

A TIE-BREAKING ALGORITHM FOR
INFINITE HORIZON OPTIMIZATION

Sarah M. Ryan
James C. Bean
Robert L. Smith

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

Technical Report 87-24
October, 1987

A Tie-Breaking Algorithm for Infinite Horizon Optimization

Sarah M. Ryan
James C. Bean
Robert L. Smith

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

October 6, 1987

1 Introduction

Dynamic sequential decision problems are an important class of optimization problems. Applications include production and inventory control, capacity expansion, and equipment replacement. In some cases the decision variables are continuous, allowing a mathematical programming formulation with a continuous solution space. Often, however, a problem calls for a model with discrete decisions as, for example, in the choice between the various pieces of equipment available to replace one currently in use. A great majority of these problems share the characteristic of an indefinite horizon. The economy, firm, or government involved has no predetermined moment of extinction. To appropriately model these problems, it is necessary to assume that the problem may continue indefinitely.

The traditional and most commonly used solution approach for such problems is to assume some finite horizon, T , and simply proceed as if the world ended there. The hope is that information beyond T will have little or no effect on the optimal solution, for at least the first few decisions, since those first decisions will be implemented immediately. A finite horizon with this guarantee is known as a *forecast horizon*. This solution method is known as a solution or forecast horizon approach (Lundin and Morton [1975]; Charnes, Drèze and Miller [1966]; Modigliani and Hohn [1955]).

Nearly all forecast horizon existence results require a unique optimal initial decision (see Bean and Smith [1984]; Hopp, Bean, and Smith [1984]; Bès and Sethi [1986], Schochetman and Smith [1987a]). Ryan and Bean [1987] show that in finite choice problems, such a requirement may be difficult to meet. In this paper, we allow that more than one optimal initial decision may exist, and develop an algorithm to select one such decision.

Forecast horizon approaches solve increasingly long finite horizon problems to approximate a sequence of decisions that is optimal over the infinite horizon. A common algorithm is to lengthen the horizon until the optimal initial decision stabilizes. When more than one optimal initial decision exists, it may happen that such a decision is optimal for some finite horizons but not others, no matter how long. In that case, the optimal initial decision may never stabilize. In this paper we describe conditions under which each

[†]This material is based on work supported by the National Science Foundation under Grants ECS-8409682 and ECS-8700836.

infinite horizon optimal initial decision is optimal for every sufficiently long finite horizon. For some types of problems, we can then select one optimal decision using a simple tie-breaking rule.

Section 2 includes a mathematical statement of the problem and assumptions. Section 3 is a discussion of the convergence of closed sets in a general compact space, and Section 4 discusses how closed set convergence can lead to an algorithm for infinite horizon optimization. In Sections 5 and 6 we present two alternate sets of conditions under which the sets of finite horizon optimal initial decisions converge to the set of infinite horizon optimal initial decisions. We illustrate the conditions with various capacity expansion models. Finally, Section 7 contains conclusions.

2 Problem Definition and Assumptions

Suppose we are faced with making a sequence of decisions over a continuous or discrete time frame. Each individual decision, denoted π_i , will be called a *policy*, and a sequence $\pi = (\pi_1, \pi_2, \dots)$ of policies constitutes a *strategy*. Let Π_n be the set of policies available after $n - 1$ decisions have been made. We assume policies may be indexed so that $\Pi_n \subset \{0, 1, 2, \dots\}$ and, for any π_n^1 and $\pi_n^2 \in \Pi_n$, $|\pi_n^1 - \pi_n^2| \leq M < \infty$, where M is uniform over n . Let $\Pi \subseteq \times_{n=1}^{\infty} \Pi_n$ be the set of all feasible strategies. Let $\theta \equiv \mathbf{0} = (0, 0, \dots)$. The strategy θ may or may not belong to Π . Associated with each strategy, π , is a cumulative net cost function $C_\pi(t)$. In order to compare costs incurred over time, we continuously discount them to time zero. Hence $\tilde{C}_\pi(r) = \int_0^\infty e^{-rt} dC_\pi(t)$ is the resulting infinite horizon discounted cost as a function of the discount rate, r . Also, $\tilde{C}_\pi(r, T) = \int_0^T e^{-rt} dC_\pi(t)$ is the discounted cost over the finite horizon T . Bean and Smith [1984] showed that infinite horizon discounted costs converge if r is larger than

$$\gamma = \sup_{\pi \in \Pi} \limsup_{t \rightarrow \infty} \frac{\ln |C_\pi(t)|}{t}.$$

The set $\{r | r > \gamma\}$ is called the *range of convergence* for interest rates.

We keep the assumptions of Bean and Smith [1984]; in particular:

$$\begin{aligned} C_\pi(t) &= \text{the cumulative net cost for strategy } \pi \text{ up to time } t \\ &= K_\pi(t) - R_\pi(t), \end{aligned}$$

$$\text{where } \left. \begin{array}{l} 0 \leq K_\pi(t) \leq B e^{\gamma t} \\ 0 \leq R_\pi(t) \leq B e^{\gamma t} \end{array} \right\} \forall t \geq T_0, \quad B, \gamma, T_0 \text{ independent of } \pi,$$

$$K_\pi(t), R_\pi(t) \text{ nondecreasing.}$$

The functions $K_\pi(t)$ and $R_\pi(t)$ represent cumulative pure cost and cumulative pure revenue, respectively. Without loss of generality, we can assume $T_0 = 0$ (if not, replace B by $B e^{\gamma T_0}$).

We define a metric similar to that of Bean and Smith [1984]: for $\pi, \hat{\pi} \in \Pi$,

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

where $\beta < \frac{1}{M+1}$. This metric has the property that, for any L , any two strategies which agree in the first L policies are closer than any two strategies which do not. This metric induces a topology on Π .

Given an interest rate, r , the problem we wish to solve is:

$$\min_{\pi \in \Pi} \tilde{C}_\pi(r).$$

As shown by Bean and Smith [1984], subject to an assumption that Π is complete, the minimum exists since Π is compact and $\tilde{C}_\pi(r)$ is continuous in π .

The forecast horizon approach involves solving finite horizon problems:

$$\min_{\pi \in \Pi} \tilde{C}_\pi(r, T).$$

In the remainder of the paper we assume r is fixed, and omit it from the notation. We define $\tilde{C}^*(T)$ to be the minimum finite horizon value and \tilde{C}^* to be the minimum infinite horizon value. A strategy $\hat{\pi}$ is termed *infinite horizon optimal* if it minimizes \tilde{C}_π and *finite horizon optimal* or *T-optimal* if it minimizes $\tilde{C}_\pi(T)$ for some T . Let Π^* and $\Pi^*(T)$ be the sets of optimal strategies for the infinite horizon and finite horizon problems, respectively. An *optimal initial policy* is an initial policy included in some optimal strategy. The sets of optimal initial policies for the infinite and finite horizon problems, respectively, are denoted Π_1^* and $\Pi_1^*(T)$.

A *specific forecast horizon* is a time, \hat{T} , such that for $T \geq \hat{T}$, $\Pi_1^*(T) = \Pi_1^*$, with $|\Pi_1^*| = 1$. A *general forecast horizon* is a time, \hat{T} , such that for $T \geq \hat{T}$, $\Pi_1^*(T) \subseteq \Pi_1^*$. Most forecast horizons found in the literature have been specific ones. In this paper, we seek a general forecast horizon.

