

Degeneracy in Infinite Horizon Optimization

Sarah M. Ryan
James C. Bean

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

January 27, 1987
Revised August 12, 1987

Technical Report 87-5

Degeneracy in Infinite Horizon Optimization [†]

Sarah M. Ryan
James C. Bean

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

August 12, 1987

Abstract

We consider sequential decision problems over an infinite horizon. The forecast or solution horizon approach to solving such problems requires that the optimal initial decision be unique. We show that multiple optimal initial decisions can exist in general and refer to their existence as degeneracy. We then present a conceptual cost perturbation algorithm for resolving degeneracy and identifying a forecast horizon. We also present a general near-optimal forecast horizon.

Key Words and Phrases

Sequential decision problems, infinite horizon optimization, forecast or solution horizons, degeneracy, perturbation, near-optimal forecast horizon.

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

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

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, Dreze 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 (1986)). Multiple optima may lead to cases with no forecast horizon. We refer to this as *degeneracy*. For the problem of minimizing discounted costs, this paper addresses two questions:

1. How often does degeneracy occur?
2. How can degeneracy be resolved?

Bean and Smith (1984) showed that forecast horizons exist when every optimal sequence of decisions includes the same initial decision and that uniqueness of the optimal initial decision depends on the interest rate used. We approach the first question by characterizing the interest rates that may allow degeneracy and show that, even in simple problems, every “reasonable” interest rate potentially allows degeneracy. We conclude that degeneracy is a serious theoretical problem and turn to the second question.

Recently, Bès and Sethi (1986) found an algorithm for the case of discrete time and discrete decision sets that is guaranteed to identify a forecast horizon when the optimal initial decision is unique. Schochetman and Smith (1986) extended their results to the case of continuous time and continuous decision sets. In this paper, we adapt the Bès and Sethi algorithm to the context of continuous time and discrete decision sets. More important, we present a scheme for perturbing time zero costs which guarantees that the optimal initial decision is unique. We also calculate, for any ϵ , a horizon that yields an initial decision that is part of a strategy with cost within ϵ of the optimal cost.

Section 2 includes a mathematical statement of the problem and assumptions. Section 3 is a discussion

of the relationship between degeneracy and the interest rate. Section 4 gives characteristics of the cost function important to the forecast horizon results in Section 5. Section 6 presents our method for degeneracy resolution. 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 Π_n is finite for all n . Let $\Pi \subseteq \times_{n=1}^{\infty} \Pi_n$ be the set of all feasible strategies. 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:

$$\begin{aligned} C_{\pi}(t) &= \text{the cumulative net cost for strategy } \pi \text{ up to time } t \\ &= K_{\pi}(t) - R_{\pi}(t), \end{aligned} \tag{1}$$

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

$K_{\pi}(t), R_{\pi}(t)$ 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 M by $M e^{\gamma T_0}$).

As in Bean and Smith (1984) we define the metric

$$\rho(\pi, \pi') = \sum_{n=1}^{\infty} \phi_n(\pi, \pi') 2^{-n},$$

where

$$\phi_n(\pi, \pi') = \begin{cases} 1 & \text{if the } n^{\text{th}} \text{ policies in } \pi, \pi' \text{ are different,} \\ 0 & \text{if the } n^{\text{th}} \text{ policies in } \pi, \pi' \text{ are the same.} \end{cases}$$

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:

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

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 $\tilde{C}_\pi(r)$ and *finite horizon optimal* if it minimizes $\tilde{C}_\pi(r, T)$ 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 $\Pi_1^*(T)$.

We formally define a (*weak*) *forecast horizon* as a time, \hat{T} , such that for $T \geq \hat{T}$, $\Pi_1^*(T) = \Pi_1^*$, with $|\Pi_1^*| = 1$. As mentioned earlier, Bean and Smith (1984) showed that $|\Pi_1^*| = 1$ is sufficient to ensure existence of a forecast horizon. We define *degeneracy* in infinite horizon optimization as the case when $|\Pi_1^*| > 1$.

