

Infinite Horizon Production Scheduling in Time-Varying Systems Under Stochastic Demand

Torpong Cheevaprawatdomrong, Robert L. Smith

Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, Michigan 48109,
{tonychee@yahoo.com, rlsmith@umich.edu}

We consider infinite horizon production scheduling under stochastic demand. All problem data are allowed to vary across periods, including demand distributions, costs, and revenues. A forecast horizon, when it exists, is a finite problem horizon with the property that the corresponding first-period optimal production decision remains optimal regardless of demand and cost projections beyond this horizon. Thus, a forecast horizon allows us to reduce the amount of future data we need to forecast to solve for an optimal first decision for the infinite horizon problem. In this paper, we establish the existence of a forecast horizon under the assumptions that (1) costs and revenues are time-varying linear, and (2) demand is never eventually zero. A key result for establishing the existence and computation of forecast horizons is the monotonicity, and hence convergence, of optimal first-period policies as the horizon increases of finite horizon versions of the infinite horizon problem. A closed-form formula is provided for computing a forecast horizon that depends *only* on the discount factor and uniform upper and lower bounds on demand and unit production and inventory holding costs. In particular, its value is independent of, and determined in advance of, forecasting the demand distribution. We show that the effect of uncertainty in demand is to increase the forecast horizon associated with a deterministic problem by a constant plus a factor equal to one plus the ratio of these upper to lower bounds on per-period demand. The associated forecast horizon can be surprisingly short, even a few days, when the inventory costs are high.

Subject classifications: Dynamic programming: applications. Production scheduling: applications, approximations.

Area of review: Manufacturing, Service, and Supply Chain Operations.

History: Received May 2001; revision received December 2002; accepted January 2003.

1. Introduction

Production scheduling problems with long but indefinite horizons are often modeled as infinite horizon problems. In the time-varying case, a possible way to solve these problems is through use of the forecast horizon approach. This allows for the determination of a first optimal decision by considering just a finite horizon of forecasted data. That is, we attempt to find a finite horizon with the property that data beyond this horizon do not affect the optimality of the first-period policy. This can resolve the following dilemma. On the one hand, we need a sufficient amount of data so that our decision is not shortsighted. On the other hand, we would like to minimize the amount of forecasting we need to do because it is usually expensive and difficult to justify. Smith and Zhang (1998) have recently established the existence of a forecast horizon for the production planning problem in time-varying systems with convex costs and deterministic demand. They also provide closed-form formulas for the calculation of forecast horizons. Our goal is to extend these results to the stochastic demand case.

There is a large literature focusing on the production scheduling problem with time-varying stochastic demand and stationary linear costs. Morton and Pentico (1995) provide near-myopic bounds and heuristics for finite horizon problems. Lovejoy (1992) presents bounds and stopping times for using myopic policies in the case where the

demand in each period can depend upon the previous period's demand. Morton (1978) provides optimal cost and policy bounds that are monotone in horizon length, as well as conditions under which policies for each period of a finite horizon problem converge to a unique optimal policy of the infinite horizon problem. Kleindorfer and Kunreuther (1978) prove that optimal policies are monotone in demand, and provide bounds on policies as well as an algorithm for finding a forecast horizon when it exists. Karlin (1960a, b) shows, in the case where the purchasing cost is stationary convex, that optimal policies are monotone in demand, and optimal policies are myopic if the demand is stochastically increasing.

In the case where costs are nonstationary, Alden and Smith (1992) provide value error bounds for solving finite horizon solutions of the underlying Markov decision-process problem (see also Hernandez-Lerma and Lasserre 1990 for an extension of this work to Markov control policies). Sethi and Cheng (1997) extend (S, s) policy optimality to the Markovian demand infinite horizon case. For nonstationary linear costs, Federgruen and Tzur (1996) provide an algorithm for detecting a minimal forecast horizon, if it exists, in a problem with restricted ordering policies. Veinott (1965) presents conditions under which myopic policies are optimal. Sobel (1981) discusses conditions for a Markov decision process to have a myopic optimal solution. Topkis (1998) proves that in a finite horizon prob-

lem, optimal productions are monotone in demand as costs remained fixed. Zhang (1998) establishes the existence of a solution horizon by using monotonicity of the optimal policy under the assumption that the demand is known in each period before we make the decision for that period.

Garcia and Smith (2000b) establish existence and discovery of forecast horizons for a general dynamic optimization problem by assuming that there exists a forecast index such that the first-period optimal actions are monotonically increasing in that index. Garcia and Smith (2000a) use a similar approach to that in Garcia and Smith (2000b) to establish the existence and discovery of forecast horizons in the stochastic demand production scheduling problem under the assumption that for any fixed finite planning horizon, there exist first-period optimal solutions that are monotone with respect to stochastic demand. Such policy monotonicity has been established under many conditions (see, for example, Karlin 1960a, Kleindorfer and Kunreuther 1978, Morton 1978, and Topkis 1998). However, a key claim in Garcia and Smith (2000b), that the optimal first production of the N -horizon problem is equal to that of the $(N + 1)$ -horizon problem with zero demand in the last period, is not true for stochastic demand problems. This fact invalidates the opportunity to conclude policy monotonicity by embedding an N -horizon problem within an $(N + 1)$ -horizon problem, thus rendering their existence proof invalid for the stochastic demand problem. Also, the optimal first decision of the N -horizon problem is not necessarily a lower bound to that of the infinite horizon problem, so that the algorithm proposed may prematurely stop with a horizon that fails to be a forecast horizon.

In this paper, we generalize Smith and Zhang (1998) to the case where demand is stochastic. We allow costs to be time varying but linear. Also, we do not make the assumption that demand is known prior to deciding on production, as in Zhang (1998). We present a forward algorithm that is guaranteed to find a forecast horizon. In particular, we prove that it will finitely terminate under the regularity condition that minimum demand in a period is strictly positive infinitely often. We also provide a simple formula for the ratio of a forecast horizon length for the stochastic demand problem to that of the deterministic demand problem in terms of lower bounds on the minimum per-period demands and upper bounds on the maximum per-period demands. This result, together with the closed-form upper bound for a forecast horizon for the deterministic demand problem presented in Smith and Zhang (1998), provides a closed-form forecast horizon for the stochastic demand problem.

This paper is organized as follows. The infinite horizon production scheduling problem is formulated in §2. Section 3 introduces a finite horizon version of the original infinite horizon problem that provides a lower bound on the first period's optimal production decision. The existence of a forecast horizon is established by showing that

the initial policies of the finite horizon problems monotonically increase, and hence eventually agree with an infinite horizon optimal policy as the horizon lengthens. Section 4 introduces another finite horizon version (the natural truncated version) of the infinite horizon problem and establishes that policies monotonically decrease as the horizon lengthens. An algorithm is provided for detecting when these upper and lower bounds agree, thus resulting in discovery of an infinite horizon optimal initial production level. A computable bound on the value of a horizon for which this stopping criterion is met is provided. This forecast horizon is provided as a closed-form formula for which illustrative numerical values are provided. The resulting horizons can be surprisingly short, thus simultaneously resulting in minimal data one needs to forecast to compute an optimal first-period production decision, as well as resulting in an efficient forward algorithm for its computation.