3 Hausdorff Convergence

Bean and Smith [1984] derive a specific forecast horizon from the convergence of $\pi^*(T)$ to π^* , assuming π_1^* is unique. In this paper, we allow Π_1^* to contain more than one element and seek a general forecast horizon from the convergence of $\Pi^*(T)$ to Π^* . In this section we describe the nature of the convergence we seek for the optimal sets of strategies. First, we present results concerning the convergence of closed sets in a general compact space. The next definitions are adapted from Schochetman and Smith [1987b].

The space (Π, ρ) is a compact metric space. Let $\mathcal{C}(\Pi)$ be the set of all non-empty closed (hence compact) subsets of Π . Define

$$d(\pi, P) = \min_{\phi \in P} \rho(\pi, \phi), \text{ for } \pi \in \Pi, P \in \mathcal{C}(\Pi),$$

$$h(Q, P) = \max_{\pi \in Q} d(\pi, P), \text{ for } Q, P \in \mathcal{C}(\Pi).$$

$$D(Q, P) = \max(h(Q, P), h(P, Q)), \text{ for } Q, P \in \mathcal{C}(\Pi).$$

The function D is known as the Hausdorff metric on $\mathcal{C}(\Pi)$. The space $(\mathcal{C}(\Pi), D)$ is a compact metric space.

Let $\{P_i\}$ be a sequence in $\mathcal{C}(\Pi)$. As in Schochetman and Smith [1987b], define $\liminf P_i$ (resp. $\limsup P_i$) to be the set of all π in Π such that every neighborhood of π is *eventually* (resp. *frequently*) intersected by P_i . In general, $\liminf P_i \subseteq \limsup P_i$. A *selection* on $\mathcal{C}(\Pi)$ is defined to be a mapping $S : \mathcal{C}(\Pi) \mapsto \Pi$ such that $S(P) \in P$ for $P \in \mathcal{C}(\Pi)$. Let p be an arbitrary point in Π , and, for $P \in \mathcal{C}(\Pi)$, let $S_p(P)$ be any point $\pi \in P$ such that

$$\rho(p, \pi) = d(p, P).$$

The function $S_p(P)$ is called the *nearest-point selection defined by p*. If p is such that there exists a *unique* π in P as above, then p is a *uniqueness point* for P . In this case, π uniquely minimizes $\rho(p, \phi)$ over all $\phi \in P$. Let P^1 denote the *uniqueness set* for P , i.e., the set of all uniqueness points for P . If $p \in P^1$ and S_p is a nearest-point selection defined by p , then $S_p(P)$ is uniquely determined in P .

Schochetman and Smith prove the following theorem:

Theorem 1 *Let $\{P_i\}$ be a collection of closed sets in Π , and P be a closed set in Π . Then the following are equivalent:*

- (i) $P = \liminf_{i \rightarrow \infty} P_i = \limsup_{i \rightarrow \infty} P_i$.
- (ii) $\limsup_{i \rightarrow \infty} P_i \subseteq P$ and for each $\pi \in P$, $\exists \pi^i \in P_i$ for all i such that $\pi^i \rightarrow \pi$ as $i \rightarrow \infty$.
- (iii) $P_i \rightarrow P$ in the Hausdorff metric.
- (iv) $S_p(P_i) \rightarrow S_p(P)$ as $i \rightarrow \infty$ for all nearest-point selections S_p on $\mathcal{C}(\Pi)$ defined by $p \in P^1$.

To apply Theorem 1 to our problem, set P equal to Π^* , the set of infinite horizon optimal strategies, and let $P(T_i)$ be the set of strategies that are optimal in some sense (to be defined below) for a finite horizon T_i . Assume $T_i \rightarrow \infty$. In this context, $\limsup_{i \rightarrow \infty} P(T_i)$ is the set of algorithmically “optimal” strategies, or “optimal” strategies which can be identified by solving finite horizon problems. Whether or not they are truly optimal depends on the definition of $P(T_i)$. Results (i) and (ii) state that the algorithmically “optimal” strategies are optimal and, further, each optimal strategy can be approximated by solving *any* sequence of increasingly long finite horizon problems. Result (iii) states that finite horizon “optimal” sets converge to the infinite horizon optimal set. This result is similar to one found by McKenzie [1976]. Finally, result (iv) suggests an algorithm for choosing one out of the possibly many optimal strategies. For $p \in (\Pi^*)^1$, S_p selects a sequence of finite-horizon “optimal” strategies that converges to the corresponding infinite-horizon optimal strategy. From (iv), we can derive a forecast horizon by using the ρ -metric as in Bean and Smith [1984]. Convergence in this metric implies agreement in early policies. Therefore, since a convergent sequence of strategies eventually share the same initial policy, a forecast horizon can be identified by extending the sequence far enough.

4 Application to Infinite Horizon Optimization

In this section we lay some groundwork for applying Theorem 1 to infinite horizon optimization. We show that the finite horizon and infinite horizon sets of strategies involved in solving an infinite horizon problem are closed. Next we establish that the point $\theta \equiv \mathbf{0}$ is a uniqueness point for any closed set of strategies. Finally, we demonstrate that any set of strategies can be ordered according to their distance, measured by ρ , from θ . The ordering can be determined simply by examining the sequence of policies which compose each strategy.

Definition: *Let $P(T) \subseteq \Pi$. We say that $P(T)$ is a T -horizon set if membership in $P(T)$ is determined by data up to time T (inclusive).*

Lemma 2 *The sets Π^* and $P(T)$, where $P(T)$ is any T -horizon set, are closed.*

Proof: Π^* is closed since \tilde{C}_π is continuous in π and Π is compact. For $P(T)$, note that there are a finite number of partial strategies up to time T . If $P(T)$ is finite then it is closed. Suppose that $P(T)$ is infinite and that π is a cluster point of $P(T)$. Let $\{\pi^n\}_{n=1}^\infty$ be such that $\pi^n \in P(T)$ for all n and $\pi^n \rightarrow \pi$ 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 $\pi \in P(T)$. Thus $P(T)$ contains all its cluster points. Therefore $P(T)$ is closed. ■

Lemma 3 *For any closed set $P \subseteq \Pi$, $\theta \equiv \mathbf{0}$ is a uniqueness point for P , that is, $\theta \in P^1$.*

Proof: We will show that:

1. $\rho(\theta, \pi)$ attains its minimum on P ,

2. The function $\rho(\theta, \pi)$ of π is one-to-one.

(1) and (2) imply that $\rho(\theta, \pi)$ achieves its minimum at a unique point of P . Therefore, $\theta \in P^1$.

