

**COST-BASED MEANS FOR RESOLVING DEGENERACY IN
INFINITE HORIZON OPTIMIZATION***

**Sarah M. Ryan
James C. Bean**

**Technical Report 86-41
December 2, 1986**

***This material is based on work supported by the National Science Foundation
under Grant No. ECS-8409682.**

Cost-Based Means for Resolving Degeneracy in Infinite Horizon Optimization [†]

Sarah M. Ryan
James C. Bean

December 2, 1986

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, the situation 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

[†]This material is based on work supported by the National Science Foundation under Grant No. ECS-8409682.

Cost-Based Means for Resolving Degeneracy in Infinite Horizon Optimization [†]

Sarah S. McAllister
James C. Bean

December 2, 1986

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, the situation 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

[†]This material is based on work supported by the National Science Foundation under Grant No. ECS-8409682.

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 planning or forecast horizon approach (Lundin and Morton [1975]; Charnes, Dreze and Miller [1966]; Modigliani and Hohn [1955]).

We define *degeneracy* as the absence of any forecast horizon. Bean and Smith [1984] showed that forecast horizons exist when every optimal sequence of decisions includes the same initial decision. 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.

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 , may also be called a *policy*, and a sequence $\pi = (\pi_1, \pi_1, \dots)$ of policies constitutes a *strategy*. Π_n is the set of decisions available after $n - 1$ decisions have been made. $\Pi = \times_{n=1}^{\infty} \Pi_n$ is 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. $\tilde{C}_\pi(r) = \int_0^\infty e^{-rt} dC_\pi(t)$ is the resulting infinite horizon discounted cost as a function of the discount rate, r . $\tilde{C}_\pi(r, T) = \int_0^T e^{-rt} dC_\pi(t)$ 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 discount rates.

Given a discount rate, r , the problem we wish to solve is:

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

As shown by Bean and Smith [1984], 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 . Π^* and $\Pi^*(T)$ are the sets of optimal strategies for the infinite horizon and finite horizon problems, respectively. An *optimal initial decision* is an initial decision included in some optimal strategy. The sets of optimal initial decisions for the infinite and finite horizon problems, respectively, are denoted Π_1^* and $\Pi_1^*(T)$.

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

$C_\pi(t)$ = the cumulative net cost for strategy π up to time t

$$= K_\pi(t) - R_\pi(t),$$

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

$K_\pi(t), R_\pi(t)$ nondecreasing.

Without loss of generality, we can assume $T_0 = 0$ (if not, replace M by $M e^{\gamma T_0}$).

We also assume that Π_n is finite for all n . Bès and Sethi ([1986]) proved that this assumption is necessary for Π to be compact. Compactness allows the existence of an optimal strategy (Bean and Smith [1984]).

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 decisions are closer than any two strategies which do not.

A *forecast horizon* is a point in time such that if the time horizon is truncated at or beyond this time, an optimal initial policy for the truncated problem will be optimal for the infinite horizon problem as well. Forecast horizon approaches consist of solving successively longer finite horizon problems until a forecast horizon is discovered. Bean and Smith [1984] showed that uniqueness of the optimal initial decision is sufficient for a forecast horizon to exist. Schochetman and Smith [1986] found a stronger result in the case of continuous decision sets. They define an *algorithmically optimal (AO)* strategy to

be an optimal strategy that is the limit of optimal strategies for finite horizon problems. This concept was originally introduced by Hopp, Bean, and Smith [1984] for discrete time and finite decision sets. There such strategies were called *periodic forecast horizon optimal* strategies. Schochetman and Smith demonstrated that the set of AO strategies can be considerably smaller than the set of all optimal strategies. An AO initial decision is an initial decision included in an AO strategy. Schochetman and Smith showed that uniqueness of the AO initial decision is *necessary* and sufficient for the existence of forecast horizons.

We define *degeneracy* in infinite horizon optimization as the presence of multiple AO initial decisions, i.e., the absence of any forecast horizon.

3 Cost Function Characteristics

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

$$a(t) = \frac{rM}{r - \gamma} e^{-(r-\gamma)t}$$

Note that $a(t) \rightarrow 0$ as $t \rightarrow \infty$.