3 Degeneracy and the Interest Rate

Whether or not a problem is degenerate is a function of the interest rate chosen. In this section we investigate the interest rate's effect on the number of optimal initial policies. If an interest rate allows multiple optimal initial policies, we call it a *degenerate rate*.

Since $\tilde{C}_\pi(r)$ is a Laplace-Stieltjes transform, it is analytic (Widder (1946)). From analytic function theory (see Bartle (1964)) we know that for any two strategies π^1 and π^2 , $\tilde{C}_{\pi^1}(r) = \tilde{C}_{\pi^2}(r)$ for at least countably infinite r in any closed interval in (γ, ∞) if and only if $\tilde{C}_{\pi^1}(r) = \tilde{C}_{\pi^2}(r)$ for all $r > \gamma$. It follows that if

$\tilde{C}_{\pi_1}(\hat{r}) \neq \tilde{C}_{\pi_2}(\hat{r})$ for some $\hat{r} > \gamma$ then $\tilde{C}_{\pi_1}(r) = \tilde{C}_{\pi_2}(r)$ for at most finitely many r in every closed interval in (γ, ∞) . Using this fact, Bean and Smith (1984) proved the following:

Theorem 1 *If the set of potentially optimal strategies is at most countable, then degeneracy can occur for at most a countable number of interest rates in the range of convergence.*

When Theorem 1 holds, the set of degenerate rates has measure zero. Hence the probability of selecting a degenerate rate is zero. If degenerate rates are also isolated in the range of convergence, then if a degenerate rate is encountered we can apply an arbitrarily small perturbation to the interest rate to find a nondegenerate rate, which will yield a forecast horizon. A special class of problems in which degenerate rates are guaranteed to be isolated is described in Ryan and Bean (1986).

Bean and Smith (1984) include a discussion of cases in which the set of potentially optimal strategies is countable. However, in general this set is uncountable. If two policies are available at each policy epoch, each strategy corresponds to an infinite sequence of ones and zeros. The set of all such sequences is uncountable (Rudin (1976)).

To examine the case of uncountable strategies, consider a case where decisions are made periodically and each policy incurs a discrete cost at the beginning of the period in which it is implemented. Thus a strategy π incurs the sequence of costs $\{c_1^\pi, c_2^\pi, \dots\}$. Letting $\alpha = e^{-r}$, we can write the discounted cost as $\tilde{C}_\pi(\alpha) = \sum_{n=1}^{\infty} c_n^\pi \alpha^n$. We further assume that $c_n^\pi \in \{0, 1, 2, \dots, L\}$ for each $\pi \in \Pi$, where L is a large integer. In this case $\gamma = 0$.

Let $C = \{c = \{c_n\}_{n=1}^{\infty} \mid 0 \leq c_n \leq L, c_n \text{ integer } \forall n\}$. Given α , let $f(c) = \sum_{n=1}^{\infty} c_n \alpha^n$. We can show the following:

Lemma 2 *If $\frac{1}{L+1} \leq \alpha < 1$, then the image of C under the function f is $[0, \frac{L\alpha}{1-\alpha}]$.*

Proof:

Claim (1): If $c \in C$, then $f(c) \in [0, \frac{L\alpha}{1-\alpha}]$

Proof:

$$\sum_{k=n}^m c_k \alpha^k \leq L \sum_{k=n}^m \alpha^k \leq L \sum_{k=n}^{\infty} \alpha^k = \frac{L\alpha^n}{1-\alpha}$$

For every $\epsilon > 0$, $\exists N$ such that $\frac{L\alpha^n}{1-\alpha} < \epsilon$, $\forall n \geq N$. Therefore the series converges by the Cauchy criterion.

For any c ,

$$0 \leq \sum_{n=1}^{\infty} c_n \alpha^n \leq \frac{L\alpha}{1-\alpha}.$$

Claim (2): $f : C \mapsto [0, \frac{L\alpha}{1-\alpha}]$ is onto.

