smith) have found solution horizons by showing that if $\Pi^* = \{\pi^*\}$, then an arbitrary choice $\pi^*(T) \in \Pi^*(T)$ converges to π^* as $T \to \infty$. If optimal strategies are not unique, a natural extension is to seek convergence of the *sets* $\Pi^*(T)$ to the *set* Π^*. We have found such convergence to be rare, and have studied instead the convergence of $\bar{\Pi}(T)$ to Π^*. To lay a foundation for this development, we first review and apply some concepts of set convergence for closed subsets of a compact metric space.

Definition: *Let* $\Pi(T) \subseteq \Pi$. *We say that* $\Pi(T)$ *is a* $T-$horizon set *if membership in* $\Pi(T)$ *is determined by decisions made at or before time* T *(inclusive).*

Schochetman and Smith introduce set convergence in infinite horizon optimization, using the Hausdorff metric for closed and hence compact sets. As this metric is built up from the strategy metric, ρ, set convergence is closely related to individual strategy convergence. To establish and then exploit convergence of a sequence of finite horizon closed sets $\{\Pi(T_n)\}$ to the optimal set Π^*, we will use the following results for the discrete case.

R1. *Set convergence means that subsequential limits are limits.* Define $\liminf \Pi(T_n)$ to be the set of all $\pi \in \Pi$ such that there exists $\pi^n \in \Pi(T_n)$ for all n such that $\pi^n \to \pi$. Define $\limsup \Pi(T_n)$ to be the set of all $\pi \in \Pi$ such that there exists a subsequence

$\{T_k\}$ of $\{T_n\}$ and $\pi^k \in \Pi(T_k)$ such that $\pi^k \to \pi$. Then $\liminf \Pi(T_n) \subseteq \limsup \Pi(T_n)$. Schochetman and Smith show that $\Pi(T_n) \to \Pi^*$ in the Hausdorff metric if and only if $\limsup \Pi(T_n) = \liminf \Pi(T_n) = \Pi^*$.

R2. *Set convergence is equivalent to early decision agreement.* $\Pi(T_n) \to \Pi^*$ if and only if, for any L, there exists N_L such that if $n \geq N_L$, then

1. for any $\pi \in \Pi^*$, there exists $\pi^n \in \Pi(T_n)$ such that $\pi_k^n = \pi_k$, $1 \leq k \leq L$, and
2. for any $\pi^n \in \Pi(T_n)$, there exists $\pi \in \Pi^*$ such that $\pi_k = \pi_k^n$, $1 \leq k \leq L$.

In particular, if $\Pi(T_n) \to \Pi^*$ and $\pi_k = \pi_k^*$ for all $\pi \in \Pi^*$, $1 \leq k \leq L$, then there exists N_L such that if $n \geq N_L$, then $\pi_k^n = \pi_k^*$, $1 \leq k \leq L$, for all $\pi^n \in \Pi(T_n)$.

R3. *Set convergence implies convergence of nearest-point selections.* Let p be a point for which there is a unique $\pi \in \Pi^*$ minimizing $\rho(p, \pi)$. Then p is called a *uniqueness point* for Π^*. Let $s_p(\Pi')$ be a strategy minimizing $\rho(p, \pi)$ over $\pi \in \Pi'$. The function $s_p(\cdot)$ is called a *nearest-point selection.* From Schochetman and Smith, $\Pi(T_n) \to \Pi^*$ implies $s_p(\Pi(T_n)) \to s_p(\Pi^*)$.

In the following, we show that the finite horizon and infinite horizon sets of strategies involved in solving an infinite horizon problem are necessarily closed. Next we define a unique lexicographic order for strategies and show that it is identical to the one obtained by ordering according to distance from the point $\theta \equiv \mathbf{0}$, measured by ρ. These results, together with the continuity of ρ, imply that θ is a uniqueness point for Π^* and interpret $s_\theta(\cdot)$ as a simple tie-breaking rule, analogous to one used for resolving degeneracy in linear programming.

Lemma 1 *The sets Π^* and $\Pi(T)$, where $\Pi(T)$ is any $T-$horizon set, are closed.*

Proof: Π^* is closed since f_π is continuous in π and Π is compact. For $\Pi(T)$, note that there are a finite number of partial strategies up to time T. If $\Pi(T)$ is finite then it is closed. Suppose that $\Pi(T)$ is infinite and that π is a cluster point of $\Pi(T)$. Let $\{\pi^n\}_{n=1}^\infty$ be such that $\pi^n \in \Pi(T)$ for all n and $\pi^n \to \pi$ in ρ. Then $\{\pi^n\}_{n=1}^\infty$ is Cauchy, which implies that for some N, $\{\pi^n,\ n \geq N\}$ are in agreement with one another and with π up to time T. Therefore $\pi \in \Pi(T)$. Thus $\Pi(T)$ contains all its cluster points, and $\Pi(T)$ is closed. ∎

Definition: *Let* $a = (a_1, a_2, \ldots)$ *and* $b = (b_1, b_2, \ldots)$. *We say that* $a \prec b$ *(*a is lexicographically smaller than b*) if and only if* $a \neq b$ *and, if* n_0 *is the smallest* n *such that* $a_n \neq b_n$, *then* $a_{n_0} < b_{n_0}$.