1. We begin by showing that $\rho(\theta, \pi)$ is continuous on Π . Suppose $\pi^n \rightarrow \pi$. By the triangle inequality,

$$\rho(\theta, \pi) \leq \rho(\theta, \pi^n) + \rho(\pi^n, \pi),$$

and

$$\rho(\theta, \pi^n) \leq \rho(\theta, \pi) + \rho(\pi^n, \pi).$$

Therefore,

$$|\rho(\theta, \pi) - \rho(\theta, \pi^n)| \leq \rho(\pi^n, \pi) \rightarrow 0.$$

Hence, $\rho(\theta, \pi)$ is continuous on Π . Now, since P is closed and $P \subseteq \Pi$, it follows that P is compact. Therefore, $\rho(\theta, \pi)$ attains its minimum on P . ■

2. Suppose $\pi^1, \pi^2 \in \Pi$ and $\pi^1 \neq \pi^2$. Then

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

Since $\pi^1 \neq \pi^2$, there exists n such that $\pi_n^1 \neq \pi_n^2$. Let n_0 be the smallest such n . Assume without loss of generality that $\pi_{n_0}^1 > \pi_{n_0}^2$. Then $\pi_{n_0}^1 - \pi_{n_0}^2 \geq 1$. By assumption, $\pi_n^1 - \pi_n^2 \geq -M$ for $n \geq n_0$. Then

$$\begin{aligned} \rho(\theta, \pi^1) - \rho(\theta, \pi^2) &= \sum_{n=n_0}^{\infty} \beta^n (\pi_n^1 - \pi_n^2) \\ &\geq \beta^{n_0} - M \sum_{n=n_0+1}^{\infty} \beta^n = \beta^{n_0} \left(1 - \frac{M\beta}{1-\beta} \right) > 0 \text{ since } \beta < \frac{1}{M+1}. \end{aligned}$$

Hence, $\rho(\theta, \pi)$ is one-to-one on Π . ■

Lemmas 2 and 3 imply that we can apply Theorem 1 to the question of whether any T -horizon set, $\mathcal{H}(T)$, converges to Π^* . We can use θ as a uniqueness point for a nearest-point selection. The next result interprets the selection $S_\theta(\cdot)$ as a simple tie-breaking rule, similar to one used for resolving degeneracy in linear programming (Murty [1983]). For finite choice infinite horizon optimization problems, we can define a lexicographic ordering of strategies, and use it to distinguish between alternate optimal strategies.

Definition: Let $a = (a_1, a_2, \dots)$ and $b = (b_1, b_2, \dots)$. 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 $a_{n_0} < b_{n_0}$.

Lemma 4 For any closed set $P \in \Pi$, $S_\theta(P)$ is the lexicographic minimum element of P .

Proof: We will show that $\rho(\theta, \pi^1) < \rho(\theta, \pi^2)$ if and only if $\pi^1 \prec \pi^2$. Suppose $\pi^1 \prec \pi^2$. Then for some n_0 we have

$$\pi_n^1 = \pi_n^2 \text{ for } n < n_0 \text{ and } \pi_{n_0}^1 < \pi_{n_0}^2.$$

Then

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

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

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

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

Thus the nearest-point selection defined by θ is simply the lexicographic minimum strategy. To summarize, we have:

Theorem 5 *If $\{P(T)\}$ is some collection of T -horizon sets, then $P(T) \rightarrow \Pi^*$ iff the lexicographic minimum element of $P(T)$ converges to the lexicographic minimum element of Π^* .*

Proof: Follows from Theorem 1 and Lemma 4. \blacksquare

Assuming $P(T) \rightarrow \Pi^*$, theorem 5 suggests the following algorithm: solve increasingly long finite horizon problems, identifying for each T , $P(T)$ and its lexicographic minimum element. The sequence of strategies thus generated converges to the lexicographic minimum infinite horizon optimal strategy.

5 Weakly Reachable Problems

Having established that Hausdorff convergence is equivalent to the existence of a tie-breaking algorithm, we would like to establish conditions under which Hausdorff convergence takes place. Then we can exploit the equivalence to derive an algorithm with a tie-breaking rule for identifying a forecast horizon and an optimal initial policy. To establish Hausdorff convergence, we define the finite horizon sets $\{P(T)\}$ and specify assumptions on the problem's structure. In this section, we give a rather broad definition for $P(T)$ and a condition under which $P(T) \rightarrow \Pi^*$. In the following section, we discuss a more restricted class of problems in which a tighter definition of $P(T)$ allows a stopping rule for the tie-breaking algorithm. In that section we also give a necessary and sufficient condition for Hausdorff convergence to take place.

We now turn from our focus on strategies to consider the dynamic programming network defined by the problem. In modeling a problem, we may consider states to be defined by strategies. In the extreme case, there may be a distinct state at time T for each different strategy up to that time. Often, however, more than one strategy will reach the same state at time T . If there is only one possible state at time T , that is, the state depends only on the time and not on the sequence of decisions up to that time, the state is called a *regeneration point*. Alternatively, we can define strategies in terms of states. For deterministic problems, strategies can be viewed as paths through the dynamic programming network, where a *path* is a sequence of states.

The concept of *reachability* of states has been used to prove forecast horizon results (McKenzie [1976], Lasserre, Bès, and Roubellat [1984], Bean and Smith [1986]). Loosely speaking, a state is reachable if it is not too far from an optimal path. Assume that states can be numbered so that states associated with later decision epochs have higher numbers. Let there be a root node for the dynamic programming network, representing the initial state. A strategy covering the infinite horizon can be represented by the infinite sequence of states it traverses from the root node through the infinite horizon, and its discounted cost represented as the length of that path. The following definitions are adapted from Bean and Smith [1986]:

Definition: $\bar{C}(s)$ = the shortest distance from the root node to node s .

$\bar{C}'(s)$ = the shortest distance from node s through the infinite horizon.

$\bar{C}(|s)$ = the length of a shortest path from the root node covering the infinite horizon which passes through node s .

The principle of optimality may be stated as:

$$\bar{C}(|s) = \bar{C}(s) + \bar{C}'(s).$$

Definition: A sequence of states $\{s_n\}_{n=1}^{\infty}$, $s_n \rightarrow \infty$, is weakly reachable from an optimal path if for all $\epsilon > 0$, there exists an N such that for $n \geq N$, there exists a $j(n)$ and an s'_n such that there is a path from $s_{j(n)}^*$ to s'_n at cost c_n with

$$\bar{C}(|s_n) \leq \bar{C}(|s'_n) + \epsilon,$$

where $s_{j(n)}^*$ is a state on some path $\{s_k^*\}$ optimal for the infinite horizon problem, $s_{j(n)}^* \rightarrow \infty$, and $c_n \rightarrow 0$ as $n \rightarrow \infty$.

Bean and Smith [1986] prove the following:

Theorem 6 $\bar{C}(s_i) \rightarrow \tilde{C}^*$ if and only if $\{s_i\}$ is weakly reachable.

To apply Theorem 1 to finite choice infinite horizon optimization, we need to define a finite horizon set of strategies for each finite horizon, T . Let $S(T)$ be some set of dynamic programming states at time T . Let

$$\bar{\Pi}^*(T) = \{\pi | \pi \text{ is optimal for some } s \in S(T) \text{ at time } T\}.$$

We would like to establish result (ii) of Theorem 1. We accomplish this in two steps.

Lemma 7 $\limsup \bar{\Pi}^*(T) \subseteq \Pi^*$ if and only if all infinite horizon state paths composed of states in $\{S(T)\}$ are weakly reachable.