Proof: Given $x \in [0, \frac{L\alpha}{1-\alpha}]$, if $x = \frac{L\alpha}{1-\alpha}$ then set $c_n = L \forall n$. $f(c) = x$.

Suppose $x < \frac{L\alpha}{1-\alpha}$.

Let $c_1 =$ the largest integer $\leq L$ such that $c_1 \alpha < x$.

Given c_1 , let $c_2 =$ the largest integer $\leq L$ such that $c_1 \alpha + c_2 \alpha^2 < x$.

In general, given c_1, c_2, \dots, c_{n-1} , let $c_n =$ the largest integer $\leq L$ such that $\sum_{k=1}^n c_k \alpha^k < x$. We

claim $\sum_{n=1}^{\infty} c_n \alpha^n = x$.

Suppose $\sum_{n=1}^{\infty} c_n \alpha^n < x$. Note that $x < \frac{L\alpha}{1-\alpha} \Rightarrow c_n < L$ for some n . Suppose that for some k , $c_k < L$ and $c_n = L, \forall n > k$. Then

$$\sum_{n=k+1}^{\infty} c_n \alpha^n = \frac{L\alpha^{k+1}}{1-\alpha} \geq \alpha^k \text{ since } \alpha \geq \frac{1}{L+1}.$$

Then $\sum_{n=1}^{k-1} c_n \alpha^n + (c_k + 1)\alpha^k > x$ which contradicts the construction of c_k . Hence for any k , $\exists n > k$ with $c_k < L$.

Now,

$$\sum_{n=1}^{\infty} c_n \alpha^n < x \Rightarrow \sum_{n=1}^{\infty} c_n \alpha^n + b = x, \quad b > 0.$$

There exists n such that $\alpha^n < b$. Let n_0 be the smallest such n . Let n_1 be the smallest $n > n_0$ such that $c_n < L$. Then

$$\sum_{n=1}^{\infty} c_n \alpha^n + \alpha^{n_1} < x,$$

and c_{n_1} could be incremented without $\sum_{n=1}^{n_1} c_n \alpha^n$ exceeding x . This contradicts the construction of c_{n_1} . Therefore $\sum_{n=1}^{\infty} c_n \alpha^n \geq x$.

Now suppose $\sum_{n=1}^{\infty} c_n \alpha^n > x$. Then $\sum_{n=1}^{\infty} c_n \alpha^n - d = x, \quad d > 0$. Then $\sum_{n=1}^N c_n \alpha^n + \sum_{n=N+1}^{\infty} c_n \alpha^n - d = x$. But $d - \sum_{n=N+1}^{\infty} c_n \alpha^n \geq d - \frac{L\alpha^{N+1}}{1-\alpha} > 0$ for N sufficiently large. Hence $\sum_{n=1}^N c_n \alpha^n > x$ for N sufficiently large which contradicts the construction of c_N . Hence $\sum_{n=1}^{\infty} c_n \alpha^n \leq x$.

Therefore $\sum_{n=1}^{\infty} c_n \alpha^n = x$. ■

Theorem 3 If $\frac{1}{L+1} < \alpha < 1$ then there exist infinitely many pairs $c^1, c^2 \in C$ such that $c_1^1 \neq c_1^2$ and $f(c^1) = f(c^2)$.

Proof: Let $c^K = \{c_1, \overbrace{0, 0, \dots, 0}^K, c_{K+2}, \dots\}$, $c_1 > 0$. Then $f(c^K) \in [c_1\alpha, c_1\alpha + \frac{L\alpha^{K+2}}{1-\alpha}]$. Let $B = \{b \in C | b_1 = c_1 - 1\}$. For $b \in B$, $f(b) = (c_1 - 1)\alpha + \sum_{n=2}^{\infty} b_n \alpha^n$. By a variation of Lemma 2,

$$\{f(b) | b \in B\} = \left[(c_1 - 1)\alpha, (c_1 - 1)\alpha + \frac{L\alpha^2}{1 - \alpha} \right]$$