2. The Infinite Horizon Problem (L)

We extend the production scheduling problem formulation in Chapter 6 in Denardo (1982) to the case where the costs as well as the demand distributions are allowed to vary across periods. We begin each period by observing an integer-valued beginning inventory level for each period $n = 1, 2, 3, \dots$. We then decide the integer-valued production level to bring inventory up to its after-production level. Production is assumed to be instantaneous. After that, we pay the inventory holding cost to carry the after-production inventory to the end of the period, at which time we meet the stochastic demand. We satisfy as much demand as possible, but no backlogging is allowed, so that we are in the sales-lost case. After satisfying demand, if there is inventory left, then it becomes the beginning inventory of the next period.

For each period $n \geq 1$,

i_n = the beginning inventory i_n for period n where $i_n \geq 0$ integer and the beginning inventory of period 1, i_1 , is given and fixed.

j_n = the after-production inventory j_n for period n where $j_n \geq i_n$ integer.

c_n = the unit production cost c_n for period n where $c_n > 0$.

h_n = the unit inventory holding cost h_n for period n where $h_n > 0$.

r_n = the unit sale price r_n for period n where $r_n > 0$.

D_n = the random integer demand D_n for period n .

\underline{d}_n = a lower bound \underline{d}_n on demand for period n so that $D_n \geq \underline{d}_n$ for all n with probability one.

\bar{d}_n = an upper bound \bar{d}_n on demand for period n so that $D_n \leq \bar{d}_n$ for all n with probability one.

α = the discount factor, $0 \leq \alpha < 1$.

ASSUMPTION 1. *Production levels are uniformly bounded; i.e., there exists $\bar{p} < \infty$ such that $0 \leq j_n - i_n < \bar{p}$ for all n .*

ASSUMPTION 2. *Demands are uniformly bounded; i.e., there exists $\bar{d} < \infty$ such that $\bar{d}_n < \bar{d}$ for all n .*

ASSUMPTION 3. *Marginal production and inventory holding costs and revenues are uniformly bounded; i.e., there exist $\bar{c} < \infty$, $\bar{h} < \infty$, and $\bar{r} < \infty$ such that $c_n < \bar{c}$, $h_n < \bar{h}$, and $r_n < \bar{r}$ for all n .*

ASSUMPTION 4. *For each period n , it is profitable to produce and hold to satisfy the demand in the same period; i.e., $\alpha r_n > c_n + h_n$. In addition, it is not profitable to lower the next period production cost by not satisfying current demand; i.e., $r_n > c_{n+1}$.*

Let $Z_n(i_n, j_n)$ denote the present value at the beginning of period n of the expected net profit in period n , if we begin period n observing a beginning inventory i_n and decide to produce to bring the inventory level up to j_n , $j_n \geq i_n$. We have

$$Z_n(i_n, j_n) = -c_n(j_n - i_n) - h_n j_n + \alpha r_n E[\min\{j_n, D_n\}].$$

In addition, if $D_n = d_n$, then $i_{n+1} = (j_n - d_n)^+$, where $(x)^+ = \max\{0, x\}$.

A strategy $\pi = (\pi_1, \pi_2, \pi_3, \dots)$ is a sequence of policies, one for each period where a policy π_n for period n is a column vector of after-production inventory decisions, one for each feasible state, i.e., possible beginning-inventory level. We denote by $\pi_n(i_n)$ the after-production inventory level given by strategy π when we are in state i_n beginning period n . Π denotes the set of all feasible strategies. By Assumption 1, the set of feasible policies is finite in each period n , so that Π is compact in the product topology of componentwise convergence.

Given a strategy $\pi \in \Pi$, let $P^m(\pi)$ denote the probability transition matrix for inventory levels beginning in period 1 and ending at the beginning of period m under π . Let $Z_n(\pi)$ denote the column vector whose i th component is given by $Z_n(i, \pi_n(i))$, i.e., the conditional expected present value beginning period n of the net profit of producing to inventory level $\pi_n(i)$, given that the process enters inventory level i at the beginning of period n . Letting $Z^m(\pi)$ denote the expected total discounted net profit incurred in periods 1 through m under strategy π , we have

$$Z^m(\pi) = \sum_{n=1}^m \alpha^{n-1} P^n(\pi) Z_n(\pi).$$

Finally, let $Z(\pi)$ be the present worth of expected net profit over the infinite horizon under strategy π . Then,

$$Z(\pi) = \sum_{n=1}^{\infty} \alpha^{n-1} P^n(\pi) Z_n(\pi).$$

We would like to find a feasible strategy that maximizes the expected total discounted net profit over the infinite horizon. Our infinite horizon problem (L) is then

$$\max_{\pi \in \Pi} \{Z(\pi)\}.$$

Because all per-period net profits are uniformly bounded, the discount factor is less than one, and the set of feasible strategies is compact, the maximum above is attained (Bean and Smith 1984).

3. Existence of Forecast Horizons

3.1. The Finite Horizon Lower Bound Problem (L^N)

In this subsection, we construct an N -horizon lower bound problem (L^N) for (L) , and formulate it as a dynamic programming problem. We call this problem a lower bound problem because it can be shown in later subsections that in each period, the smallest optimal decision for each state in this N -horizon lower bound problem is a lower bound for optimal decisions of an infinite horizon problem.

The N -horizon lower bound problem (L^N) is an N -horizon problem in which the problem data from periods 1 through $N - 1$ is the same as that of the infinite horizon problem. However, the undiscounted net profit in period N is replaced by a terminal value depending only on the state i_N that we enter period N , independent of the infinite horizon problem. We will assume without loss of generality that excess inventory over demand is to be sold at the next-period production cost, which is then recovered by production in the next period. This pair of cost and revenue flows will exactly cancel each other out. In short, we must pay the production cost for the beginning inventory of each period.

A strategy $\pi^N = (\pi_1^N, \pi_2^N, \dots, \pi_{N-1}^N)$ is a sequence of policies, one for each period up to period $N - 1$. Because the net profit for period N is independent of the decision taken, we do not include a policy for period N . However, it is sometimes convenient to view π^N as its $N - 1$ policies followed by an arbitrary extending sequence of feasible policies over the infinite horizon. We take these to be policies resulting in zero production. (The context within which we invoke the symbol π^N should make clear which interpretation we are placing on it.) A policy π_n^N for period n is a column vector of decisions, one for each state. Thus, $\pi_n^N(i_n)$ is the decision provided by strategy π^N if we are in state i_n in period n . Π^N is the set of all possible strategies, which is finite by Assumption 1.

Given a strategy π^N , let $P^m(\pi^N)$ denote the probability transition matrix starting in period 1 and ending at the beginning of period m under π^N , $m = 1, 2, \dots, N$.