A strategy $\tilde{\pi}$ *is the* lexico minimum *of* Π' *if* $\tilde{\pi} \prec \pi$ *for all* $\pi \in \Pi'$, $\pi \neq \tilde{\pi}$.

Lemma 2 *For any closed set* $\Pi' \in \Pi$, $s_\theta(\Pi')$ *is the lexico minimum element of* Π' *and therefore is unique.*

Proof: We will show that:

1. $\rho(\theta, \pi)$ attains its minimum on $\pi \in \Pi'$,
2. $\rho(\theta, \pi^1) < \rho(\theta, \pi^2)$ if and only if $\pi^1 \prec \pi^2$.

(1) and (2) imply that $\rho(\theta, \pi)$ attains its minimum at the unique lexico minimum point of Π'.

1. Follows immediately from the continuity of ρ and compactness of Π.
2. Suppose $\pi^1 \prec \pi^2$. Then for some n_0 we have

$$\pi^1_n = \pi^2_n \text{ for } n < n_0 \text{ and } \pi^1_{n_0} < \pi^2_{n_0}.$$

By assumption, $\pi_n^1 - \pi_n^2 \geq -M$ for $n \geq n_0$. Then

$$\rho(\theta, \pi^2) - \rho(\theta, \pi^1) = \sum_{n=n_0}^{\infty} \beta^n (\pi_n^2 - \pi_n^1)$$

$$\geq \beta^{n_0} - \frac{M\beta^{n_0+1}}{1-\beta} > 0 \text{ since } \beta < \frac{1}{M+1}.$$

Hence,

$$\pi^1 \prec \pi^2 \Rightarrow \rho(\theta, \pi^1) < \rho(\theta, \pi^2).$$

Now suppose $\pi^1 \not\prec \pi^2$. Either $\pi^1 = \pi^2$ or $\pi^1 \succ \pi^2$. If $\pi^1 = \pi^2$ then $\rho(\theta, \pi^1) = \rho(\theta, \pi^2)$. Suppose $\pi^1 \succ \pi^2$. Then by the same argument as above, $\rho(\theta, \pi^1) > \rho(\theta, \pi^2)$. Hence

$$\pi^1 \not\prec \pi^2 \Rightarrow \rho(\theta, \pi^1) \geq \rho(\theta, \pi^2). \blacksquare$$

Lemmas 1 and 2 imply that we can apply result R3 whenever any sequence of T_n–horizon sets, $\{\Pi(T_n),\ n \in N\}$, converges to Π^*. To summarize, we have:

Theorem 3 *If $\{\Pi(T_n)\}$ is some collection of T_n–horizon sets, then the lexico minimum element of $\Pi(T_n)$ converges to the lexico minimum element of Π^* whenever $\Pi(T_n) \rightarrow \Pi^*$.*

Proof: Follows from result R3 and Lemma 2. ∎

Assuming $\Pi(T_n) \rightarrow \Pi^*$, Theorem 3 suggests the following algorithm:

> Solve increasing length finite horizon problems, identifying for each T_n the lexico minimum element of $\Pi(T_n)$.

The sequence of strategies thus generated converges to the lexico minimum infinite horizon optimal strategy.

3 A Tie-Breaking Algorithm

Having seen that a tie-breaking algorithm can follow from set convergence, we now establish conditions for set convergence to occur. Various concepts of *reachability* (McKenzie 1976) have been used to identify solution horizons in the presence of a unique optimum (Lasserre 1986, Bean and Smith 1986). We show that a reachability condition guarantees the convergence of the efficient sets $\bar{\Pi}(T)$ to Π^* even in the absence of a unique optimum. This convergence leads to a tie-breaking algorithm for identifying solution horizons. Further, reachability sometimes can be determined simply from the problem structure.

Let $g(i)$ be the minimum cost from the root node to node i and $g'(i)$ be the minimum cost from i through the infinite horizon. Let $g(|i)$ be the cost of a minimum cost path from the root node over the infinite horizon which passes through node i. By the principle of optimality, $g(|i) = g(i) + g'(i)$.