Now,

$$\alpha > \frac{1}{L+1} \Rightarrow \alpha < \frac{L\alpha^2}{1 - \alpha}$$

and

$$\alpha + \frac{L\alpha^{K+2}}{1 - \alpha} < \frac{L\alpha^2}{1 - \alpha} \text{ for } K > \frac{\log \frac{(L+1)\alpha - 1}{L}}{\log \alpha} - 1.$$

Hence for K sufficiently large we have

$$\left[c_1\alpha, c_1\alpha + \frac{L\alpha^{K+2}}{1 - \alpha} \right] \subset \left[(c_1 - 1)\alpha, (c_1 - 1)\alpha + \frac{L\alpha^2}{1 - \alpha} \right].$$

Hence $f(c^K) \in \{f(b) | b \in B\}$. Therefore, $\exists b \in B$ such that $f(b) = f(c^K)$. This holds for each value of K sufficiently large. Therefore there exist at least countably infinitely many such pairs for *each* value of $\alpha > \frac{1}{L+1}$. ■

Theorem 3 implies that for *any* interest rate r , $0 < r < L$, infinitely many pairs of strategies can attain the same cost, with disagreement in the initial policy. It remains unknown whether such pairs of strategies can tie for optimal for many interest rates. But at this point, we must assume that degeneracy is a serious theoretical problem in infinite horizon optimization. In the remainder of the paper we develop ways of resolving degeneracy when it occurs.

4 Cost Function Characteristics

This section contains results that will be used to derive forecast horizon results in Section 5. They were derived by Bès and Sethi (1986) for the case of discrete time and discrete policy sets and extended to the case of continuous time and continuous policy sets by Schochetman and Smith (1986). We refer the reader to Ryan and Bean (1986) for details on their adaptation to the case of continuous time and discrete policy sets.

Using the exponential bound on cumulative costs, we can derive bounds on discounted costs.

Definition: $a(T) = \frac{rM}{r-\gamma}e^{-(r-\gamma)T}$ is an upper bound on $\int_T^\infty e^{-rt}dC_\pi(t)$ for any $\pi \in \Pi$, where M is defined in (1).

Note that $a(T) \rightarrow 0$ as $T \rightarrow \infty$ for $r > \gamma$.

Lemma 4 For $\pi \in \Pi$, $r > \gamma$, and $T \geq 0$,

- (a) $|\tilde{C}_\pi(r) - \tilde{C}_\pi(r, T)| \leq a(T)$. In particular, it follows that as $T \rightarrow \infty$, $|\tilde{C}_\pi(r) - \tilde{C}_\pi(r, T)| \rightarrow 0$ uniformly with respect to $\pi \in \Pi$.
- (b) For $S \geq T$, $|\tilde{C}_\pi(r, S) - \tilde{C}_\pi(r, T)| \leq a(T)$.
- (c) $|\tilde{C}^*(r) - \tilde{C}^*(r, T)| \leq a(T)$.

Lemma 5 Let $\{T_n\}$ be a sequence of positive times such that $T_n \rightarrow \infty$ as $n \rightarrow \infty$. For each $n = 1, 2, 3, \dots$, let $\pi^n \in \Pi^*(T_n)$. Then \exists a subsequence $\{\pi^k\}$ of $\{\pi^n\}$ and an element $\pi^* \in \Pi^*$ with the property that $\pi^k \rightarrow \pi^*$ as $k \rightarrow \infty$. This property holds for any convergent subsequence of $\{\pi^n\}$.

5 Forecast Horizon Results

The existence of a forecast horizon and the algorithm for identifying it are based on the bounds on the discounted costs derived in Section 4. Lemma 6 allows us to restrict the set of potentially optimal strategies based on their minimum cost up to a finite time horizon. Theorem 7 expresses a stopping criterion for identifying a forecast horizon. Theorems 8 and 9 state that the stopping criterion will be satisfied if the optimal initial policy is unique. For the remainder of the paper we assume $r > \gamma$ is fixed.