We now define the terminal value function for period N . We set the undiscounted profit entering state i in period N to be $-(c_N + \bar{h}/(1 - \alpha))i$, independent of the decision taken in period N . Note that this represents the buy-back cost c_N of inventory “sold” at the end of the previous period plus the cost of carrying that inventory indefinitely over the infinite horizon. This cost, together with no prospect of selling this inventory, discourages inventory buildup for period N , thus lowering the optimal production level for period 1 as intended.

Let $Z^N(\pi^N)$ denote the expected discounted net profit incurred in period 1 through N under π^N . Then,

$$Z^N(\pi^N) = \sum_{n=1}^{N-1} \alpha^{n-1} P^n(\pi^N) Z_n(\pi^N) + \alpha^{N-1} P^N(\pi^N) \tilde{Z},$$

where \tilde{Z} is the vector whose i th component is the terminal value $-(c_N + \bar{h}/(1 - \alpha))i$.

We would like to find a feasible strategy that maximizes the expected discounted net profit for this problem over its horizon N . The N -horizon lower bound problem (L^N) is then

$$\max_{\pi^N \in \Pi^N} \{Z^N(\pi^N)\}.$$

Because expected discounted net profit is finite and the set of feasible strategies is finite, the maximum above is again attained. Because (L^N) is a finite horizon problem, we can solve (L^N) by employing the following dynamic programming (DP) functional equations.

Let $F_n^N(y)$ denote the maximum expected discounted net profit from period n through N if we choose to produce the amount of product y in period n . In addition, let $V_n^N(x)$ be the maximum expected discounted net profit from period n through N if we begin period n at inventory level x , or equivalently (because the previous periods “sold” ending inventory is produced back), if we need to produce at least x ; i.e.,

$$V_n^N(x) = \max_{y: y \geq x} \{F_n^N(y)\}, \quad n < N \quad \text{and} \quad (1)$$

$$V_n^N(x) = -\left(c_N + \frac{\bar{h}}{1 - \alpha}\right)x, \quad (2)$$

where for all $n < N$,

$$F_n^N(y) = M_n(y) + \alpha E[V_{n+1}^N(y - D_n)^+] \quad \text{and} \quad (3)$$

$$M_n(y) = -(c_n + h_n)y + \alpha r_n E[\min\{y, D_n\}] + \alpha c_{n+1} E[(y - D_n)^+] \quad (4)$$

$$= (\alpha r_n - c_n - h_n)y - \alpha(r_n - c_{n+1})E[(y - D_n)^+] \quad (5)$$

because $\min\{y, D_n\} = y - (y - D_n)^+$. We can interpret $M_n(y)$ as the expected discounted net profit incurred during period n if we decide to produce the amount y at the beginning of period n .

3.2. The Optimal Policy Structure for the N -Horizon Lower Bound Problem (L^N)

In this subsection, we show that optimal policies are in the form of a threshold produce-up-to level. We now recall Lemma 5 in Supplement 2 in Denardo (1982).

LEMMA 1. *Let h be a convex function on the interval S , and let g be a convex function on an interval T that contains $\{h(x) : x \in S\}$. If g is nondecreasing on T , then $g[h(x)]$ is convex on S .*

PROOF. See Denardo (1982). \square

LEMMA 2. *If $h(x)$ is convex and $g(x)$ is nonincreasing concave, then $g[h(x)]$ is concave.*

PROOF. Note that $-g(x)$ is nondecreasing convex and invoke Lemma 1. \square

LEMMA 3. *If X is a random variable and $f(y, x)$ is concave in y for each fixed value x , then $E[f(y, X)]$ is concave.*

PROOF. Follows from the fact that $E[f(y, X)]$ is a convex combination of concave functions $f(y, x)$. \square

LEMMA 4. *For all $n < N$, $M_n(y)$ is concave.*

PROOF. For all $n < N$, to show that $M_n(y)$ is concave, it suffices to show that each of its terms in (5) is concave. Its first term $(\alpha r_n - c_n - h_n)y$ is linear, and thus concave. By Lemma 3, $E[(y - D_n)^+]$ is concave because $[y - x]^+$ is convex in y for each x . By Assumption 4, $r_n \geq c_{n+1}$, and we get $-\alpha(r_n - c_{n+1}) \leq 0$. Therefore, its second term $-\alpha(r_n - c_{n+1})E[(y - D_n)^+]$ is concave. Hence, $M_n(y)$ is concave in y . \square

THEOREM 1. *In the N -horizon lower bound problem (L^N), there exist integers \underline{S}_n^N and \bar{S}_n^N for all $n \leq N$ such that the following equivalence holds:*

$$F_n^N(s) = \max_{y: y \geq 0} \{F_n^N(y)\} \quad \text{and} \quad (6)$$

$$V_n^N(x) = \begin{cases} F_n^N(s), & x \leq s, \\ F_n^N(x), & x > s, \end{cases} \quad (7)$$

if and only if $\underline{S}_n^N \leq s \leq \bar{S}_n^N$.

PROOF. We will prove this theorem by inductively showing that for all n , $V_n^N(x)$ is nonincreasing concave. From (2), $V_N^N(x)$ is nonincreasing concave.

Now suppose for some $n < N$, $V_{n+1}^N(x)$ is nonincreasing concave. To show that $F_n^N(y) = M_n(y) + \alpha E[V_{n+1}^N(y - D_n)^+]$ is concave, it suffices to show that each component is concave. $M_n(y)$ is concave by Lemma 4. For each fixed x , because $(y - x)^+$ is convex and $V_{n+1}^N(x)$ is nonincreasing concave by assumption, Lemma 2 implies that $V_{n+1}^N(y - x)^+$ is concave. Thus, $E[V_{n+1}^N(y - D_n)^+]$ is concave by Lemma 3. Hence, $F_n^N(y)$ is concave.

Because $F_n^N(y)$ is concave in y and, by Assumption 1, the maximum in (6) is over a finite set of integers, we have that there exist integers \underline{S}_n^N and \bar{S}_n^N such that s satisfies $\underline{S}_n^N \leq s \leq \bar{S}_n^N$ if and only if s satisfies (6). Because $F_n^N(y)$ is concave, it is nondecreasing when $y \leq s$ and nonincreasing when $y > s$. Thus, (7) follows. To complete the induction, from (7), we have that $V_n^N(x)$ is nonincreasing concave. \square

From Theorem 1, we conclude that \underline{S}_n^N and \bar{S}_n^N are threshold values representing, respectively, the smallest and largest produce-up-to quantities optimal for entering period n in the N -horizon lower bound problem (L^N).

3.3. Policy Monotonicity (L^N)

In this subsection, we prove that both the minimum and the maximum optimal produce-up-to levels of the N -horizon lower bound problem (L^N) are monotone nondecreasing in the horizon length N .