Proof: By definition, $\pi \in \limsup \bar{\Pi}^*(T)$ if and only if there exists $\{T_i\}_{i=1}^{\infty}$, $T_i \rightarrow \infty$, and $\pi^i \in \bar{\Pi}^*(T_i)$ such that $\pi^i \rightarrow \pi$. Say π^i passes through state s_i at time T_i . By Theorem 4.5 of Schochetman and Smith [1987a], $\tilde{C}_{\pi^i}(T_i) \rightarrow \tilde{C}_{\pi}$. We have $\bar{C}(s_i) = \tilde{C}_{\pi^i}(T_i)$. Therefore, by Theorem 6, we have

$$\tilde{C}_{\pi^i}(T_i) = \bar{C}(s_i) \rightarrow \tilde{C}^* = \tilde{C}_{\pi}$$

if and only if $\{s_i\}$ is weakly reachable. ■

Lemma 8 If $S(T)$ is the set of all possible states at time T , then $\Pi^* \subseteq \liminf \bar{\Pi}^*(T)$.

Proof: Suppose $\pi \in \Pi^*$. Then, by the principle of optimality, π is optimal for its state at each time T . Thus, $\pi \in \bar{\Pi}^*(T)$ for all $T > 0$. Hence, there exists $\pi^T \in \bar{\Pi}^*(T)$ such that $\pi^T \rightarrow \pi$. Therefore, $\pi \in \liminf \bar{\Pi}^*(T)$. ■

We now have the following theorem.

Theorem 9 If $S(T)$ is the set of all possible states at time T and if all infinite horizon state paths composed of states in $\times S(T)$ are weakly reachable, then $\bar{\Pi}^*(T) \rightarrow \Pi^*$.

5.1 Capacity Expansion Illustration

In this section we give an example in which all dynamic programming states are weakly reachable. The problem is therefore amenable to solution by the algorithm suggested by Theorem 5.

Consider the problem of optimally deploying facilities to meet growing demands for capacity. As in Bean and Smith [1985], assume we are given a cumulative continuous demand, $D(t)$, for capacity at time t . Demand can be satisfied by any of a finite set of distinct replicable facility types indexed $1, 2, \dots, n$. Installing facility j incurs a discrete fixed cost $F_j > 0$ and provides X_j units of capacity. Assume that $X_1 < X_2 < \dots < X_n$ and, without loss of optimality, that $F_1 < F_2 < \dots < F_n$. No undercapacity is allowed, and we assume for convenience that variable operating costs are unaffected by the choice of facilities. The problem is to choose an optimal sequence of facilities to meet the demand for capacity over an infinite horizon.

We can assume without loss of optimality that a new facility is never deployed until all existing capacity is exhausted. A point in time when it is necessary to add capacity is called an *installation epoch*. Bean and Smith [1985] show that, for this problem, installation epochs are regeneration points. We define a dynamic programming state to include the time, T , and the amount of excess capacity at T . At an installation epoch, excess capacity equals zero.

The cumulative cost function for any strategy is a step function with a discontinuity at each installation epoch. Assume this function is left-continuous so that the fixed cost for a facility is added just after its installation epoch. Also note that if T_1 and T_2 are two installation epochs (possibly for different strategies) and $T_1 < T_2$, then the optimal cost for T_1 is less than or equal to the optimal cost for T_2 . Otherwise, the cumulative demand up to T_1 could be satisfied by the decision sequence leading up to T_2 at lower cost. Finally, assume $D(t)$ is continuous, nonnegative, and eventually bounded by an exponential function with rate g . Bean and Smith [1985] prove that if costs are discounted at a rate $r > g$, then infinite horizon discounted costs are finite.

Lemma 10 *If states are defined as $(T, \text{excess capacity at } T)$, then all paths are weakly reachable.*

Proof: Choose a state path $\{s_n\}_{n=1}^{\infty}$. Let N be such that for $n \geq N$, $\bar{C}'(s_n) \leq \epsilon$ (N exists for interest rates $r > g$). Then $\bar{C}'(s) \leq \epsilon$ for all s at times beyond T where T is the time associated with s_n . Let t_n be the last installation epoch leading to s_n and t_{n+1} be the first installation epoch after s_n on a strategy that goes through s_n . Given some optimal path $\{s_n^*\}$, let

$$j(n) = \max\{j | T_{s_j^*} < t_n, s_j^* \text{ an installation epoch}\}$$

$$s'_n = \min\{s_n^* | T_{s_n^*} < t_{n+1}, s_n^* \text{ an installation epoch}\}$$

(See Figure 1). Then

$$c_n = \text{cost}(s_{j(n)}^*, s'_n) \rightarrow 0$$

by the convergence of f . Hence

$$\bar{C}(|s_n) - \bar{C}(|s'_n) = \bar{C}(s_n) - \bar{C}(s'_n) + \bar{C}'(s_n) - \bar{C}'(s'_n).$$

But

$$\bar{C}(s_n) - \bar{C}(s'_n) \leq 0 \text{ since } \bar{C}(s_n) \leq \bar{C}(t_{n+1}) \leq \bar{C}(s'_n),$$

and

$$\bar{C}'(s_n) - \bar{C}'(s'_n) < \epsilon \text{ by the choice of } N.$$

Therefore $\{s_n\}$ is weakly reachable. ■

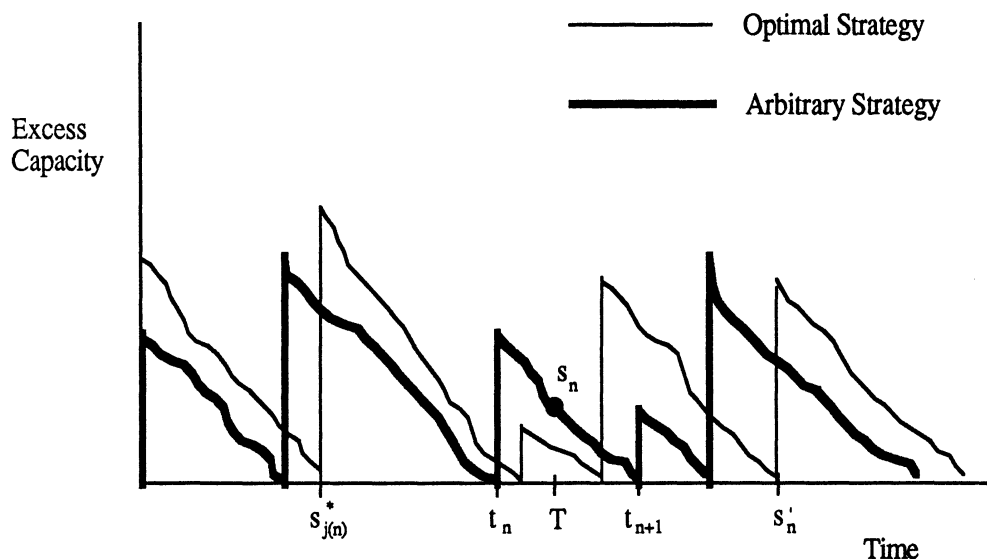


Figure 1. Weak Reachability in Capacity Expansion Problems.

Bean and Smith [1985] present an algorithm for identifying a specific forecast horizon for this problem assuming nondegeneracy. By Theorem 5 and Lemma 10, we have an approach for identifying a general forecast horizon even in the presence of more than one optimal initial policy. In this instance, weak reachability follows from the regeneration point properties of optimal strategies. We first had to establish, by insight into the problem's structure, that it is sufficient to consider only strategies in which capacity is added only when all existing capacity is used up. A weak reachability result could be established similarly in a Wagner-Whitin type of production planning problem, in which production occurs only when inventory is zero. We expect that proving weak reachability for other problems will require some measure of insight into special problem characteristics.