The following theorem is proved in Schochetman and Smith [1986]:

Theorem 1 For $\pi \in \Pi$, $r > \gamma$, and $T \geq 0$,

- (i) $0 \leq \tilde{K}_\pi(r) - \tilde{K}_\pi(r, T) \leq a(T)$
- (ii) $0 \leq \tilde{R}_\pi(r) - \tilde{R}_\pi(r, T) \leq a(T)$
- (iii) $|\tilde{C}_\pi(r) - \tilde{C}_\pi(r, T)| \leq a(T)$

In particular, it follows that as $T \rightarrow \infty$, $\left| \tilde{C}_\pi(r) - \tilde{C}_\pi(r, T) \right| \rightarrow 0$ uniformly with respect to $\pi \in \Pi$.

Corollary 2 For $S \geq T$, $\left| \tilde{C}_\pi(r, S) - \tilde{C}_\pi(r, T) \right| \leq a(T)$.

Proof: From Theorem 1,

$$0 \leq \tilde{K}_\pi(r, S) - \tilde{K}_\pi(r, T) \leq \tilde{K}_\pi(r) - \tilde{K}_\pi(r, T) \leq a(T),$$

$$0 \leq \tilde{R}_\pi(r, S) - \tilde{R}_\pi(r, T) \leq \tilde{R}_\pi(r) - \tilde{R}_\pi(r, T) \leq a(T).$$

Now,

$$\tilde{C}_\pi(r, S) - \tilde{C}_\pi(r, T) = \tilde{K}_\pi(r, S) - \tilde{K}_\pi(r, T) - (\tilde{R}_\pi(r, S) - \tilde{R}_\pi(r, T)).$$

Thus

$$-a(T) \leq -(\tilde{R}_\pi(r, S) - \tilde{R}_\pi(r, T)) \leq \tilde{C}_\pi(r, S) - \tilde{C}_\pi(r, T) \leq \tilde{K}_\pi(r, S) - \tilde{K}_\pi(r, T) \leq a(T). \blacksquare$$

Theorem 3 $|\tilde{C}^*(r) - \tilde{C}^*(r, T)| \leq a(T)$.

Proof: From Theorem 1,

$$\tilde{C}_\pi(r) \leq \tilde{C}_\pi(r, T) + a(T), \quad \forall \pi \in \Pi.$$

Taking the infimum on the left, we get

$$\tilde{C}^*(r) \leq \tilde{C}_\pi(r, T) + a(T), \quad \forall \pi \in \Pi.$$

Taking the infimum on the right yields

$$\tilde{C}^*(r) - \tilde{C}^*(r, T) \leq a(T)$$

Similarly, from Theorem 1,

$$\tilde{C}_\pi(r) \geq \tilde{C}_\pi(r, T) - a(T), \quad \forall \pi \in \Pi.$$

Taking the infimum on the right and then the left yields

$$\tilde{C}^*(r) - \tilde{C}^*(r, T) \geq -a(T). \blacksquare$$

4 Continuity Results

In order to demonstrate the existence of forecast horizons and develop an algorithm for finding them, we will require discounted costs to be continuous. Schochetman and Smith [1986] prove the following lemma:

Lemma 4 *Fix $T \geq 0$. Suppose the real sequence $\{r_k\}$ converges to r monotonically from the left or right. Then the sequence of functions $\{e^{-r_k t}\}$ of t converges uniformly to the function e^{-rt} on $[0, T]$.*

Schochetman and Smith [1986] prove the following theorem for continuous strategy spaces subject to an additional assumption on the cost functions. We extend it to the discrete case.