Let $f_n^N(x)$ denote the marginal optimal expected discounted net profit of producing x units with no beginning inventory at the beginning of period n ; i.e.,

$$f_n^N(x) = F_n^N(x) - F_n^N(x-1), \quad x \geq 1.$$

Also, let $v_n^N(x)$ denote the marginal optimal expected discounted net profit of starting period n with beginning inventory x ; i.e.,

$$v_n^N(x) = V_n^N(x) - V_n^N(x-1), \quad x \geq 1.$$

We will set $v_n^N(x) = V_n^N(x)^+ - V_n^N(x-1)^+$. We have, for all $n < N$,

$$\begin{aligned} f_n^N(x) &= F_n^N(x) - F_n^N(x-1) \\ &= M_n(x) + \alpha E[V_{n+1}^N(x - D_n)^+] - M_n(x-1) \\ &\quad + \alpha E[V_{n+1}^N(x-1 - D_n)^+] \\ &= m_n(x) + \alpha E[v_{n+1}^N(x - D_n)], \end{aligned} \quad (8)$$

where

$$m_n(x) = M_n(x) - M_n(x-1) \quad (9)$$

$$\begin{aligned} &= -(c_n + h_n)x + \alpha r_n E[\min\{x, D_n\}] \\ &\quad + \alpha c_{n+1} E[(x - D_n)^+] - (c_n + h_n)(x-1) \\ &\quad - \alpha r_n E[\min\{x-1, D_n\}] - \alpha c_{n+1} E[(x-1 - D_n)^+] \\ &= -(c_n + h_n) + \alpha r_n w_n(x) + \alpha c_{n+1} u_n(x) \quad \text{by (4)} \quad (10) \end{aligned}$$

$$= \alpha r_n - c_n - h_n - \alpha(r_n - c_{n+1})u_n(x) \quad \text{by (5)}, \quad (11)$$

where

$$w_n(x) = E[\min\{x, D_n\} - \min\{x-1, D_n\}] \geq 0 \quad \text{and} \quad (12)$$

$$u_n(x) = E[(x - D_n)^+ - (x-1 - D_n)^+] \leq 1. \quad (13)$$

We can interpret $w_n(x)$ as the marginal expected demand satisfied in period n if we decide to produce x . In addition, we can interpret $u_n(x)$ as the marginal expected beginning inventory of period $n+1$ if we decide to produce x in period n .

Because $F_n^N(x)$ is concave from the proof of Theorem 1, $f_n^N(x)$ is nonincreasing. Also, by Theorem 1, $\underline{s}_n^N \leq s \leq \bar{s}_n^N$ just in case s attains the maximum in (6). Hence,

$$\begin{aligned} f_n^N(x) &> 0 \quad \forall 0 < x \leq \underline{s}_n^N, \\ f_n^N(x) &= 0 \quad \forall \underline{s}_n^N < x \leq \bar{s}_n^N, \\ f_n^N(x) &< 0 \quad \forall \bar{s}_n^N < x. \end{aligned} \quad (14)$$

In addition, from (7), we have for $\underline{s}_n^N \leq s \leq \bar{s}_n^N$,

$$v_n^N(x) = \begin{cases} 0, & x \leq s, \\ f_n^N(x), & x > s. \end{cases} \quad (15)$$

LEMMA 5. For all x , $v_n^N(x) \leq 0$, $v_n^N(x) \leq f_n^N(x)$, and $v_n^N(x)$ is nonincreasing in x . Furthermore,

$$v_n^N(x) = \begin{cases} 0, & \text{when } f_n^N(x) \geq 0, \\ f_n^N(x), & \text{when } f_n^N(x) < 0. \end{cases} \quad (16)$$

PROOF. From (14) and (15), we have (16). Thus, $v_n^N(x) \leq 0$ and $v_n^N(x) \leq f_n^N(x)$ for all x . Furthermore, we know from the proof of Theorem 1 that $F_n^N(x)$ is concave in x . Hence, $f_n^N(x)$ is nonincreasing in x . Thus, $v_n^N(x)$ is also nonincreasing in x from (15). \square

LEMMA 6. $v_N^{N+1}(x) \geq v_N^N(x)$ for all N and x .

PROOF. Case I: $x \leq \bar{s}_N^{N+1}$. From (15) and (2),

$$v_N^{N+1}(x) = 0 \geq -\left(c_N + \frac{\bar{h}}{1-\alpha}\right) = v_N^N(x).$$

Case II: $x > \bar{s}_N^{N+1}$. We have

$$\begin{aligned} v_N^{N+1}(x) &= f_N^{N+1}(x) \quad \text{from (15)} \\ &= -(c_N + h_N) + \alpha r_N w_N(x) + \alpha c_{N+1} u_N(x) \\ &\quad - \alpha \left(c_{N+1} + \frac{\bar{h}}{1-\alpha}\right) u_N(x) \\ &\quad \text{from (8), (10), (13), and (2)} \end{aligned}$$

$$\geq -(c_N + h_N) - \frac{\alpha \bar{h} u_N(x)}{1-\alpha} \quad \text{by (12)}$$

$$\geq -(c_N + h_N) - \frac{\alpha \bar{h}}{1-\alpha} \quad \text{by (13)}$$

$$\geq -(c_N + \bar{h}) - \frac{\alpha \bar{h}}{1-\alpha}$$

$$= -\left(c_N + \frac{\bar{h}}{1-\alpha}\right)$$

$$= v_N^N(x). \quad \square$$

THEOREM 2. For all N and $n < N$, \bar{s}_n^N and \underline{s}_n^N are monotone nondecreasing in N .

PROOF. Proof by induction. By Lemma 6, $v_N^{N+1}(x) \geq v_N^N(x)$ for all x . Suppose for some $n < N$, $v_{n+1}^{N+1}(x) \geq v_{n+1}^N(x)$ for all x . From (8),

$$\begin{aligned} f_n^{N+1}(y) - f_n^N(y) &= m_n(y) + \alpha E[v_{n+1}^{N+1}(y - D_n)] - m_n(y) - \alpha E[v_{n+1}^N(y - D_n)] \\ &= \alpha E[v_{n+1}^{N+1}(y - D_n) - v_{n+1}^N(y - D_n)]. \end{aligned} \quad (17)$$

By the induction hypothesis, $v_{n+1}^{N+1}(y - D_n) - v_{n+1}^N(y - D_n) \geq 0$ with probability one. Thus, for all y ,

$$f_n^{N+1}(y) \geq f_n^N(y). \quad (18)$$

Then, from (14), $f_n^{N+1}(\underline{s}_n^N) \geq f_n^N(\underline{s}_n^N) > 0$. We have then, by (14), $\underline{s}_n^{N+1} \geq \underline{s}_n^N$.

In addition, $\bar{s}_n^{N+1} \geq \bar{s}_n^N$ because if $\bar{s}_n^{N+1} < \bar{s}_n^N$, then by (14), $f_n^{N+1}(\bar{s}_n^N) < 0$. However, by (18), $f_n^{N+1}(\bar{s}_n^N) \geq f_n^N(\bar{s}_n^N) \geq 0$ by (14) and we have reached a contradiction.

To complete the induction, it is left to show $v_n^{N+1}(x) \geq v_n^N(x)$ for all x .