Denote the set of initial policies as $\Pi_1 = \{\pi_1^1, \pi_1^2, \dots, \pi_1^k\}$ for some k . Let ${}_i\tilde{C}^*(T) = \min\{\tilde{C}_\pi(T) | \pi \in \Pi, \pi_1 = \pi_1^i\}$ and ${}_i\tilde{C}^* = \min\{\tilde{C}_\pi | \pi \in \Pi, \pi_1 = \pi_1^i\}$. These minima exist since $\{\pi \in \Pi | \pi_1 = \pi_1^i\}$ is compact. Also, let $\hat{C}(T) = \min\{{}_i\tilde{C}^*(T) | i : \pi_1^i \notin \Pi_1^*(T)\}$ be the minimum T -horizon cost given a suboptimal initial policy and $\hat{C} = \min\{{}_i\tilde{C}^* | i : \pi_1^i \notin \Pi_1^*\}$ be its infinite horizon counterpart.

Lemma 6 If for any strategy $\hat{\pi}$, $\tilde{C}_{\hat{\pi}}(T) - \tilde{C}^*(T) > 2a(T)$ then $\hat{\pi} \notin \Pi^*(S)$ for all $S \geq T$.

Proof: For each $S \geq T$, and $\forall \pi \in \Pi$,

$$\tilde{C}_\pi(S) \leq \tilde{C}_\pi(T) + a(T), \text{ by Lemma 4(b).}$$

Then by taking infimums on the right and then the left,

$$\tilde{C}^*(S) \leq \tilde{C}^*(T) + a(T).$$

We also have

$$\tilde{C}_{\hat{\pi}}(S) \geq \tilde{C}_{\hat{\pi}}(T) - a(T) \text{ by Lemma 4(c).}$$

Then, by subtracting,

$$\tilde{C}_{\hat{\pi}}(S) - \tilde{C}^*(S) \geq \tilde{C}_{\hat{\pi}}(T) - \tilde{C}^*(T) - 2a(T) > 0. \blacksquare$$

We generalize two theorems of Bès and Sethi (1986) to continuous time.

Theorem 7 *If $\hat{C}(T) - \tilde{C}^*(T) > 2a(T)$ and $\Pi_1^*(T) = \{\pi_1^*(T)\}$ then $\Pi_1^*(S) = \{\pi_1^*(T)\}$ for all $S \geq T$.*

Proof: Choose $S > T$, and consider any $\pi \in \Pi$ with $\pi_1 \neq \pi_1^*(T)$. By definition, $\tilde{C}_{\pi}(T) \geq \hat{C}(T)$. Therefore,

$$\tilde{C}_{\pi}(T) - \tilde{C}^*(T) > 2a(T).$$

Then by Lemma 6, $\pi \notin \Pi^*(S)$. This is true for each $\pi \in \Pi$ with $\pi_1 \neq \pi_1^*(T)$. Hence $\Pi_1^*(S) \subseteq \Pi_1^*(T)$. \blacksquare

Theorem 8 *If $\Pi_1^* = \{\pi_1^*\}$ then $\exists T$ such that $\forall S \geq T$, $\Pi_1(S) = \{\pi_1^*\}$.*

Proof: Suppose not. Then \exists a sequence $\{T_n\}$ such that $T_n \rightarrow \infty$ and one can find $\pi^n \in \Pi^*(T_n)$ with $\pi_1^n \neq \pi_1^*$. By Lemma 5, \exists a subsequence $\{\pi^k\}$ of $\{\pi^n\}$ such that $\pi^k \rightarrow \hat{\pi} \in \Pi^*$. Then, for k sufficiently large, $\pi_1^k = \hat{\pi}_1$. But $\hat{\pi}_1 = \pi_1^*$, which is a contradiction. \blacksquare

Now let $b(t)$ be any function satisfying $b(t) > 0$, for $t > 0$, and $b(t) \rightarrow 0$ as $t \rightarrow \infty$. Bès and Sethi (1986) prove the following result for discrete time and discrete policy sets. Schochetman and Smith (1986) extend it to continuous time and continuous policy spaces. The version here applies to continuous time and discrete policy sets.