Theorem 5 *The functions $\tilde{C}_\pi(r, T)$, $\tilde{K}_\pi(r, T)$, $\tilde{R}_\pi(r, T)$ of (π, r) are continuous.*

Proof: Suppose $\pi^n \rightarrow \pi$ in Π and $r_n \rightarrow r$ in $(0, \gamma)$. Suppose the theorem is false. Then $\exists \epsilon > 0$ and a subsequence $\{(\pi^m, r_m)\}$ of $\{(\pi^n, r_n)\}$ such that

$$\left| \tilde{K}_{\pi^m}(r_m, T) - \tilde{K}_\pi(r, T) \right| \geq \epsilon, \quad m = 1, 2, \dots$$

Since $r_m \rightarrow r$, \exists a subsequence that converges from either the left or the right. There exists a subsequence of this that converges *monotonically* to r from either the left or the right. Call this subsequence $\{r_k\}$. Then we have $\{(\pi^k, r_k)\} \rightarrow \{(\pi, r)\}$ and

$$\left| \tilde{K}_{\pi^k}(r_k, T) - \tilde{K}_{\pi}(r, T) \right| \geq \epsilon, \quad k = 1, 2, \dots$$

Since $r^k \rightarrow r$ monotonically, we have $e^{-r_k t} \rightarrow e^{-r t}$ uniformly on $[0, T]$ by Lemma 4.

Hence

$$\max_{0 \leq t \leq T} \left| e^{-r_k t} - e^{-r t} \right| \rightarrow 0 \text{ as } k \rightarrow \infty.$$

Then

$$\begin{aligned} \left| \tilde{K}_{\pi^k}(r_k, T) - \tilde{K}_{\pi}(r, T) \right| &= \left| \int_0^T e^{-r_k t} dK_{\pi^k}(t) - \int_0^T e^{-r t} dK_{\pi}(t) \right| \\ &\leq \left| \int_0^T (e^{-r_k t} - e^{-r t}) dK_{\pi^k}(t) \right| + \left| \int_0^T e^{-r t} dK_{\pi^k}(t) - \int_0^T e^{-r t} dK_{\pi}(t) \right| \end{aligned}$$

Let $\text{Var}(f(\cdot), [a, b])$ denote the total variation of the function f on the interval $[a, b]$.

Then

$$\left| \int_0^T (e^{-r_k t} - e^{-r t}) dK_{\pi^k}(t) \right| \leq \max_{0 \leq t \leq T} (e^{-r_k t} - e^{-r t}) \text{Var}(K_{\pi^k}(\cdot), [0, T]) \rightarrow 0$$

since

$$\max_{0 \leq t \leq T} (e^{-r_k t} - e^{-r t}) \rightarrow 0$$

and

$$\text{Var}(K_{\pi^k}(\cdot), [0, T]) \leq M e^{\gamma t}, \quad \forall k.$$

For the second term, note that since $\pi^k \rightarrow \pi$, π^k and π agree over an ever increasing time frame. Let T_k be the time covered by the decisions on which they agree. $T_k \rightarrow \infty$

as $k \rightarrow \infty$. Then $T_k \geq T$ for $k \geq K_T$. Hence

$$\left| \int_0^T e^{-rt} dK_{\pi^k}(t) - \int_0^T e^{-rt} dK_{\pi}(t) \right| = 0, \quad k \geq K_T.$$

Thus

$$\left| \tilde{K}_{\pi^k}(r_k, T) - \tilde{K}_{\pi}(r, T) \right| \rightarrow 0 \text{ as } k \rightarrow \infty.$$

This is a contradiction. The claim for $\tilde{R}_{\pi}(t)$ is proved similarly and thus follows for $\tilde{C}_{\pi}(t)$. ■

The next two theorems are proved as in Schochetman and Smith [1986] using Theorems 1 and 5.

Theorem 6 *The functions $\tilde{C}_{\pi}(r)$, $\tilde{K}_{\pi}(r)$, $\tilde{R}_{\pi}(r)$ of (π, r) are continuous.*

Theorem 7 *Suppose $\pi^n \rightarrow \pi$, $r_n \rightarrow r$, and $T_n \rightarrow \infty$ as $n \rightarrow \infty$. Then $\tilde{C}_{\pi^n}(r_n, T_n) \rightarrow \tilde{C}_{\pi}(r)$ as $n \rightarrow \infty$. The same is true for K and R .*

The following lemma is proved in Bean and Smith [1984], but is not stated explicitly there. We prove it here for the sake of completeness.