6 Regeneration Point Problems

In this section we address Hausdorff convergence for problems with regeneration point structure. We give an algorithm with a stopping rule which is guaranteed to settle on an infinite horizon optimal initial policy in finite time. This is the first implementable algorithm guaranteed to converge in finite time without the restriction that π_1^* be unique. Next we present a necessary and sufficient condition for Hausdorff convergence to occur in regeneration point problems. This condition is a type of reachability stronger than those defined by Bean and Smith [1986] and McKenzie [1976]. Finally, we illustrate the reachability condition by means of two examples. The knapsack problem solved by Shapiro and Wagner [1967] satisfies the reachability condition. The Shapiro-Wagner algorithm is actually a special case of our tie-breaking algorithm. We also present an example of capacity expansion under exponential demand which has regeneration points but fails to pass the reachability test.

Recall that a decision epoch is a regeneration point if the optimal strategy beyond T is independent of the sequence of policies up to time T . A problem has regeneration point structure, or is a regeneration point problem, if the decision epochs for optimal policies are regeneration points. We can draw two conclusions which will be useful for proving results in this section. First, suppose T is a regeneration point for $\pi^* \in \Pi^*$. By the principle of optimality, π^* is optimal to its state at time T and optimal from its state at time T through the infinite horizon. It follows that (i) any strategy truncation ending at T which is T -optimal can

be continued by π^* beyond T . Otherwise, since optimality implies feasibility, the optimal strategy beyond T depends on the policies before T . Also, (ii) $\pi^* \in \Pi^*(T)$. If not, there is a less costly sequence of policies up to time T which can be continued by π^* beyond T . This strategy has lower cost than π^* .

6.1 Tie-Breaking Algorithm

For regeneration point problems, we define the finite horizon set $P(T)$ for each T to be $\Pi^*(T)$, the set of T -optimal strategies. Assume that $\Pi^*(T) \rightarrow \Pi^*$ in the Hausdorff metric. The following algorithm can be used to solve any regeneration point problem. Let τ be the maximum duration of any policy.

Algorithm

1. Let $\{T_n\}_{n=1}^{\infty}$ be the sequence composed of all decision epochs for all strategies. Set $n \leftarrow 1$.
2. Solve the T_n -horizon problem to obtain $[S_{\theta}(\Pi^*(T_n))]_1$, the initial policy of the lexico minimum T_n -optimal strategy.
3. If $[S_{\theta}(\Pi^*(T_n))]_1 = [S_{\theta}(\Pi^*(T_k))]_1$ for all k such that $T_n - \tau \leq T_k \leq T_n$, then stop. Otherwise, set $n \leftarrow n + 1$ and go to step 2.

Theorem 11 *If the algorithm is applied to a regeneration point problem then*

- (i) *The algorithm will stop in finite time.*
- (ii) *When the algorithm stops, $[S_{\theta}(\Pi^*(T_n))]_1 = [S_{\theta}(\Pi^*)]_1$.*

Thus, the algorithm selects an optimal initial policy in finite time.

Proof (i): Since $\Pi^*(T) \rightarrow \Pi^*$, we know that $S_{\theta}(\Pi^*(T)) \rightarrow S_{\theta}(\Pi^*)$. Hence there exists T^* such that for $T \geq T^*$, $\rho(S_{\theta}(\Pi^*(T)), S_{\theta}(\Pi^*)) < \beta$. Therefore, for $T \geq T^*$, $[S_{\theta}(\Pi^*(T))]_1 = [S_{\theta}(\Pi^*)]_1$.

(ii): Suppose not, i.e., suppose the algorithm stops at the finite horizon T_n and

$$[S_{\theta}(\Pi^*(T_n))]_1 = [S_{\theta}(\Pi^*(T_k))]_1 \text{ for all } T_n - \tau \leq T_k \leq T_n,$$

but

$$[S_{\theta}(\Pi^*)]_1 \neq [S_{\theta}(\Pi^*(T_n))]_1.$$

For any $\pi^* \in \Pi^*$, π^* has a decision epoch in $[T_n - \tau, T_n]$. Call it T' . Now, $\Pi_1^*(T') \subseteq \Pi_1^*$ by the regeneration point property. Then $[S_{\theta}(\Pi^*(T'))]_1 \in \Pi_1^*$. Hence $[S_{\theta}(\Pi^*(T_n))]_1 \in \Pi_1^*$. Then, since $[S_{\theta}(\Pi^*)]_1 \neq [S_{\theta}(\Pi^*(T_n))]_1$, we have

$$[S_{\theta}(\Pi^*)]_1 \prec [S_{\theta}(\Pi^*(T_n))]_1.$$

But $S_{\theta}(\Pi^*)$ has a decision epoch $T'' \in [T_n - \tau, T_n]$. Then $S_{\theta}(\Pi^*) \in \Pi^*(T'')$. Therefore, $S_{\theta}(\Pi^*) \not\prec S_{\theta}(\Pi^*(T''))$. Hence,

$$[S_{\theta}(\Pi^*)]_1 \not\prec [S_{\theta}(\Pi^*(T''))]_1 = [S_{\theta}(\Pi^*(T_n))]_1.$$

This is a contradiction. ■

6.2 Necessary and Sufficient Condition for Hausdorff Convergence

Let ${}_L\Pi^*$ be the set of truncations of strategies in Π^* to the first L policies. Let ${}_L\Pi^*(T)$ be the analogous set for $\Pi^*(T)$. Since ${}_L\Pi^*$ and ${}_L\Pi^*(T)$ are finite, they are compact. To prove the main result in this section, we need the following lemma.

Lemma 12 $\Pi^* \subseteq \liminf \Pi^*(T) \iff$ for every L there exists T_L such that for $T \geq T_L$, ${}_L\Pi^* \subseteq {}_L\Pi^*(T)$.

Proof: Suppose $\Pi^* \subseteq \liminf \Pi^*(T)$. Then for each $\pi \in \Pi^*$, for each T , there exists $\pi^T \in \Pi^*(T)$ such that $\pi^T \rightarrow \pi$. Choose L and consider the truncation ${}_L\pi$ of $\pi \in \Pi^*$. Since $\pi^T \rightarrow \pi$, there exists T_L such that for $T \geq T_L$, ${}_L\pi^T = {}_L\pi$. Thus ${}_L\pi \in {}_L\Pi^*(T)$. But ${}_L\pi$ was an arbitrary element of ${}_L\Pi^*$.

Now suppose $\Pi^* \not\subseteq \liminf \Pi^*(T)$. Then for some $\pi \in \Pi^*$, for some $\epsilon > 0$, for every T , there exists $T' \geq T$ such that $\rho(\pi, \pi^{T'}) > \epsilon$ for each $\pi^{T'} \in \Pi^*(T')$. Then for some K , for every T , there exists $T' \geq T$ such that for every $\pi^{T'} \in \Pi^*(T')$, $\pi_n^{T'} \neq \pi_n$ for some $n \leq K$. In other words, for some K , for every T , there exists $T' \geq T$ such that ${}_K\pi \notin {}_K\Pi^*(T')$. The result holds by contraposition. ■

Our necessary and sufficient condition stipulates that for any pair of infinite horizon optimal strategies, there exists, indefinitely far into the future, a sequence of policies that connect the two strategies optimally. That is, one can follow the first strategy up to an arbitrary decision epoch, then implement a sequence of policies to arrive with optimal cost at a decision of the second strategy. A rigorous definition follows. The terminology “ T on π ” means T is a decision epoch for strategy π .