Case I: $x \leq \bar{S}_n^N \leq \bar{S}_n^{N+1}$. By (15),

$$v_n^{N+1}(x) = 0 = v_n^N(x).$$

Case II: $\bar{S}_n^N < x \leq \bar{S}_n^{N+1}$. By (14), (15), and the fact that $f_n^N(x)$ is nonincreasing from concavity of $F_n^N(x)$,

$$v_n^{N+1}(x) = 0 > f_n^N(\bar{S}_n^N + 1) \geq f_n^N(x) = v_n^N(x).$$

Case III: $x > \bar{S}_n^{N+1} \geq \bar{S}_n^N$. By (15) and (18),

$$v_n^{N+1}(x) = f_n^{N+1}(x) \geq f_n^N(x) = v_n^N(x). \quad \square$$

3.4. Optimal Policy and Value Convergence for the N -Horizon Lower Bound Problem (L^N)

In this section, we show that optimal policies and values of the N -horizon lower bound problem converge to those of the infinite horizon problem. We begin with a lemma.

LEMMA 7. *If $\pi_n^N \rightarrow \pi_n$ as $N \rightarrow \infty$ for all n when $\pi^N \in \Pi^N$, $N = 1, 2, \dots$, then $Z^N(\pi^N) \rightarrow Z(\pi)$ as $N \rightarrow \infty$.*

PROOF. See the Appendix. \square

Note that $\pi_n^N \rightarrow \pi_n$ as $N \rightarrow \infty$ for all n is equivalent to writing $\pi^N \rightarrow \pi$ as $N \rightarrow \infty$ in the product topology of componentwise convergence. Because π^N feasibly extended is in Π for all N and Π is compact and in particular closed, we conclude that π^N being feasible for all N implies its limit π must be in Π , i.e., feasible.

We turn to showing optimal policy convergence. Because Theorem 2 states that S_n^N and \bar{S}_n^N are monotone increasing in N and bounded above by Assumption 1, $\bar{S}_n \equiv \lim_{N \rightarrow \infty} \bar{S}_n^N$ and $\underline{S}_n \equiv \lim_{N \rightarrow \infty} S_n^N$ must exist for all n . Because for all $n = 1, 2, \dots$, both \bar{S}_n^N and S_n^N , $N = n + 1, n + 2, \dots$, is a converging sequence of integers, there exist N_n for each $n = 1, 2, \dots$ such that

$$\bar{S}_n^N = \bar{S}_n \quad \text{and} \quad \underline{S}_n^N = \underline{S}_n \quad \forall N \geq N_n.$$

Then, for all S_n with $\underline{S}_n \leq S_n \leq \bar{S}_n$, we have $S_n = \lim_{N \rightarrow \infty} S_n^N$, where $S_n^N \equiv \underline{S}_n^N + S_n - \underline{S}_n$ for all $N \geq N_n$. Note that $\underline{S}_n^N \leq S_n^N \leq \bar{S}_n^N$ for all $N \geq N_n$ and, hence, is an optimal produce-up-to threshold for the N -horizon lower bound problem, $N \geq N_n$. We conclude that all thresholds S_n between \bar{S}_n and \underline{S}_n are limits of sequence of optimal N -horizon lower bound problem thresholds. It remains to show that these thresholds are infinite horizon optimal, i.e., that optimal policy convergence holds.

THEOREM 3. *Optimal policy convergence holds; i.e., if $\underline{S}_n^N \leq S_n^N \leq \bar{S}_n^N$ for all N and $S_n^N \rightarrow S_n$ as $N \rightarrow \infty$, then S_n is an optimal produce-up-to threshold for the infinite horizon problem (L).*

PROOF. Let $^*\pi_n^N$ represent the policy of the N -horizon lower bound problem corresponding to produce-up-to level S_n^N and let $^*\pi_n$ correspond to the policy for the infinite horizon problem of produce-up-to level S_n . Hence, $^*\pi^N \rightarrow ^*\pi$

as $N \rightarrow \infty$. We have already noted that because $\pi_n^N \rightarrow \pi_n$, the limit policy $^*\pi$ must be infinite horizon feasible; i.e., $^*\pi \in \Pi$. Let $\pi \in \Pi$ and note that π is feasible for the N -horizon problem so that

$$Z^N(^*\pi^N) \geq Z^N(\pi).$$

Hence,

$$\lim_{N \rightarrow \infty} Z^N(^*\pi^N) \geq \lim_{N \rightarrow \infty} Z^N(\pi). \quad (19)$$

But by Lemma 7, because $^*\pi^N \rightarrow ^*\pi$,

$$\lim_{N \rightarrow \infty} Z^N(^*\pi^N) = Z(^*\pi) \quad (20)$$

and also by Lemma 7, because $\hat{\pi}^N \rightarrow \pi$ where $\hat{\pi}^N \equiv \pi$, we get

$$\lim_{N \rightarrow \infty} Z^N(\pi) = Z(\pi). \quad (21)$$

Hence, from (19), (20), and (21) we conclude

$$Z(^*\pi) \geq Z(\pi)$$

for all $\pi \in \Pi$. Hence, $^*\pi$ is infinite horizon optimal. \square

Note from (20) in the proof of the previous theorem we get

$$\lim_{N \rightarrow \infty} Z^N = ^*Z$$

where $^*Z^N \equiv Z^N(^*\pi^N)$ is the optimal value for the N -horizon lower bound problem and $^*Z \equiv Z(^*\pi)$ is the optimal value for the infinite horizon problem. That is, we have demonstrated optimal value convergence as well as optimal policy convergence.

Also note that we have proven that *the infinite horizon problem (L) has an optimal produce-up-to threshold structure*, inherited from the finite horizon problems.

3.5. Existence of Forecast Horizons for (L)

In this subsection, we constructively prove existence of a forecast horizon by showing that a horizon sufficiently long to generate a cumulative demand exceeding the largest optimal initial production level must be a forecast horizon. In particular, we show that \bar{N} is a forecast horizon where

$$\bar{N} = \min \left\{ N : \bar{S}_1^N < \sum_{k=1}^{N-1} \underline{d}_k, N = 2, 3, \dots \right\}. \quad (22)$$

Roughly speaking, if we think of inventory being depleted under a FIFO policy, no demand beyond horizon \bar{N} is optimally satisfied by production in period 1. In this sense, \bar{N} is similar to the deterministic case of Smith and Zhang (1998), where a forecast horizon is the longest horizon over which it is optimal to carry a unit of inventory. The following assumption assures that \bar{N} defined above exists and is finite.

ASSUMPTION 5. Demand is strictly positive infinitely often; i.e., $\underline{d}_n > 0$ infinitely often.

LEMMA 8. In the N -horizon lower bound problem (L^N), for all $1 \leq n \leq N$, if $0 \leq x \leq \underline{d}_n$, then $f_n^N(y)$ for all $y \leq x$ depends only on the problem data in period n .

PROOF. For all $x \leq \underline{d}_n$, from (8), (11), (13), and the fact that $x \leq \underline{d}_n$, we have $f_n^N(x) = \alpha r_n - c_n - h_n$, which depends only on the problem data in period n . \square

Note the importance of the assumption of lost sales in concluding Lemma 8.

LEMMA 9. In the N -horizon lower bound problem (L^N), for all n , if $x > \underline{d}_n$, it is sufficient to know the problem data in period n and to know $f_{n+1}^N(y)$ for all $y \leq x - \underline{d}_n$ to determine the value of $f_n^N(y)$ for all $y \leq x$.