Lemma 8 *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 $\{\pi^n\}$.*

Proof: Since Π is compact, \exists a convergent subsequence $\{\pi^k\}$ of $\{\pi^n\}$. Let $\pi' = \lim_{k \rightarrow \infty} \pi^k$. By definition, $\tilde{C}_{\pi^k}(T_k) \leq \tilde{C}_{\pi'}(T_k)$. Taking limits of both sides,

$$\lim_{k \rightarrow \infty} \tilde{C}_{\pi^k}(T_k) \leq \lim_{k \rightarrow \infty} \tilde{C}_{\pi'}(T_k) = \tilde{C}_{\pi'}.$$

By Theorem 7,

$$\lim_{k \rightarrow \infty} \tilde{C}_{\pi^k}(T_k) = \tilde{C}_{\pi'}.$$

But $\pi^* \in \Pi^* \Rightarrow \tilde{C}_{\pi'} = \tilde{C}_{\pi^*}$. Hence $\pi' \in \Pi^*$. ■

Theorem 9 $\tilde{C}^*(r) = \lim_{T \rightarrow \infty} \tilde{C}^*(r, T)$.

Proof: Follows from Theorem 7 and Lemma 8 as proved in Schochetman and Smith [1986]. ■

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 3. Lemma 10 allows us to restrict the set of potentially optimal initial decisions based on their minimum cost up to a finite time horizon. Theorem 11 expresses a stopping criterion for identifying a forecast horizon. Theorems 12 and 13 state that the stopping criterion will be satisfied if the optimal initial decision is unique. For the remainder of the paper we assume $r > \gamma$ is fixed.

Define:

$${}_i\tilde{C}^*(T) = \min\{\tilde{C}_{\pi}(T) | \pi \in \Pi, \pi_1 = \pi_1^i\}.$$

${}_i\tilde{C}^* = \min\{\tilde{C}_{\pi} | \pi \in \Pi, \pi_1 = \pi_1^i\}$. These minima exist since $\{\pi \in \Pi | \pi_1 = \pi_1^i\}$ is compact.

$$\hat{C}(T) = \min\{{}_i\tilde{C}^*(T) | i : \pi_1^i \notin \Pi_1^*(T)\}.$$

$$\hat{C} = \min\{{}_i\tilde{C}^* | i : \pi_1^i \notin \Pi_1^*\}.$$

Lemma 10 *If for any strategy $\hat{\pi}$, $\tilde{C}_\#(T) - \tilde{C}^*(T) > 2a(T)$ then $\hat{\pi}_1 \notin \Pi_1^*(S)$ for all $S \geq T$.*

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

$$\tilde{C}_\pi(S) \leq \tilde{C}_\pi(T) + a(T), \text{ by Corollary 2.}$$

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

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

We also have

$$\tilde{C}_\#(S) \geq \tilde{C}_\#(T) - a(T) \text{ by Corollary 2.}$$

Then, by subtracting,

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

We generalize two theorems of Bès and Sethi [1985] to continuous time.

Theorem 11 *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: Suppose for $S > T$, $\exists \pi^S \in \Pi$ with $\pi_1^S \neq \pi_1^*(T)$. By definition, $\tilde{C}_{\pi^S}(T) \geq \hat{C}(T)$. Therefore,

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

Then by Lemma 10, $\pi_1^S \notin \Pi_1^*(S)$. Hence $\Pi_1^*(S) \subseteq \Pi_1^*(T)$. \blacksquare

Theorem 12 *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 8, \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. ■

Now let $b(t)$ be any function satisfying $b(t) > 0$, for $t \geq 1$, and $b(t) \rightarrow 0$ as $t \rightarrow \infty$. Bès and Sethi [1986] prove the following result for discrete time and discrete decision sets. Schochetman and Smith [1986] extend it to continuous time and continuous decision spaces. The version here applies to continuous time and discrete decision sets.