As in Bean and Smith [1986], define *weak reachability* as follows:

Definition: *A sequence of nodes $\{i_n\}$, $i_n \to \infty$, is* weakly reachable *if, for all $\epsilon > 0$, there exists an N_ϵ such that for all $n \geq N_\epsilon$, there is a $j(n)$, an i'_n, and a path from $i^*_{j(n)}$ to i'_n at cost c_n with*

$$g(|i_n) \leq g(|i'_n) + \epsilon,$$

where $i^*_{j(n)}$ is a node for some optimal strategy, $i^*_{j(n)} \to \infty$, and $c_n \to 0$.

Bean and Smith show that $g(i_n) \to f$ as $n \to \infty$ if and only if $\{i_n\}$ is weakly reachable. Let $i(\pi, T)$ be the node for π at its first decision epoch greater than or equal to T. Thus $i(\pi, T)$ represents the head of the decision arc for π that ends at or just after time T. Let $\{T_n\}$ be any sequence of finite horizons such that $T_n \to \infty$.

Theorem 4 *If all sequences $\{i_k\} \subseteq \mathcal{N}$, $i_k \to \infty$, are weakly reachable then $\bar{\Pi}(T_n) \to \Pi^*$.*

Proof: We will show that

1. $\limsup \bar{\Pi}(T_n) \subseteq \Pi^*$,
2. $\Pi^* \subseteq \liminf \bar{\Pi}(T_n)$.

(1) and (2) imply that $\limsup \bar{\Pi}(T_n) = \liminf \bar{\Pi}(T_n)$. The result then follows from result R1 in Section 2.

1. Suppose $\pi \in \limsup \bar{\Pi}(T_n)$. Then there exists $\{T_k\} \subseteq \{T_n\}$ and $\pi^k \in \bar{\Pi}(T_k)$ such that $\pi^k \to \pi$. Let $i_k = i(\pi, T_k)$. Then $i_k \to \infty$ and $g(i_k) = f_{\pi^k}(T_k)$. From the uniform convergence of f_{π^k}, we have

$$f_{\pi^k}(T_k) \to f_\pi.$$

But

$$g(i_k) \to f$$

since $\{i_k\}$ is weakly reachable. Hence $f_\pi = f$, and $\pi \in \Pi^*$.

2. Suppose $\pi \in \Pi^*$. By the principle of optimality, π is optimal to each node it traverses. Thus $\pi \in \hat{\Pi}(i_n)$ for each n where $i_n = i(\pi, T_n)$. Hence, $\pi \in \bar{\Pi}(T_n)$ for all n and $\pi \in \liminf \bar{\Pi}(T_n)$. ∎

Lemma 5 *If $g(i) \leq g(j)$ whenever $T_i < T_j$, then all sequences $\{i_n\} \subseteq \mathcal{N}$, $i_n \to \infty$, are weakly reachable.*

Proof: Choose any $\{i_n\} \subseteq \mathcal{N}$, $i_n \to \infty$, and any $\epsilon > 0$. Let $\{i_n^*\}$ be $\{i_n(\pi)\}$ for some $\pi \in \Pi^*$. Let

$$j(n) = \max\{j | T_{i_j^*} \leq T_{i_n}\},$$

and

$$i'_n = i^*_{j(n)+1}.$$

Then $c_n = C(i^*_{j(n)}, i^*_{j(n)+1}) \to 0$ as $n \to \infty$ since $f < \infty$. By the uniform convergence of f_π over $\pi \in \Pi$, there exists N_ϵ such that $|g'(i_n)| < \epsilon/2$ and $|g'(i'_n)| < \epsilon/2$. Then for $n \geq N_\epsilon$,

$$g(|i_n) - g(|i'_n) = g(i_n) - g(i'_n) + g'(i_n) - g'(i'_n) < \epsilon,$$

since

$$g(i_n) - g(i'_n) \leq 0$$

by hypothesis and

$$g'(i_n) - g'(i'_n) \leq |g'(i_n)| + |g'(i'_n)| < \epsilon.$$

Hence $g(|i_n) \leq g(|i'_n) + \epsilon$. ∎

Theorem 6 *If $g(i) \leq g(j)$ whenever $T_i < T_j$, then $\bar{\Pi}(T_n) \to \Pi^*$.*

Efficient set convergence, therefore, follows from a simple structural property: that the minimum cost to a node is monotonic in the time to reach that node. In general, this property will only hold for a restricted class of problems where the decision network has been pruned sufficiently. For example, in the next section, we show that it holds in a general model of capacity expansion. However, in general weak reachability must be established directly (see Bean and Smith [1986] for several applications where weak reachability holds). In the meantime, either this result or Theorem 4, together with Theorem 3, yields the following corollary. Let $\bar{\pi}(T_n)$ and π^* be the lexico minimum elements of $\bar{\Pi}(T_n)$ and Π^*, respectively.