PROOF. For all $x > \underline{d}_n$ and all $y \leq x$, from (8),

$$\begin{aligned} f_n^N(y) &= m_n(y) + \alpha E[v_{n+1}^N(y - D_n)] \\ &= m_n(y) + \alpha E[v_{n+1}^N(y - D_n)^+ - V_{n+1}^N(y - 1 - D_n)^+] \\ &= m_n(y) + \alpha \sum_{k=\underline{d}_n}^{\bar{d}_n} [V_{n+1}^N(y - k)^+ - V_{n+1}^N(y - 1 - k)^+] \\ &\quad \cdot P(D_n = k). \end{aligned}$$

Because $k \geq y$ implies $V_{n+1}^N(y - k)^+ = V_{n+1}^N(y - 1 - k)^+ = V_{n+1}^N(0)$, we have

$$\begin{aligned} f_n^N(y) &= m_n(y) + \alpha \sum_{k=\underline{d}_n}^{y-1} [V_{n+1}^N(y - k)^+ - V_{n+1}^N(y - 1 - k)^+] \\ &\quad \cdot P(D_n = k) \\ &= m_n(y) + \alpha \sum_{k=\underline{d}_n}^{y-1} v_{n+1}^N(y - k) P(D_n = k). \end{aligned}$$

Setting $k' = y - k$, we have

$$f_n^N(y) = m_n(y) + \alpha \sum_{k'=1}^{y-\underline{d}_n} v_{n+1}^N(k') P(D_n = y - k'). \quad (23)$$

Lemma 5 implies that $v_{n+1}^N(k)$ depends on $f_{n+1}^N(k)$ and $k \leq y - \underline{d}_n \leq x - \underline{d}_n$, and thus the result follows. \square

LEMMA 10. For all $y \leq \bar{S}_1^{\bar{N}} + 1$ and $N \geq \bar{N}$, $f_1^N(y)$ is independent of the problem data beyond $\bar{N} - 1$.

PROOF. For all $N \geq \bar{N}$, by applying Lemma 9 recursively, starting with $n = 1$ and ending with $n = \bar{N} - 1$, we have that, in general, it is sufficient to know the problem data in period $1, 2, \dots, N - 2$ and to know $f_{\bar{N}-1}^N(y)$ for all $y \leq \bar{S}_1^{\bar{N}} + 1 - \sum_{k=1}^{\bar{N}-2} \underline{d}_j$ to compute the value of $f_1^N(y)$ for all $y \leq \bar{S}_1^{\bar{N}} + 1$.

However, by the definition of \bar{N} , $\bar{S}_1^{\bar{N}} < \sum_{k=1}^{\bar{N}-1} \underline{d}_k$, which implies $\bar{S}_1^{\bar{N}} + 1 - \sum_{k=1}^{\bar{N}-2} \underline{d}_k \leq \underline{d}_{\bar{N}-1}$. Thus, by Lemma 8, $f_{\bar{N}-1}^N(y)$ for all $y \leq \bar{S}_1^{\bar{N}} + 1 - \sum_{k=1}^{\bar{N}-2} \underline{d}_j$ depends only on the problem data in period $\bar{N} - 1$. Hence, for all $y \leq \bar{S}_1^{\bar{N}} + 1$, $f_1^N(y)$ is independent of the problem data beyond $\bar{N} - 1$. \square

THEOREM 4. We have $\underline{S}_1^N = \bar{S}_1^{\bar{N}}$ and $\bar{S}_1^N = \bar{S}_1^{\bar{N}}$ for all $N \geq \bar{N}$; i.e., \bar{N} is a forecast horizon.

PROOF. From (14), $f_1^N(\bar{S}_1^{\bar{N}} + 1) < 0$. For all $N \geq \bar{N}$, by Lemma 10, $f_1^N(y)$ for all $y \leq \bar{S}_1^{\bar{N}} + 1$ is independent of the problem data beyond \bar{N} . Thus, $f_1^N(y) = f_1^{\bar{N}}(y)$ for all $y \leq \bar{S}_1^{\bar{N}} + 1$. Because from (14) $\bar{S}_1^{\bar{N}}$ and $\underline{S}_1^{\bar{N}}$ can be computed if $f_1^{\bar{N}}(y)$, $y \leq \bar{S}_1^{\bar{N}} + 1$ is known, $\bar{S}_1^N = \bar{S}_1^{\bar{N}}$ and $\underline{S}_1^N = \underline{S}_1^{\bar{N}}$ independently of the information beyond \bar{N} . \square

4. Computing Forecast Horizons for (L)

4.1. The N -Horizon Upper Bound Problem

In this subsection, we define an N -horizon upper bound problem (U^N) and formulate it as a dynamic programming problem. Analogous to the lower bound problem (L^N), we call (U^N) an upper bound problem because it will be shown later that the largest optimal decision for each state in each period of the N -horizon upper bound problem is an upper bound to the corresponding optimal decisions of the infinite horizon problem.

An N -horizon upper bound problem (U^N) is an N -horizon problem in which the problem data from periods 1 through $N - 1$ are the same as that of the infinite horizon problem except that there is a terminal value function which gives the undiscounted net profit for period N depending on the state of period N , i_N .

We now define the terminal value function \hat{Z} for the N -horizon upper bound problem (U^N). That is, let \hat{Z} represent the column vector in which the i th component of the vector \hat{Z}_i , is the net profit in period N given that we enter state i at the beginning of period N . For all i , we set

$$\hat{Z}_i = 0.$$

Note that without loss of optimality it will never prove worthwhile for a problem with horizon exceeding N to generate an inventory level greater than that yielded by (U^N). The reason is the c_N received per unit inventory sold at the end of period $N - 1$ in (U^N) places a marginal value on ending inventory of, at most, c_N . It is thus no more costly to produce additional inventory in period N at its cost c_N per unit. The problem (U^N) thus results in a maximal beginning-inventory period N and in a corresponding maximal production level beginning period 1, which upper bounds production for the infinite horizon problem. In summary, as we shall show below, *the ordinary N -horizon truncation of the infinite horizon problem upper bounds optimal first-period production.*

Let $\hat{Z}^N(\pi^N)$ denote the expected present value of the net profit incurred in periods 1 through N under π^N in the N -horizon upper bound problem (U^N). Then,

$$\begin{aligned} \hat{Z}^N(\pi^N) &= \sum_{n=1}^{N-1} \alpha^{n-1} P^n(\pi^N) Z_n(\pi^N) + \alpha^{N-1} P^N(\pi^N) \hat{Z} \\ &= \sum_{n=1}^{N-1} \alpha^{n-1} P^n(\pi^N) Z_n(\pi^N). \end{aligned}$$

We would like to find a feasible strategy that maximizes the expected total discounted net profit over the horizon N . The N -horizon upper bound problem (U^N) is as follows:

$$\max_{\pi^N \in \Pi^N} \{\hat{Z}^N(\pi^N)\}.$$

Because the set of feasible strategies is finite, the maximum above is attained.