Theorem 9 *If $|\Pi_1^*| = 1$, then there exists $T > 0$ such that $\hat{C}(T) - \tilde{C}^*(T) > b(T)$.*

Proof: See Ryan and Bean (1986).

We can sum up our results so far:

Theorem 10 *If Π_1^* is a singleton then $\exists T > 0$ such that T is a forecast horizon.*

Proof: Follows from Theorems 7, 8, 9, with $b(t) = a(t)$. ■

Theorem 10 was also proved by Bean and Smith (1984). But from Theorems 7, 8, and 9 we also get an algorithm, generalized from Bés and Sethi (1986), guaranteed to find a forecast horizon whenever $|\Pi_1^*| = 1$:

Algorithm

- (1) *Choose any sequence $\{T_n\}_{n=1}^\infty$ such that $T_n \rightarrow \infty$ as $n \rightarrow \infty$. Set $n \leftarrow 1$.*
- (2) *Solve the T_n -horizon problem to get $\Pi_1^*(T_n)$, $\tilde{C}^*(T_n)$, and $\hat{C}(T_n)$.*
- (3) *If $|\Pi_1^*(T_n)| = 1$ and $\hat{C}(T_n) - \tilde{C}^*(T_n) > 2a(T_n)$ then STOP: T_n is a forecast horizon. Else set $n \leftarrow n + 1$ and go to (1).*

Theorem 11 *If Π_1^* is a singleton then the Algorithm will terminate with an infinite-horizon optimal initial policy for finite T_n .*

Proof: Follows from Theorems 7, 8, 9. ■

6 Degeneracy Resolution

The results in Section 5 show that we can find a forecast horizon if the optimal initial policy is unique. Degeneracy can occur if there is more than one optimal initial policy. In this section we present two approaches to resolving degeneracy that follow from the results of the previous section. One is an exact algorithm for finding an optimal initial policy by perturbing costs. Together with the algorithm in Section 5, this constitutes the first method guaranteed to terminate in finite time when solving any of a general class of infinite horizon problems. The second approach is the calculation of a forecast horizon with the guarantee that any finite horizon optimal first policy has infinite horizon discounted cost nearly as small as that of the infinite horizon optimal first policy.

We make some additional definitions:

Let $G = \{l : \pi_1^l \in \Pi_1^*\}$, the set of indices for optimal initial policies. Let $\epsilon = (\epsilon_1, \epsilon_2, \dots, \epsilon_k)$ be a perturbation vector and set $C_\pi^\epsilon(t) = C_\pi(t) + \epsilon_i$, $t > 0$, if $\pi_1 = \pi_1^i$; that is, add an instantaneous, discrete cost

of ϵ_i for π_1^i at time 0. Define $\tilde{C}_\pi^\epsilon(r)$, $\tilde{C}^{\epsilon*}$, ${}_i\tilde{C}^{\epsilon*}$, \hat{C}^ϵ as for $C_\pi(t)$. Let $\Pi^{\epsilon*} = \{\pi : \pi = \arg \min_{\pi \in \Pi} \tilde{C}_\pi^\epsilon(r)\}$ be the set of optimal strategies for the perturbed problem and let $\Pi_1^{\epsilon*}$ be the set of initial policies corresponding to $\pi \in \Pi^{\epsilon*}$. Set $G^\epsilon = \{l : \pi_1^l \in \Pi_1^{\epsilon*}\}$, the set of indices for optimal policies for the perturbed problem. Finally, let $\hat{\epsilon} = \hat{C} - \tilde{C}^*$ be the minimum penalty for choosing a suboptimal initial policy in the original problem.