Corollary 7 *If either of the conditions of Theorem 4 or Theorem 6 are satisfied, then $\bar{\pi}(T_n) \to \pi^*$.*

We can now present an algorithm for the case when efficient sets converge to the optimal set. Let $\tilde{\Pi}(T)$ be the set $\bar{\Pi}(T)$ with, for any node n, ties between strategies optimal up to n broken by choosing the lexico minimum. Note that $\bar{\pi}(T)$ is its lexico minimum.

Tie-Breaking Algorithm

1. Let $\{T_n\}_{n=1}^{\infty}$ be any sequence of finite horizons with $T_n \to \infty$. Set $n \leftarrow 1$.

2. Solve the T_n–horizon problem to obtain $\tilde{\Pi}(T_n)$.

3. If for each $\pi(T_n) \in \tilde{\Pi}(T_n)$ and for all $1 \leq k \leq L$, $\pi_k(T_n) = \bar{\pi}_k(T_n)$, then stop. Otherwise, set $n \leftarrow n+1$ and go to step 2.

Theorem 8 *Suppose the condition of Theorem 4 or 6 is satisfied. Then*

(i) *If each strategy in Π^* has the same first L decisions, $\{\pi_1^*, \pi_2^*, \ldots, \pi_L^*\}$, then the algorithm will stop in finite time.*

(ii) *If the Algorithm stops at T_n, then $\pi_k^* = \bar{\pi}_k(T_n)$, $1 \leq k \leq L$, where π^* is the lexico minimum of Π^*.*

Proof:

(i) By hypothesis, $\bar{\Pi}(T_n) \to \Pi^*$. Then, from result R2 in Section 2, we have $\pi_k^n = \pi_k^*$, $1 \leq k \leq L$, for each $\pi^n \in \bar{\Pi}(T_n)$ and each $\pi \in \Pi^*$. Thus, since $\pi(T_n) \in \tilde{\Pi}(T_n) \subseteq \bar{\Pi}(T_n)$ and $\bar{\pi}(T_n) \in \bar{\Pi}(T_n)$, we have $\pi_k(T_n) = \bar{\pi}_k(T_n) = \pi_k^*$, $1 \leq k \leq L$.

(ii) For simplicity, we present the proof for $L = 1$. By the principle of optimality, $\pi^* \subseteq \bar{\Pi}(T_n)$. Let n^* be the node for π^* at or just beyond time T_n. Then $\bar{\pi}_1(T_n)$ and π_1^* both initiate paths that are optimal to n^*. By the definition of $\tilde{\Pi}(T_n)$, $\bar{\pi}_1(T_n) \leq \pi_1^*$. Now suppose $\bar{\pi}_1(T_n) < \pi_1^*$. Then a new strategy, $\hat{\pi}$, could be formed by following $\bar{\pi}(T_n)$ to n^* and continuing with π^* beyond n^*. Then $\hat{\pi} \in \Pi^*$ and $\hat{\pi} \prec \pi^*$. This contradicts the definition of π^*. ∎

Though uniqueness of the optimal initial decisions is sufficient for the algorithm to stop, it is not necessary. The next subsection gives an example in which the algorithm stops in the presence of infinitely many optimal strategies. Further, the minimal solution horizon may be shorter than that discovered by the Algorithm. We can distinguish between two types of horizons. A *forecast horizon* for the first L decisions is a finite horizon at which the algorithm's stopping rule is satisfied. It is called a forecast horizon since data beyond that time are irrelevant to the optimality of the first L decisions. A *solution horizon* for L is a time T such that if $T' > T$, then $\bar{\pi}_k(T') = \pi^*_k$, $1 \leq k \leq L$, for some $\pi^* \in \Pi^*$. From Theorem 8, the existence of a forecast horizon follows from uniqueness of the first L optimal decisions. From Corollary 7, a solution horizon exists regardless of uniqueness. For long enough T, the first L optimal decisions of $\bar{\pi}(T)$ agree with those of π^*. However, in practice the solution horizon may be discovered in retrospect, once a forecast horizon has been found. The minimal solution horizon for L is no longer than the corresponding forecast horizon. Due to the lexico minimum selection rule, whether it is shorter depends on the (arbitrary) assignment of indices to decisions. If leading optimal decisions have low indices, the solution horizon is likely to be relatively short; if they have high indices, the two horizons will be the same. Similarly, the extraction of $\tilde{\Pi}(T)$ from $\bar{\Pi}(T)$ also depends on the indices assigned to decisions. The minimal forecast horizon might be identified by testing the stopping rule for every possible numbering of initial decisions in each iteration. We did not implement this enhancement in our computations.