By dynamic programming, we can solve (U^N) by solving the optimality equations, $\hat{V}_n^N = \max_{y: y \geq x} \{\hat{F}_n^N(y)\}$, $n \leq N$ as in (1), with (2) replaced by

$$\hat{V}_N^N(x) = 0, \quad 0 \leq x < \bar{d}, \tag{24}$$

which implies

$$\hat{v}_N^N(x) = 0, \quad 0 < x < \bar{d}. \tag{25}$$

Note that hat ($\hat{\cdot}$) in functions indicates that the functions are for the N -horizon upper bound problem (U^N).

THEOREM 5. *In the N -horizon upper bound problem (U^N), there exist integers \underline{T}_n^N and \bar{T}_n^N for all $n \leq N$ such that the following equivalence holds:*

$$\hat{F}_n^N(s) = \max_{y: y \geq 0} \{\hat{F}_n^N(y)\} \quad \text{and} \tag{26}$$

$$\hat{V}_n^N(x) = \begin{cases} \hat{F}_n^N(s), & x \leq s, \\ \hat{F}_n^N(x), & x > s, \end{cases} \tag{27}$$

if and only if $\underline{T}_n^N \leq s \leq \bar{T}_n^N$.

PROOF. Because $\hat{V}_n^N(x)$ is nonincreasing concave, the rest of the proof is similar to that of Theorem 1. \square

LEMMA 11. $\hat{v}_N^{N+1}(x) \leq \hat{v}_N^N(x)$ for all $x > 0$ and $N = 1, 2, \dots$

PROOF. For all $N = 1, 2, \dots$ and $x > 0$, we have by Lemma 5 that

$$\hat{v}_N^{N+1}(x) \leq 0 = \hat{v}_N^N(x). \quad \square$$

THEOREM 6. *For all N and $n < N$, \bar{T}_n^N and \underline{T}_n^N are monotone nonincreasing in N .*

PROOF. The result and proof for the lower bound problem goes through for the upper bound problem if we replace N by $N + 1$ and $N + 1$ by N throughout. For example, $V_n^N(x)$ becomes $\hat{V}_n^{N+1}(x)$, and $v_n^N(x)$ becomes $\hat{v}_n^{N+1}(x)$. Then, by Lemma 11,

$$v_N^N(x) = \hat{v}_N^{N+1}(x) \leq \hat{v}_N^N(x) = v_N^{N+1}(x).$$

The rest of the proof of Theorem 2, then, still holds. In Theorem 2, we have proved that for all $n \leq N$,

$$\bar{S}_n^N \leq \bar{S}_n^{N+1} \quad \text{and} \quad \underline{S}_n^N \leq \underline{S}_n^{N+1},$$

which implies that

$$\bar{T}_n^{N+1} \leq \bar{T}_n^N \quad \text{and} \quad \underline{T}_n^{N+1} \leq \underline{T}_n^N$$

for all $n \leq N$, which in turn implies that \bar{T}_n^N and \underline{T}_n^N are monotone nonincreasing in N . \square

We now introduce optimality equations for the infinite horizon problem. Let $F_n(y)$ denote the maximum expected discounted net profit from period n on in the infinite horizon problem if we choose to produce y in period n . In addition, let $V_n(x)$ be the maximum expected discounted net profit from period n on in the infinite horizon problem if we begin period n at inventory level x . Similarly, $f_n(y) \equiv F_n(y) - F_n(y - 1)$, $y > 0$ and $v_n(x) \equiv V_n(x) - V_n(x - 1)$, $x > 0$. Furthermore, let \underline{S}_n and \bar{S}_n denote the minimum and the maximum produce-up-to levels for the infinite horizon problem, respectively. We have

$$V_n(x) = \max_{y: y \geq x} \{F_n(y)\}, \tag{28}$$

$$F_n(y) = M_n(y) + \alpha E[V_{n+1}(y - D_n)^+] \tag{29}$$

for all $n = 1, 2, 3, \dots$

LEMMA 12. $\underline{S}_n^N \leq \underline{S}_n \leq \underline{T}_n^N$ and $\bar{S}_n^N \leq \bar{S}_n \leq \bar{T}_n^N$, $N \geq n$, $n = 1, 2, 3, \dots$

PROOF. Omitted. \square

4.2. Algorithm for Detecting Forecast Horizons

We begin with a forward algorithm that solves the N -horizon problems (L^N) and (U^N) for ever greater horizons N until and if we get agreement in the largest and smallest corresponding initial production levels for the first period.

Solution Algorithm

Step 0. $N = 1$.

Step 1. Solve (L^N) and (U^N). If $\bar{S}_1^N = \bar{T}_1^N$ and $\underline{S}_1^N = \underline{T}_1^N$, stop.

Otherwise, let $N = N + 1$, then go to Step 1.

The following theorem assures us that if the algorithm stops, then we have indeed solved for the infinite horizon problem's optimal first decision. Moreover, *the optimality of this initial production level is unaffected by data beyond the horizon N^* for which the stopping condition was met.*

THEOREM 7. *If the algorithm stops at $N = N^*$, then N^* is a forecast horizon.*

PROOF. If the algorithm stops at $N = N^*$, by Theorem 2, Theorem 6, and Lemma 12, we have $\bar{S}_1^N = \bar{T}_1^N$ and $\underline{S}_1^N = \underline{T}_1^N$ for all $N \geq N^*$ independently of the information beyond N^* . \square

THEOREM 8. *The algorithm finitely terminates; i.e., the stopping condition is eventually met.*

PROOF. It is sufficient to show that $\bar{S}_1^{\bar{N}} = \bar{T}_1^{\bar{N}}$ and $\underline{S}_1^{\bar{N}} = \underline{T}_1^{\bar{N}}$ where \bar{N} is defined by (22). By Theorem 10, for all $y \leq \bar{S}_1^{\bar{N}} + 1$, $f_1^{\bar{N}}(y)$ is independent of the problem data beyond $\bar{N} - 1$. However, the data from periods 1 through $\bar{N} - 1$ of $(L^{\bar{N}})$ and $(U^{\bar{N}})$ are the same. Thus, $f_1^{\bar{N}}(y) = \hat{f}_1^{\bar{N}}(y)$ for all $y \leq \bar{S}_1^{\bar{N}} + 1$, which implies that $\bar{S}_1^{\bar{N}} = \bar{T}_1^{\bar{N}}$ and $\underline{S}_1^{\bar{N}} = \underline{T}_1^{\bar{N}}$. \square

THEOREM 9. *If $\underline{T}_n^{N'} \leq s \leq \bar{S}_n^{N'}$ for some N' , then s is an infinite horizon optimal produce-up-to level for period n .*

PROOF. Consider all $N \geq N'$. By Theorem 6, \underline{T}_n^N is monotonically decreasing in N . Therefore,

$$\underline{T}_n^N \leq s.$$

By Lemma 12, \underline{T}_n^N is bounded below by \underline{S}_n^N . Thus,

$$\underline{S}_n^N \leq s. \quad (30)$$

By Theorem 2, \bar{S}_n^N is monotonically increasing in N . Therefore,

$$s \leq \bar{S}_n^N. \quad (31)$$

From (30) and (31), for all $N \geq N'$,

$$\underline{S}_n^N \leq s \leq \bar{S}_n^N.$$

Thus, s is an optimal produce-up-to level for period n for all lower bound problems with horizon $N \geq N'$. From the discussion in §3.4, we can construct a converging sequence of optimal strategies for N -horizon lower bound problems, $N = 2, 3, 4, \dots$, such that s is an optimal produce-up-to level for period n of all lower bound problems with horizon $N \geq N'$. Then, s must be an optimal produce-up-to level of the limit of this converging sequence. If s were not infinite horizon optimal, this would contradict Theorem 3. Thus, s is an infinite horizon optimal produce-up-to level for period n . \square

Theorem 9 is useful because it may allow us to obtain in advance an infinite horizon optimal policy not only for the first period, but also for future periods, *before* the stopping rule in the Solution Algorithm is met.