Definition: A strategy π^1 is optimally connected to a strategy π^2 if for any T on π^1 , there exists $T' \geq T$ on π^1 and S on π^2 , $S \geq T$, such that there exists a sequence of policies $(\hat{\pi}_1, \hat{\pi}_2, \dots, \hat{\pi}_n)$ with $\tilde{C}_{\pi^2}(S) = \tilde{C}_{\pi^1}(T') + \{\text{discounted cost of } (\hat{\pi}_1, \hat{\pi}_2, \dots, \hat{\pi}_n) \text{ when implemented at time } T'\}$, where $(\hat{\pi}_1, \hat{\pi}_2, \dots, \hat{\pi}_n)$ continues at least up to S when implemented at T' .

We assume that, in the event that $(\hat{\pi}_1, \hat{\pi}_2, \dots, \hat{\pi}_n)$ continue beyond S , the policies of π^2 which would be implemented starting at S may still be implemented at S . An example is capacity expansion (Bean and Smith [1985]), where $(\hat{\pi}_1, \hat{\pi}_2, \dots, \hat{\pi}_n)$ may install more capacity than is actually needed up to time S . In that case, we assume the excess capacity is ignored, and a new facility can be installed at time S . We say that all strategies are optimally connected if every strategy is optimally connected to every other. The following is a statement of our necessary and sufficient condition for Hausdorff convergence.

Theorem 13 For a regeneration point problem, $\Pi^*(T) \rightarrow \Pi^*$ if and only if all strategies in Π^* are optimally connected.

Proof: As shown by Schochetman and Smith [1987b], $\limsup \Pi^*(T) \subseteq \Pi^*$. We will show that optimal connectedness is necessary and sufficient for $\Pi^* \subseteq \liminf \Pi^*(T)$.

Suppose $\Pi^* \subseteq \liminf \Pi^*(T)$. Then by Lemma 12, for any $\pi^* \in \Pi^*$ and any L , ${}_L\pi^* \in {}_L\Pi^*(T)$ for all T sufficiently large. In particular, there exists T_L^* such that for any $\hat{\pi} \in \Pi^*$, ${}_L\Pi^* \in {}_L\Pi^*(\hat{T}_i)$ where \hat{T}_i is a decision epoch for $\hat{\pi}$, $\hat{T}_i > T_L^*$. Now suppose that not all strategies are optimally connected. Then for some π^* and $\hat{\pi} \in \Pi^*$, for some T on π^* , for every $T' \geq T$ on π^* and every $S \geq T$ on $\hat{\pi}$, $\tilde{C}_{\hat{\pi}}(S) < \tilde{C}_{\pi^*}(T') + \{\text{cost of any sequence of policies starting at } T' \text{ and continuing at least up to } S.\}$ Suppose T is the $(m+1)$ st decision epoch. Then ${}_m\pi^* \notin \Pi^*(\hat{T}_i)$ where \hat{T}_i is a decision epoch for $\hat{\pi}$, $\hat{T}_i \geq T$. This is a contradiction.

Now suppose every optimal strategy is optimally connected. We will show that for every L , there exists T_L such that for $T \geq T_L$, ${}_L\Pi^* \subseteq {}_L\Pi^*(T)$. Then $\Pi^*(T) \rightarrow \Pi^*$ by Lemma 12.

Suppose not. Then for some L , there exists $\{T_i^L\}_{i=1}^\infty$, $T_i \rightarrow \infty$, such that ${}_L\Pi^* \not\subseteq {}_L\Pi^*(T_i^L)$, i.e., for each i , there exists ${}_L\pi^i \in {}_L\Pi^*$ such that ${}_L\pi^i \notin {}_L\Pi^*(T_i^L)$. By compactness, there exists a convergent

subsequence of $\{ {}_L\pi^i \}$, call it $\{ {}_L\pi^j \}$, such that ${}_L\pi^j \rightarrow {}_L\pi$. That is, ${}_L\pi^j = {}_L\pi$ for j sufficiently large and ${}_L\pi \in {}_L\Pi^*$. Let $\tilde{\pi}$ be some continuation of ${}_L\pi$ which is contained in Π^* .

For each j choose an element of $\Pi^*(T_j^L)$, call it $\hat{\pi}^j$. There exists a convergent subsequence $\{\hat{\pi}^k\}$ of $\{\hat{\pi}^j\}$ such that $\hat{\pi}^k \rightarrow \hat{\pi} \in \Pi^*$. We have ${}_L\hat{\pi}^k = {}_L\hat{\pi}$ for k sufficiently large. Summarizing, we have,

$${}_L\tilde{\pi} \in {}_L\Pi^*, \quad {}_L\tilde{\pi} \notin {}_L\Pi^*(T_j^L), \quad \forall j,$$

and

$${}_L\hat{\pi} \in {}_L\Pi^*, \quad {}_L\hat{\pi} \in {}_L\Pi^*(T_k^L), \quad \forall k,$$

where $\{T_k^L\} \subseteq \{T_j^L\}$. There exists some \tilde{T}_L on $\tilde{\pi}$ such that L policies have been implemented by $\tilde{\pi}$ before \tilde{T}_L . By assumption, there exists \tilde{T}' on $\tilde{\pi}$, $\tilde{T}' \geq \tilde{T}_L$, and \hat{T} on $\hat{\pi}$, $\hat{T} \geq \tilde{T}'$, and a sequence of policies $(\pi_1^*, \pi_2^*, \dots, \pi_n^*)$, which, when implemented at \tilde{T}' , extend at least up to \hat{T} and satisfy $\tilde{C}_{\tilde{\pi}}(\hat{T}) = \tilde{C}_{\tilde{\pi}}(\tilde{T}') + \{\text{cost of } (\pi_1^*, \pi_2^*, \dots, \pi_n^*) \text{ when implemented at } \tilde{T}'\}$. We can form a new strategy: follow $\tilde{\pi}$ up to \tilde{T}' , then $(\pi_1^*, \pi_2^*, \dots, \pi_n^*)$ between \tilde{T}' and \hat{T} , then the policies of $\hat{\pi}$ from \hat{T} up to

$$\bar{T} \equiv \min\{T_k^L | k \text{ such that } T_k^L \geq \hat{T}\}.$$

This strategy is optimal for \bar{T} . Then ${}_L\hat{\pi} \in {}_L\Pi^*(\bar{T})$, which is a contradiction. ■

6.3 Examples

The knapsack problem solved by Shapiro and Wagner [1967] is one example in which a forecast horizon exists regardless of degeneracy. Their paper gives an algorithm together with a stopping rule for solving infinite horizon problems which is guaranteed to be satisfied in finite time. They make no assumption concerning the number of optimal initial policies. In fact, the knapsack problem is an example of Hausdorff convergence in a regeneration point problem and the Shapiro-Wagner algorithm is a special case of our tie-breaking algorithm. We can show, by extending the results of Shapiro and Wagner, that the knapsack problem satisfies the necessary and sufficient condition for Hausdorff convergence. Therefore, the tie-breaking algorithm is guaranteed to converge.