Theorem 13 *If $|\Pi_1^*| = 1$, then there exists a positive integer N satisfying $\hat{C}(N) - \tilde{C}^*(N) > b(N)$.*

Proof: Since $|\Pi_1^*| = 1$, $\exists \pi^* \in \Pi^*$ and T_0 such that $\Pi_1^* = \{\pi_1^*\}$ and $\pi_1^*(T) = \pi_1^*$, $T \geq T_0$ (by Theorem 12). Suppose $\hat{C}(T) - \tilde{C}^*(T) \leq b(T)$, $\forall T \geq T_0$ (in particular, for all integers $N \geq T_0$). For each N , $\exists \pi^N$ such that $\tilde{C}_{\pi^N} = \hat{C}(N)$ and $\pi_1^N \neq \pi_1^*$. Then

$$0 \leq \tilde{C}_{\pi^N} - \tilde{C}^*(N) \leq b(N) \quad \forall \text{ integers } N \geq T_0.$$

Since Π is compact, we can extract from $\{\pi^N\}$ a subsequence $\{\pi^{N_n}\}$ such that $\pi^{N_n} \rightarrow \pi$. Then $\exists n_0$ sufficiently large such that $n \geq n_0$ implies $N_n \geq T_0$ so that

$$0 \leq \tilde{C}_{\pi^{N_n}}(N_n) - \tilde{C}^*(N_n) \leq b(N_n).$$

But $b(N_n) \rightarrow 0$, $\tilde{C}_{\pi^{N_n}}(N_n) \rightarrow \tilde{C}_\pi$ by Theorem 7, and $\tilde{C}^*(N_n) \rightarrow \tilde{C}^*$ by Theorem 9. Then $\tilde{C}_\pi = \tilde{C}^*$. Hence $\pi \in \Pi^*$ and $\pi_1 = \pi_1^*$. Since $\lim_{n \rightarrow \infty} \pi^{N_n} = \pi$, $\exists n_1$ such that $\forall n \geq n_1$, $\rho(\pi^{N_n}, \pi) < \frac{1}{2}$. Thus $\pi_1^{N_n} = \pi_1 = \pi_1^*$, which is a contradiction. ■

We can sum up our results so far:

Theorem 14 *If Π_1^* is a singleton then $\exists T > 0$ such that T is a forecast horizon.*

Proof: Follows from Theorems 11, 12, 13, with $b(t) = a(t)$. ■

Theorem 14 was also proved by Bean and Smith [1984]. But from Theorems 11, 12, and 13 we also get an algorithm guaranteed to find a forecast horizon whenever $|\Pi_1^*| = 1$:

Algorithm

- (1) *Set $T \leftarrow 1$.*
- (2) *Solve the T -horizon problem to get $\Pi_1^*(T)$, $\tilde{C}^*(T)$, and $\hat{C}(T)$.*
- (3) *If $|\Pi_1^*(T)| = 1$ and $\hat{C}(T) - \tilde{C}^*(T) > 2a(T)$ then STOP: T is a forecast horizon. Else set $T \leftarrow T + 1$ and go to (1).*

Theorem 15 *If Π_1^* is a singleton then the Algorithm will terminate with an infinite-horizon optimal initial decision for finite T .*

Proof: Follows from Theorems 11, 12, 13. ■

6 Degeneracy Resolution

The results in Section 5 show that we can find a forecast horizon if the optimal initial decision is unique. Degeneracy can occur if there is more than one optimal initial decision. 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 decision 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 decision has infinite horizon discounted cost nearly as small as that of the infinite horizon optimal first decision.

We make some additional definitions:

$$G = \{l : \pi_1^l \in \Pi_1^*\}.$$

$\epsilon = (\epsilon_1, \epsilon_2, \dots, \epsilon_k)$ is a perturbation vector.

Let $C_\pi^\epsilon(t) = C_\pi(t) + \epsilon_i$, $t > 0$, if $\pi_1 = \pi_1^i$; that is, add an instantaneous, discrete cost of ϵ_i for π_1^i at time 0.

$\tilde{C}_\pi^\epsilon(r)$, $\tilde{C}^{\epsilon*}$, \hat{C}^ϵ are defined as for $C_\pi(t)$.