3.1 Regeneration Point Problems

The tie-breaking algorithm is simplified when applied to regeneration point problems. Recall that a decision epoch, T, is a regeneration point if the feasible decisions and costs beyond T are independent of the sequence of decisions up to time T. A problem formulation has regeneration point structure, or is a regeneration point problem, if the decision epochs for all strategies are regeneration points. Then the decision network may be aggregated so that there is at most one node for each point in time. Further, without loss of optimality, the network may be pruned so that $g(i) \leq g(j)$ for $T_i < T_j$. Examples include the discounted knapsack problem of Shapiro and Wagner and production planning problems with the Wagner-Whitin property.

Assume that the maximum time between any two nodes on the same path is bounded by τ. Then $\tilde{\Pi}(T) = \bigcup_{T' \in [T, T+\tau]} \Pi^*(T')$. Also, letting $\pi^*(T)$ be the lexico minimum strategy in $\Pi^*(T)$, it follows that $\tilde{\Pi}(T) = \{\pi^*(T'), T' \in [T, T+\tau]\}$. The tie-breaking algorithm stops when $(\pi_1^*(T'), \pi_2^*(T'), \ldots, \pi_L^*(T'))$ is the same for each $T' \in [T, T+\tau]$. This stopping rule is similar to that of Bean and Smith [1984], but will be satisfied more often due to the tie-breaking scheme.

Shapiro and Wagner showed that the infinite horizon discounted knapsack problem is solved by following a "turnpike" of least average cost decisions. If there is more than one such decision, then Π^* is the (infinite) set of all possible sequences of them. However, in the knapsack problem, one can show that $\Pi^*(T) \rightarrow \Pi^*$ (Ryan 1988). Then by Result R3 in Section 2, $\pi^*(T) \rightarrow \pi^*$, therefore $\pi_1^*(T) = \pi_1^*$ for all T sufficiently large, where $\pi^*(T)$ and π^* are the lexico minimum $T-$horizon and infinite horizon optimal solutions, respectively. Then $\pi_1^*(T') = \pi_1^*(T)$ for all T sufficiently large and $T' \in [T, T+\tau]$. Hence, for T sufficiently large, all solutions in $\tilde{\Pi}(T)$ have the same initial decision, so that the tie-breaking algorithm's stopping rule is satisfied for $L = 1$. In fact, the tie-breaking algorithm is a generalization of the Shapiro-Wagner algorithm.

4 Application to Capacity Expansion

We now apply the convergence results and algorithm to a general capacity expansion model. As in Bean and Smith [1985], suppose we are given a continuous demand function $D(t)$, which may be satisfied by any of a finite set of replicable facilities, indexed by $i = 1, \dots, n$. Facility i incurs a fixed installation cost F_i and provides X_i units of capacity. No undercapacity is allowed and costs are discounted. The case when $D(t) = dt$ is equivalent to the discounted knapsack problem.

We can assume without loss of optimality that a new facility is never deployed until all existing capacity is exhausted. Therefore, restrict Π to contain only strategies with this characteristic. The nodes in the decision network then represent installation epochs. Bean and Smith further argue that if $T_i < T_j$ and $g(i) > g(j)$ for some pair of nodes i and j, then i, resulting from a situation of low capacity at high cost, cannot be an installation epoch for an optimal strategy. Once all such dominated nodes are pruned from the network, the condition of Theorem 6 is immediately satisfied. Therefore $\bar{\Pi}(T_n) \rightarrow \Pi^*$ and $\bar{\pi}(T_n) \rightarrow \pi^*$.

4.1 Computational Experience

The Tie-Breaking Algorithm can be implemented as a traditional forward dynamic programming procedure, with little additional bookkeeping. We tested it using telephone capacity expansion problems from the literature, randomly generated problems, and a specially-designed example. Dominated nodes were pruned from the network as they were discovered by the dynamic programming procedure.

For the four telephone link capacity expansion examples in Bean and Smith [1985], the algorithm identified forecast horizons for the initial decision that were comparable to those previously computed. In addition, the first ten optimal decisions were discovered with forecast horizons of approximately 50 years or less. For the exponential demand example of Smith [1979], the forecast horizon for the first decision was eleven years, and the first nine decisions were uncovered with a 40-year horizon. Thus, for problems appearing in the literature, the tie-breaking algorithm performs as well as previous algorithms, but also uncovers more of the optimal strategy.

To further test the algorithm while simulating real capacity decision problems, we generated five random sets of nine facilities each. The capacities were generated uniformly between 0 and 100,000 units of capacity. Facility fixed costs were assigned according to the Dixon-Clapp relation (Yaged 1975):

$$F_i = KX_i^{1-\gamma},$$

where γ is an economy of scale factor, and K is a constant (set to 1 for convenience). For these randomly generated facilities, the indices assigned to decisions were unrelated to the relative sizes of the facilities. Thus, a lexico minimum selection is completely arbitrary. As noted previously, there is no way to predict whether solution horizons will be strictly shorter than the corresponding forecast horizons.