First, we define the problem. Some of the notation and conventions from Shapiro and Wagner [1967] are altered to harmonize with the rest of this paper. In particular, we reverse the policy indexing order. At any point in time, the same set of n policies, $\{u_1, u_2, \dots, u_n\}$, are available. Policy j incurs cost F_j and has duration τ_j . Expressed as a dynamic programming recursion, the finite horizon problem is

$$G(T) = \min_{\{j=1,2,\dots,n | \tau_j \leq T\}} F_j + \alpha^{\tau_j} G(T - \tau_j),$$

where $\alpha = e^{-r}$ and $G(0) = 0$. Let

$$\gamma_j = \begin{cases} \frac{(1-\alpha)F_j}{1-\alpha^{\tau_j}}, & \alpha < 1 \\ \frac{F_j}{\tau_j}, & \alpha = 1 \end{cases}$$

The infinite horizon problem can be expressed as

$$g \equiv (1-\alpha)G = \min_{j=1,2,\dots,n} \gamma_j.$$

Any policy minimizing γ_j is called a turnpike policy. Then Π^* is the set of turnpike strategies, or the set of all possible sequences of turnpike policies. Let $\{u_1, u_2, \dots, u_i\}$ be the set of turnpike policies. Assume policies are indexed in order of increasing γ_j . In the event of a tie, the lowest index is given to the policy with the smallest τ_j . Thus, u_1 is the turnpike policy with the smallest duration.

As in Shapiro and Wagner [1967], let a_j and b_j be relatively prime positive integers such that

$$a_j \tau_j = b_j \tau_1, \text{ for } j = 1, 2, \dots, t.$$

Let M^* be the smallest integer such that

$$\frac{F_1}{1 - \alpha^{\tau_1}} (1 - \alpha^{L - \tau_1 + 1}) < \frac{F_{k^*}}{1 - \alpha^{\tau_{k^*}}} (1 - \alpha^L)$$

for all $L > M^*$, where k^* is such that

$$\frac{F_{k^*}}{1 - \alpha^{\tau_{k^*}}} = \min_j \left\{ \frac{F_j}{1 - \alpha^{\tau_j}} \mid \frac{F_j}{1 - \alpha^{\tau_j}} > \frac{F_1}{1 - \alpha^{\tau_1}} \right\}.$$

Let

$$N^* = \sum_{j=2}^t (a_j - 1) \tau_j$$

and

$$L^* = \sum_{j=2}^t a_j \tau_j.$$

Shapiro and Wagner show that for $T > M^* + N^*$, there exists a strategy π in $\Pi^*(T)$ such that $\pi_1 = u_1$. Therefore, for any $T > M^* + N^*$, there exists a strategy $\pi \in \Pi^*(T)$ such that $\pi_i = u_1$ up to some epoch in $[T - M^* - N^*, T - M^* - N^* + \tau_1)$.

Lemma 14 *In the knapsack problem, all infinite horizon optimal strategies are optimally connected.*

Proof: Consider any π^1 and π^2 in Π^* . We will show that π^1 and π^2 are optimally connected. Choose some decision epoch T for π^1 . Up to time T , π^1 includes N_T^j replications of policy u_j , for $j = 1, \dots, t$. Let

$$m = \min\{i \geq 0, \text{ integer} \mid i a_j \geq N_T^j, \forall j \in \{1, \dots, t\}\}.$$

Let S be a decision epoch for π^2 , $S > M^* + N^* + mL^*$. Then there exists a strategy $\hat{\pi} \in \Pi^*(S)$ with $\hat{\pi}_i = u_1$ up to some $S' \in [S - M^* - N^*, S - M^* - N^* + \tau_1)$. We have $S' > mL^*$. In strategy $\hat{\pi}$, we can replace mb_j replications of u_1 by ma_j replications of u_j with no difference in cost. In so doing, we include at least N_T^j replications of u_j , $j = 1, 2, \dots, t$. Rearrange the initial policies so that they agree with π^1 up to time T . The resulting strategy is optimal for S and has a decision epoch at T . Its policies between T and S form an optimal connecting path between π^1 and π^2 . Therefore π^1 and π^2 are optimally connected. ■

The Shapiro-Wagner algorithm solves increasingly long finite horizon problems, identifying an optimal initial policy for each horizon. If two policies tie, the one with the smaller index is chosen. The algorithm terminates at T_S if the optimal initial policy is u_1 for all $T \in [T_S - \tau, T_S]$ where $\bar{\tau} = \max_{j=1,2,\dots,n} \tau_j$. Thus their algorithm is identical to our tie-breaking algorithm.

One instance of the knapsack problem is capacity expansion with linear demand. The policies represent available facilities to augment capacity. The problem is to choose an optimal sequence of capacity additions, subject to the constraint that cumulative demand for capacity up to time t , $D(t) = dt$, be satisfied for all t . Facility j incurs a fixed cost F_j and adds $X_j = \tau_j d$ units of capacity.

An extension of this problem is the case where $D(t) = D_0 e^{gt}$, as in Smith [1979]. The extension preserves the regeneration point property. In the case of exponential demand, however, Hausdorff convergence is not guaranteed. In the following example, Hausdorff convergence fails to take place.

Nonlinear $D(t)$ implies that policy durations are nonstationary. Smith shows that for exponential demand,

$$\tau_j(t) = \frac{1}{g} \log \left(1 + \frac{X_j}{D_0 e^{gt}} \right).$$

is the duration of policy j when installed at time t . Also, the cost of installing facility j forever starting at time t is

$$C_j(t) = e^{-rt} \sum_{n=0}^{\infty} F_j \left(1 + \frac{nX_j}{D_0 e^{gt}} \right)^{-r/g},$$

where r is the interest rate used to discount costs. Let

$$\gamma_j(t) = \frac{rF_j}{1 - e^{-r\tau_j(t)}}.$$

This is a nonstationary version of the γ_j defined by Shapiro and Wagner [1967].

Example: Let facility 1 add capacity $X_1 = 1$ and incur fixed cost $F_1 = 2$. Facility 2 adds capacity $X_2 \approx 0.1052$ and incurs fixed cost $F_2 \approx 0.3314$. The values of X_2 and F_2 are designed so that $\tau_2(0) = 1$ and $F_2 = C_1(0) - C_1(1)$. Let $\pi^1 = (1, 1, 1, \dots)$ be the strategy of installing facility 1 indefinitely. Let $\pi^2 = (2, 1, 1, 1, \dots)$ be the strategy of installing facility 2 once, then facility 1 indefinitely. Then π^1 and π^2 have the same infinite horizon discounted cost. See Figure 2. Also, one can verify that $\gamma_2(\tau_2(0)) > \gamma_1(\tau_2(0))$. By Theorem 3 of Smith [1979], it follows that π^1 and π^2 are superior to all other strategies. 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 claim that each strategy is optimal to its own decision epochs and suboptimal to the other strategy's decision epochs. It follows that $\limsup_{T \rightarrow \infty} \Pi^*(T) = \{\pi^1, \pi^2\} = \Pi^*$ but $\liminf_{T \rightarrow \infty} \Pi^*(T) = \emptyset$. Therefore, by Theorem 1, $\Pi^*(T) \not\rightarrow \Pi^*$.