Adjust M if necessary so that

$$K_\pi(t) + \hat{\epsilon} \leq Me^{\gamma t} \text{ for } t \geq 0.$$

Lemma 12 *Let ϵ be such that $\epsilon_i < \hat{\epsilon} \forall i$, and $\epsilon_i \neq \epsilon_j$ for $i \neq j$. Then $G^\epsilon \subseteq G$ and $|G^\epsilon| = 1$.*

Proof: Suppose $j \notin G$. Then for any $i \in G$, ${}_j\tilde{C}^* - {}_i\tilde{C}^* \geq \hat{\epsilon}$. Then

$${}_j\tilde{C}^{\epsilon*} - {}_i\tilde{C}^{\epsilon*} = {}_j\tilde{C}^* - {}_i\tilde{C}^* + \epsilon_j - \epsilon_i \geq \hat{\epsilon} + \epsilon_j - \epsilon_i.$$

But $0 < \epsilon_i < \hat{\epsilon} \forall i \Rightarrow |\epsilon_j - \epsilon_i| < \hat{\epsilon}$. Hence, ${}_j\tilde{C}^{\epsilon*} - {}_i\tilde{C}^{\epsilon*} > 0$. Therefore, $j \notin G^\epsilon$. Thus $G^\epsilon \subseteq G$.

Now suppose $i \in G^\epsilon$. Suppose $j \in G^\epsilon$ also. Then $i \in G$ and $j \in G$. Hence ${}_i\tilde{C}^* = {}_j\tilde{C}^*$. But

$\epsilon_i \neq \epsilon_j \Rightarrow {}_i\tilde{C}^{\epsilon*} \neq {}_j\tilde{C}^{\epsilon*}$. This is a contradiction. ■

Theorem 13 *Let costs be perturbed as in Lemma 12. Then the Algorithm will terminate in finite time.*

Proof: After the perturbation, $\Pi_1^{\epsilon*}$ is a singleton. The result follows from Theorem 11. ■

Because calculating $\hat{\epsilon}$ requires all the data over the infinite horizon, this perturbation scheme cannot be implemented. We are currently working on an algorithm to implement these ideas without calculating $\hat{\epsilon}$ directly. But as a consequence of Lemma 6, we can find a near-optimal forecast horizon.

We say that T is an ϵ -forecast horizon if $\tilde{C}_{\pi^*(S)} - \tilde{C}^* \leq \epsilon$ for all $S \geq T$ and all $\pi^*(S) \in \Pi^*(S)$.

Theorem 14 *If $T_\epsilon > \frac{1}{r-\gamma} \log \frac{4rM}{(r-\gamma)\epsilon}$ then T_ϵ is an ϵ -forecast horizon.*

Proof: For any π , $\tilde{C}_\pi \leq \tilde{C}_\pi(T) + a(T)$ by Lemma 4(a). Also, $\tilde{C}^* \geq \tilde{C}^*(T) - a(T)$ by Lemma 4(c). For any

$S \geq T$, let $\pi^S \in \Pi^*(S)$. Then $\tilde{C}_{\pi^S}(T) - \tilde{C}^*(T) \leq 2a(T)$ by Lemma 6. Then

$$\tilde{C}_{\pi^S} - \tilde{C}^* \leq \tilde{C}_{\pi^S}(T) + a(T) - \tilde{C}^*(T) + a(T) \leq 4a(T)$$

$$< \epsilon \text{ if } T > \frac{1}{r-\gamma} \log \frac{4rM}{(r-\gamma)\epsilon}. \blacksquare$$

Note that by Lemma 4(c), $\tilde{C}^* \leq a(0)$. Let $p = \frac{\epsilon}{a(0)} = \frac{\epsilon(r-\gamma)}{rM}$ be the error expressed as a proportion of the maximum possible total discounted cost. Then $T_\epsilon = \frac{1}{r-\gamma} \log \frac{4}{p}$. Thus the ϵ -forecast horizon length grows proportionally to the log of the accuracy desired and inversely proportionally to the interest rate used. These properties are consistent with those of the ϵ -forecast horizons found by Bean and Smith (1985) for capacity expansion problems and by Bean, Birge, and Smith (1984) for deterministic infinite dynamic programs.