To construct the first type of demand function tested, we generated incremental demands randomly between 10,000 and 50,000 units. Thus the average facility of 50,000 units would exhaust in one to five years. Each set of facilities was tested twice, giving a total of ten test problems. The interest rate was set to 10.5 percent, as in the literature, and the economy of scale factor to 0.5. In all ten problems, the optimal solution (up to the first ten decisions) was to always install the largest (and eventually least average cost) facility.

Figure 1 shows forecast horizons in number of decisions implemented for each L, $1 \leq L \leq 10$. The best case, worst case, and average over all ten test problems are displayed for comparison. The L^{th} set of bars from the bottom gives the horizon length needed to discover the optimal L^{th} decision. For $L = 1$, the forecast horizon ranged between 6 and 12.5 years, with an average of 9 years. In order to normalize over the random facility sizes, we report the horizons in terms of the average number of installations required to reach them efficiently. That is, the horizontal axis gives, rather than T, the average number of decisions up to time T for strategies in $\tilde{\Pi}(T)$. An alternate way of interpreting the graph is that, for example, on average for $L = 5$, when eight installations are required to reach a given time horizon efficiently, the first 5 facilities installed by any lexico minimum efficient strategy are optimal. Notice that, after a delay of approximately four installations, the forecast horizons increase with a linear slope of nearly one with the number of decisions uncovered.

In some instances, solution horizons were considerably shorter than the corresponding forecast horizons. Figure 1 also displays solution horizons for the average case.

In order to judge the algorithm's performance when the optimal strategy contains a variety of decisions, we designed a second type of random demand function. We started with a sine wave with amplitude 50,000 and period 20 years. In an application such as along a telephone link, incremental demand may be negative as traffic is diverted to other parts of the network. To this cyclic pattern we added random annual increments of between 0 and 25,000 units. The interest rate was held at 10.5 percent, but the economy of scale factor changed to 0.3 to provide a wider gap between costs of large and small facilities. Once again, each facility set was tested with two demand functions for a total of ten test problems. In the optimal strategies, large facilities alternated with smaller ones in an unpredictable sequence.

Figure 2 shows the best, average, and worst case forecast horizons for the cyclic demand

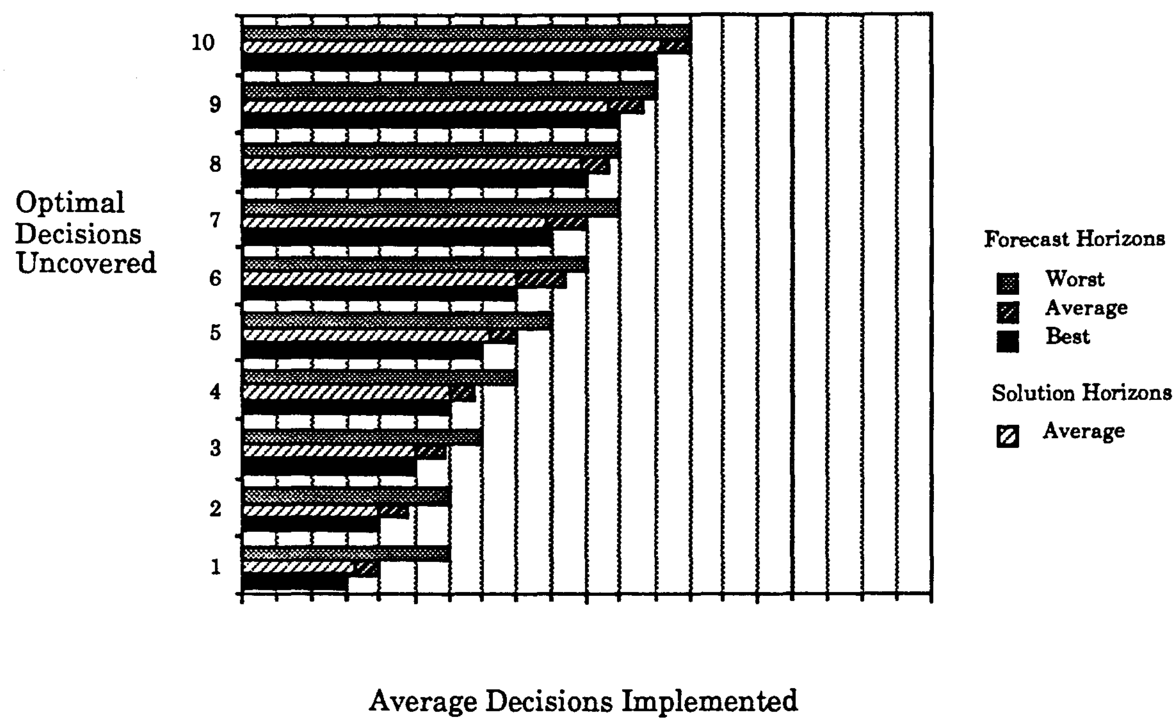

Figure 1: Forecast horizons for the linear demand function, measured by the average number of decisions implemented by efficient strategies. The best, average, and worst case forecast horizons are shown as well as solution horizons for the average case.