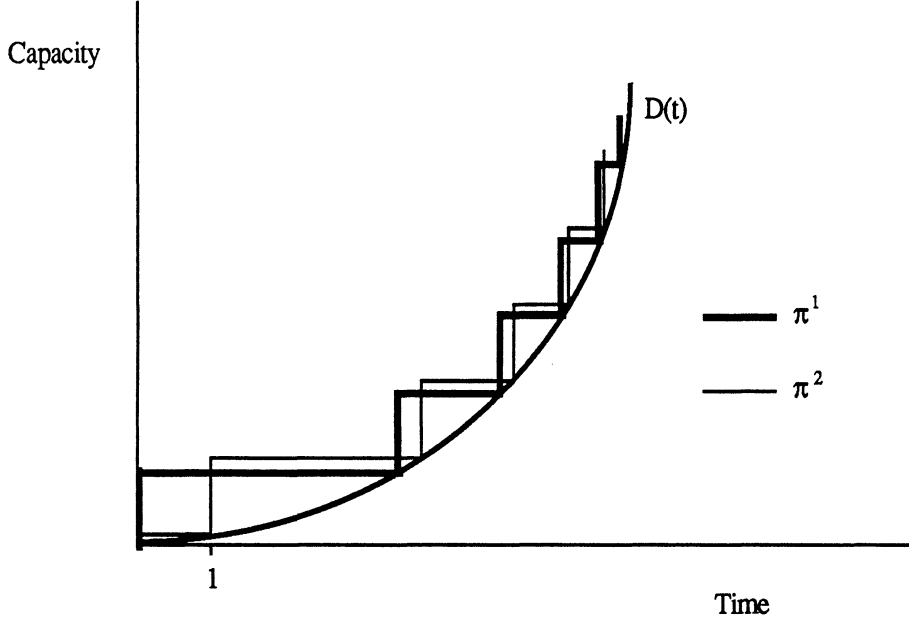


Figure 2. Example of Capacity Expansion under Exponential Demand with no Hausdorff Convergence.

Proof of Claim: At a decision epoch T for π^1 , $D(T) = X_1 k$ for some k . Since $X_2 < X_1$, the cumulative capacity installed by π^2 up to time T must be at least $X_2 + X_1 k$. Let T_n^1 be the installation epoch for the $(n + 1)^{\text{st}}$ facility of π^1 . Then

$$\begin{aligned}\tilde{C}_{\pi^1}(T) &= \sum_{n=0}^{k-1} F_1 e^{-rT_n^1} \\ &= F_1 \sum_{n=0}^{k-1} e^{-rT_0^1} \left(1 + \frac{nX_1}{D_0 e^{gT_0^1}}\right)^{-r/g} \\ &= F_1 \sum_{n=0}^{k-1} \left(1 + \frac{nX_1}{D_0}\right)^{-r/g}\end{aligned}$$

and

$$\begin{aligned}\tilde{C}_{\pi^2}(T) &= F_2 + F_1 e^{-r} \sum_{n=0}^{k-1} \left(1 + \frac{nX_1}{D_0 e^{g}}\right)^{-r/g} \\ &= F_2 + F_1 \sum_{n=0}^{k-1} \left(e^g + \frac{nX_1}{D_0}\right)^{-r/g}\end{aligned}$$

By algebra, $\tilde{C}_{\pi^1}(T) < \tilde{C}_{\pi^2}(T)$.

Similarly, up to a decision epoch S for π^2 , π^1 has installed facility $k + 1$ replications of facility 1 and

π^2 has installed one replication of facility 2 and k of facility 1. Then

$$\tilde{C}_{\pi^1}(S) = F_1 \sum_{n=0}^k \left(1 + \frac{nX_1}{D_0}\right)^{-r/g}$$

and

$$\tilde{C}_{\pi^2}(S) = F_2 + F_1 \sum_{n=0}^{k-1} \left(e^g + \frac{nX_1}{D_0}\right)^{-r/g}$$

From algebra, $\tilde{C}_{\pi^1}(S) > \tilde{C}_{\pi^2}(S)$. ■

7 Conclusions

In this paper we have abandoned the common assumption of a unique optimal initial policy in infinite horizon optimization. Instead, we have addressed the question of whether the set of finite horizon optimal strategies converges to the set of infinite horizon optimal strategies as the horizon is lengthened. We presented two alternate reachability conditions which are sufficient for set convergence. The second condition, for regeneration point problems, is also necessary for set convergence. We have shown that convergence of the optimal sets leads to a simple tie-breaking algorithm, and have given a stopping rule which is guaranteed to be satisfied in finite time. The Shapiro-Wagner algorithm for solving knapsack problems (Shapiro and Wagner [1967]) is a special case of our tie-breaking algorithm.



References

- Bean, J. and R. L. Smith [1984], "Conditions for the Existence of Planning Horizons," **Mathematics of Operations Research**, Vol. 9, pp. 391-401.
- Bean, J. and R. L. Smith [1985], "Optimal Capacity Expansion Over an Infinite Horizon," **Management Science**, Vol. 31, pp. 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 [1986], "Concepts of Forecast and Decision Horizons: Applications to Dynamic Stochastic Optimization Problems," to appear in **Mathematics of Operations Research**.
- Charnes, A., J. Drèze and M. Miller [1966], "Decision and Horizon Rules for Stochastic Planning Problems: A Linear Example," **Econometrica**, Vol. 34, pp. 307-330.
- Hopp, W., J. Bean and R. L. Smith [1984], "Non-Homogeneous Markov Decision Processes," to appear in **Operations Research**.
- Lasserre, J., C. Bès, and F. Roubellat [1984], "Detection of Planning Horizons for the General Inventory Problem with Discounted Concave Costs," Laboratoire d'Automatique et d'Analyse des systèmes du C.N.R.S., 7 Avenue du colonel Roche, 31400, Toulouse, France.
- Lundin and Morton [1975], "Planning Horizons for the Dynamic Lot Size Model: Zabel vs. Protective Procedures and Computations Results," **Operations Research**, Vol. 23, pp. 711-734.
- McKenzie, L. W. [1976], "Turnpike Theory," **Econometrica**, Vol. 44, pp.841-844.
- Modigliani, F. and F. E. Hohn [1955], "Production Planning over Time and the Nature of the Expectation and Planning Horizon," **Econometrica**, Vol. 23, pp. 46-66.
- Murty, K. G. [1983], **Linear Programming**, Wiley, New York.
- 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.
- Schochetman, I. and R. L. Smith [1987a], "Infinite Horizon Optimization," Technical Report 87-3, Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, Michigan 48109.
- Schochetman, I. and R. L. Smith [1987b], "Set Convergence in Infinite Horizon Optimization," Technical Report 87-2, Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, Michigan 48109.
- Shapiro, J. F. and H. M. Wagner [1967], "A Finite Renewal Algorithm for the Knapsack and Turnpike Models," **Operations Research** Vol. 15, pp.319-341.
- Smith, R. L. [1979], "Turnpike Results for Single Location Capacity Expansion," **Management Science**, Vol. 25, pp.474-484.