7 Conclusions

Forecast horizon existence results have been found for problems with special structure (see Gilmore and Gomory (1966), Morton (1978), and Shapiro and Wagner (1967)). But all general existence results (Bean and Smith (1984), Hopp, Bean, and Smith (1984), Bès and Sethi (1986), Schochetman and Smith (1986)) assume that the optimal initial policy is unique. The results in this paper indicate that this assumption may not be valid. An example is constructed in which every reasonable interest rate may lead to degeneracy. We have therefore turned our attention to resolving degeneracy when it occurs and present the first algorithm guaranteed to stop regardless of degeneracy. Though our method for perturbing time zero costs cannot currently be implemented, it shows that degeneracy resolution is theoretically possible.

In some situations a near-optimal initial policy will suffice. Then an ϵ -forecast horizon can be calculated using only the interest rate and parameters of an exponential function that bounds the total cost for any strategy. Stronger near-optimal horizons, as in Bean and Smith (1985), are likely to exist for specific applications.

Acknowledgment

We would like to thank Professors Robert L. Smith and Hugh L. Montgomery of The University of Michigan for their suggestions and comments.

References

- R. G. Bartle, *The Elements of Real Analysis* (Wiley, New York, 1964).
- J. Bean, J. Birge and R. L. Smith, "Aggregation in Dynamic Programming," Technical Report 84-10, Department of Industrial and Operations Engineering, The University of Michigan (Ann Arbor, MI, 1984).
- J. Bean and R. L. Smith, "Conditions for the Existence of Planning Horizons," *Mathematics of Operations Research* 9(1984)391-401.
- J. Bean and R. L. Smith, "Optimal Capacity Expansion Over an Infinite Horizon," *Management Science* 31(1985)1523-1532.
- C. Bès and S. Sethi, "Concepts of Forecast and Decision Horizons: Applications to Dynamic Stochastic Optimization Problems," unpublished manuscript, University of Toronto (Toronto, 1986).
- A. Charnes, J. Dreze and M. Miller, "Decision and Horizon Rules for Stochastic Planning Problems: A Linear Example," *Econometrica* 34(1966)307-330.
- P. C. Gilmore and R. E. Gomory, "The Theory and Computation of Knapsack Functions," *Operations Research* 14(1966)1045-1074.
- W. Hopp, J. Bean and R. L. Smith, "A New Optimality Criterion for Non-Homogeneous Markov Decision Problems," Technical Report 84-24, Department of Industrial and Operations Engineering, The University of Michigan (Ann Arbor, MI, 1984). To appear in *Operations Research*.
- R. A. Lundin and T. E. Morton, "Planning Horizons for the Dynamic Lot Size Model: Zabel vs. Protective Procedures and Computations Results," *Operations Research* 23(1975)711-734.
- F. Modigliani and F. Hohn, "Production Planning Over Time and the Nature of the Expectation and Planning Horizon," *Econometrica* 23(1955)46-66.
- T. E. Morton, "Universal Planning Horizons for Generalized Convex Production Scheduling," *Operations Research* 26(1978)1046-1058.
- W. Rudin, *Principles of Mathematical Analysis* (McGraw-Hill, New York, 1976).
- S. M. Ryan and J. Bean, "Cost-Based Means for Resolving Degeneracy in Infinite Horizon Optimization," Technical Report 86-41, Department of Industrial and Operations Engineering, The University of Michigan (Ann Arbor, MI, 1986).
- I. Schochetman and R. L. Smith, "Infinite Horizon Optimization," Technical Report, Department of Industrial and Operations Engineering, The University of Michigan (Ann Arbor, MI, 1986).
- J. F. Shapiro and H. M. Wagner, "A Finite Renewal Algorithm for the Knapsack and Turnpike Models," *Operations Research* 15(1967)319-341.
- D. V. Widder, *The Laplace Transform* (Oxford University Press, Oxford, 1946).