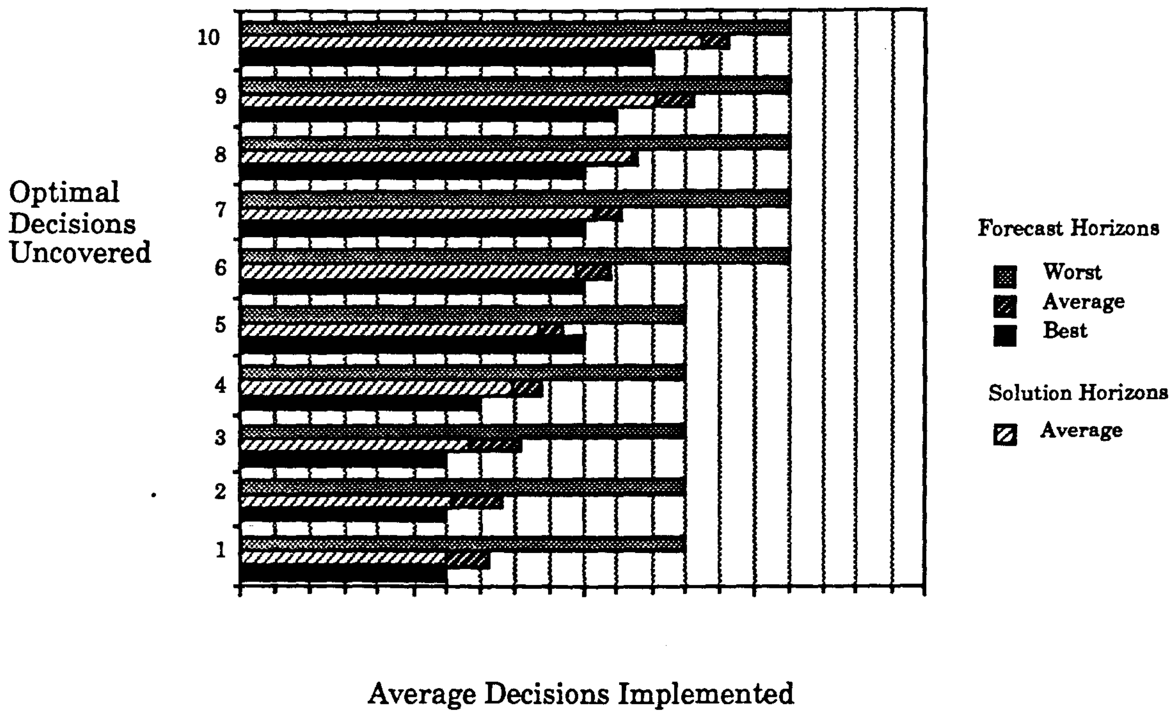

Figure 2: Forecast horizons for the cyclic demand function.

function, measured again by efficient facility installations. Though they are less smooth than for linear demand, they once again increase approximately linearly in the number of optimal decisions discovered, with a slope slightly more than one. The figure also gives a comparison of the corresponding solution and forecast horizons for the average case.

The unevenness of the successive forecast horizons can be attributed to the cyclic demand. As the horizon length, T, increased through years of declining demand, the sets $\tilde{\Pi}(T)$ typically remained unchanged. Subtracting capacity was not allowed, and there was no reason to add more. Then, as demand started to increase again, strategies in $\tilde{\Pi}(T)$ were re-evaluated. As demand grew rapidly, several successive optimal decisions might be determined all at once with a small increase in T.

The longest forecast horizons occurred with a facility set in which the capacities of the two largest facilities differed by less than 0.5 percent. Most of the computational effort was spent resolving ties between these two facilities, leading to solution and forecast horizons of nearly 50 years for the first decision. Such a situation would be unlikely in practice. For the other four facility sets, forecast horizons for $L = 1$ were between 7.5 and 28 years and the corresponding solution horizons were between 0 and 9 years.