$\Pi^{\epsilon*} = \{\pi : \pi = \arg \min_{\pi \in \Pi} \tilde{C}_\pi^\epsilon(r)\}$ = the set of optimal strategies for the perturbed problem.

$\Pi_1^{\epsilon*}$ = the set of initial decisions corresponding to $\pi \in \Pi^{\epsilon*}$.

$$G^\epsilon = \{l : \pi_1^l \in \Pi_1^{\epsilon*}\}.$$

$\hat{\epsilon} = \hat{C}^\epsilon - \tilde{C}^{\epsilon*}$ = the minimum penalty for choosing a suboptimal initial decision in the perturbed problem.

Adjust M if necessary so that

$$K_\pi(t) \leq (M + \hat{\epsilon})e^{\gamma t}$$

$$R_\pi(t) \leq (M + \hat{\epsilon})e^{\gamma t}, \text{ for } t \geq 0.$$

Lemma 16 *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 $i \in G$, $\tilde{C}^*(j) - \tilde{C}^*(i) \geq \bar{\epsilon}$. Then

$$\tilde{C}^{\epsilon*}(j) - \tilde{C}^{\epsilon*}(i) = \tilde{C}^*(j) - \tilde{C}^*(i) + \epsilon_j - \epsilon_i \geq \bar{\epsilon} + \epsilon_j - \epsilon_i.$$

But $0 < \epsilon_i < \hat{\epsilon} \forall i \Rightarrow |\epsilon_j - \epsilon_i| < \bar{\epsilon}$. Hence, $\tilde{C}^{\epsilon*}(j) - \tilde{C}^{\epsilon*}(i) > 0$. 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 $\tilde{C}^*(i) = \tilde{C}^*(j)$. But $\epsilon_i \neq \epsilon_j \Rightarrow \tilde{C}^{\epsilon*}(i) \neq \tilde{C}^{\epsilon*}(j)$. This is a contradiction. ■

Theorem 17 *Let costs be perturbed as in Lemma 16. Then the Algorithm will terminate in finite time.*

Proof: After the perturbation, $\Pi_1^{\epsilon*}$ is a singleton. The result follows from Theorem 15. ■

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 a consequence of Lemma 10, we can find a near-optimal forecast horizon.

Let $P(T) = \{i | \tilde{C}^*(T) - \tilde{C}^*(i) \leq 2a(T)\}$ be the set of potentially optimal initial decisions. By Lemma 10, if for any T , $j \notin P(T)$, then π_1^j is not an optimal initial decision for the infinite horizon problem.

We say that T as 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 18 *If $T_\epsilon \geq \frac{1}{r-\gamma} \log \frac{4rM}{(r-\gamma)\epsilon}$ then T_ϵ is an ϵ -forecast horizon.*

Proof: For any π ,

$$\tilde{C}_\pi \leq \tilde{C}_\pi(T) + a(T) \text{ by Theorem 1.}$$

Also,

$$\tilde{C}^* \geq \tilde{C}^*(T) - a(T) \text{ by Theorem 3.}$$

For any $S \geq T$, let $\pi^S \in \Pi^*(S)$. Then $\pi_1^S = \pi_1^i$ for some $i \in P(T)$ by Lemma 10.

Then

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

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

Note that by Theorem 3, $\tilde{C}^* \leq a(0)$. 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 discount rate used. These properties are consistent with those of the ϵ -forecast horizon found by Bean and Smith [1985] for capacity expansion problems.

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.
- Bès, C. and S. Sethi [1986], "Concepts of Forecast and Decision Horizons: Applications to Dynamic Stochastic Optimization Problems," Technical Report.
- Charnes, A., J. Dreze 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," Technical Report 84-24, Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, MI 48109. Submitted to **Operations Research**.
- 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.
- Modigliani, F. and F. Hohn [1955], "Production Planning Over Time and the Nature of the Expectation and Planning Horizon," **Econometrica**, Vol. 23, pp. 46-66.
- Schochetman, I. and R. L. Smith [1986], "Infinite Horizon Optimization Over Continuous Policy Spaces," in preparation.