We also tested the algorithm with an example specially constructed to defeat the stopping rule. In this example $D(t) = e^{0.1t} - 1$, a special case of exponential demand as studied by Smith [1979]. There are two facilities, with $X_1 = 2$, $F_1 = 1$, $X_2 \approx 0.1052$, and $F_2 \approx 0.3314$. The interest rate is set at 0.4. Let $\pi^1 = (1, 1, 1, \ldots)$ be the strategy of installing facility 1 indefinitely. Let $\pi^2 = (2, 1, 1, 1, \ldots)$ be the strategy of installing facility 2 once, then facility 1 indefinitely. One can show that both π^1 and π^2 are both infinite horizon optimal. Note that decision epochs for π^1 and π^2 never coincide: decision epochs for π^1 occur when $D(t) = X_2 n$ for some integer n and those for π^2 occur when $D(t) = 0$ or $X_1 + X_2 m$ for some integer m. We can also show that each strategy is optimal to its own decision epochs and nonoptimal to the other strategy's decision epochs. It follows that $\tilde{\Pi}_1(T) = \{1, 2\}$ for all T, and the tie-breaking algorithm's stopping rule is never satisfied. Details are contained in Ryan [1988].

However, the algorithm is still informative when applied to this example. The finite horizon efficient sets quickly settle into a pattern, allowing recognition that both optimal strategies are at least near-optimal. Moreover, the lexico minimum efficient strategies (underlined) converge to $(1, 1, \ldots)$. Efficient sets for selected finite horizons are shown in Table I. Notice that nonoptimal leading sequences soon disappear from $\tilde{\Pi}(T)$ as T increases.

$T=1$	$T=5$	$T=10$	$T=15$	$T=20$	$T=25$
21	2222221	2222221	21222221	2111221	21111122222222221
1	222221	222221	2122221	211121	2111112222222221
22	22221	22221	212221	21111	211111222222221
	2221	2221	21221	1111	21111122222221
	221	221	2121	2111222	2111112222221
	21	21	211		211111222221
	1	1	11		21111122221
					2111112221
					211111221
					21111121
					2111111
					111111

Table 1: Efficient sets, $\tilde{\Pi}(T)$, for an example where the stopping rule fails.

5 Conclusions

Solution horizon methods for solving infinite horizon problems have been burdened by the requirement of a unique optimum. As shown by Ryan and Bean, and illustrated by our example, such an assumption may not be true or easily verified. The tie-breaking algorithm is an efficient approach for finding solution horizons in the presence of multiple optima. When the stopping criterion is not met, the nature of the set convergence allows a decision maker to intelligently select so as to approximate them by solving finite horizon problems. In future research, we hope to develop bounds on the nearness to optimality of finite horizon efficient solutions. Such bounds will allow the derivation of near-solution horizons.

Acknowledgement

We would like to thank Greg Miller for computer implementation of the tie-breaking algorithm, funded under NSF Grant ECS-8700836.

References

Bean, J. and R. L. Smith. 1984. Conditions for the Existence of Planning Horizons. *Mathematics of Operations Research* **9**, 391-401.

Bean, J. and R. L. Smith. 1985. Optimal Capacity Expansion Over an Infinite Horizon. *Management Science* **31**, 1523-1532.

Bean, J. and R. L. Smith. 1986. Conditions for the Discovery of Solution Horizons. Technical Report 86-23, Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, Michigan 48109.

Bès, C. and S. Sethi. 1988. Concepts of Forecast and Decision Horizons: Applications to Dynamic Stochastic Optimization Problems. *Mathematics of Operations Research* **13**, 295-310.

Chand, S. and T. Morton. 1986. Minimal Forecast Horizon Procedures for Dynamic Lot Size Models. *Naval Research Logistics Quarterly* **33**, 111-122.

Hopp, W., J. Bean and R. L. Smith. 1987. A New Optimality Criterion for Nonhomogeneous Markov Decision Processes. *Operations Research* **35**, 875-883.

Lasserre, J. 1986. Detecting Planning Horizons in Deterministic Infinite Horizon Optimal Control. *IEEE Transactions on Automatic Control* **AC-31**, 70-72.

McKenzie, L. W. 1976. Turnpike Theory. *Econometrica* **44**, 841-844.

Ryan, S. M. 1988. Degeneracy in Discrete Infinite Horizon Optimization. Ph.D. dissertation, Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, Michigan.

Ryan, S. M. and J. Bean. 1987. Degeneracy in Infinite Horizon Optimization. Technical Report 87-5, Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, Michigan 48109-2117. To appear in *Mathematical Programming.*

Schochetman, I. and R. L. Smith. 1987. Infinite Horizon Optimization. Technical Report, Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, Michigan 48109-2117. To appear in *Mathematics of Operations Research.*

Shapiro, J. F. and H. M. Wagner. 1967. A Finite Renewal Algorithm for the Knapsack and Turnpike Models. *Operations Research* **15**, 319-341.

Skilton, D. K. 1985. Imbedding Posets in the Integers. *Order* **1**, 229-233.

Smith, R. L. 1979. Turnpike Results for Single Location Capacity Expansion. *Management Science* **25**, 474-484.

Yaged, B. 1975. Economies of Scale, Networks, and Network Cost Elasticity. *IEEE Transactions on Systems, Man and Cybernetics* **5**, 30-39.