4.3. Forecast Horizon Bounds

In this section, we present a formula for the ratio of the forecast horizon length of the stochastic demand problem to the deterministic demand problem in terms of lower bounds on the minimum and upper bounds on the maximum per-period demands to be encountered. This ratio, together with the closed-form formula for an upper bound on the minimal forecast horizon for the deterministic demand problem presented in Smith and Zhang (1998), results in a closed-form formula for a forecast horizon for our stochastic demand problem. Let N^* be a forecast horizon when the demand

is deterministic, as given in Smith and Zhang (1998). We present their formula here in our notation,

$$N^* = \left\lceil \log_\alpha \left\{ \frac{(1-\alpha)c_1 + \underline{h}}{(1-\alpha)\bar{c} + \underline{h}} \right\} \right\rceil,$$

where $\underline{h} = \inf_n \{h_n\} > 0$ and $\lceil x \rceil$ denotes the smallest integer strictly greater than x . One may interpret N^* as a horizon long enough so that the discounted cost of producing one unit and holding it through period N^* is greater than the discounted cost of producing this unit in period $N^* + 1$; i.e.,

$$c_1 + \sum_{k=1}^{N^*} \alpha^{k-1} \underline{h} > \alpha^{N^*} \bar{c}.$$

Suppose now there exist $\bar{d} < \infty$ and $\underline{d} < \infty$ such that $\bar{d} = \sup_n \{\bar{d}_n\}$ and $\underline{d} = \inf_n \{\underline{d}_n\}$. Let $\theta = \bar{d}/\underline{d}$. Define $N^{**} = 2 + \lceil \theta N^* \rceil$. We will show that N^{**} is a forecast horizon for our stochastic demand problem.

LEMMA 13. *For all $n \leq N$ and $x > \bar{d}_n$, $f_n^N(x) \leq -c_n - h_n + \alpha c_{n+1} + \alpha v_{n+1}^N(x - \bar{d}_n)$.*

PROOF. For all $x > \bar{d}_n$, $x - 1 \geq D_n$ with probability one. From (13), we have

$$u_n(x) = E[(x - D_n)^+ - (x - 1 - D_n)^+] = 1.$$

From (8) and (11), because $u_n(x) = 1$, we have

$$f_n^N(x) = -c_n - h_n + \alpha c_{n+1} + \alpha E[v_{n+1}^N(x - D_n)].$$

However, $v_{n+1}^N(x)$ is nonincreasing in x by Lemma 5, and $x > \bar{d}_n$ by assumption. Thus, with probability one,

$$v_{n+1}^N(x - D_n) \leq v_{n+1}^N(x - \bar{d}_n).$$

Therefore,

$$f_n^N(x) \leq -c_n - h_n + \alpha c_{n+1} + \alpha v_{n+1}^N(x - \bar{d}_n). \quad \square$$

LEMMA 14. *For all $N > N^*$, $\bar{S}_1^N \leq \sum_{k=1}^{N^*} \bar{d}_k$.*

PROOF. Proof by contradiction. Suppose for some $N \geq N^*$ we have $\bar{S}_1^N > \sum_{k=1}^{N^*} \bar{d}_k$. Because $\bar{S}_1^N > \sum_{k=1}^{N^*} \bar{d}_k > \bar{d}_1$, by Lemma 13 we get

$$f_1^N(\bar{S}_1^N) \leq -c_1 - h_1 + \alpha c_2 + \alpha v_2^N(\bar{S}_1^N - \bar{d}_1).$$

From Lemma 5, $v_2^N(\bar{S}_1^N - \bar{d}_1) \leq f_2^N(\bar{S}_1^N - \bar{d}_1)$. Therefore,

$$f_1^N(\bar{S}_1^N) \leq -c_1 - h_1 + \alpha c_2 + \alpha f_2^N(\bar{S}_1^N - \bar{d}_1).$$

Applying the same procedure with $f_1^N(\bar{S}_1^N)$ replaced by $f_2^N(\bar{S}_1^N - \bar{d}_1)$, we have

$$f_1^N(\bar{S}_1^N) \leq -c_1 - h_1 + \alpha c_2 - \alpha c_2 - \alpha h_2 + \alpha^2 c_3 + \alpha^2 f_3^N(\bar{S}_1^N - \bar{d}_1 - \bar{d}_2).$$

Applying the same procedure recursively, we have

$$f_1^N(\bar{S}_1^N) \leq -c_1 - \sum_{k=1}^{N^*} \alpha^{k-1} h_k + \alpha^{N^*} c_{N^*+1} + \alpha^{N^*-1} v_{N^*}^N \left(\bar{S}_1^N - \sum_{k=1}^{N^*} \bar{d}_k \right).$$

From Lemma 5, $v_{N^*}^N(x) \leq 0$ for all x . Thus,

$$f_1^N(\bar{S}_1^N) \leq -c_1 - \sum_{k=1}^{N^*} \alpha^{k-1} h_k + \alpha^{N^*} c_{N^*+1}.$$

By definition of N^* ,

$$f_1^N(\bar{S}_1^N) < 0.$$

This contradicts (14). Therefore, for all $N \geq N^*$, $\bar{S}_1^N \leq \sum_{k=1}^{N^*} \bar{d}_k$. \square

THEOREM 10. N^{**} is a forecast horizon for stochastic demand problems.

PROOF. Because \bar{N} is a forecast horizon for stochastic demand problems, it is sufficient to show

$$\bar{S}_1^{N^{**}} < \sum_{k=1}^{N^{**}-1} \bar{d}_k,$$

which implies by (22) that $N^{**} \geq \bar{N}$. By definition of \bar{d} ,

$$\sum_{k=1}^{N^{**}-1} \bar{d}_k \geq \sum_{k=1}^{N^{**}-1} \bar{d} = \bar{d}(N^{**} - 1). \tag{32}$$

By definition of θ ,

$$\bar{d}(N^{**} - 1) \geq \frac{\bar{d}(N^{**} - 1)}{\theta}. \tag{33}$$

By definition of \bar{d} , $\sum_{k=1}^{N^*} \bar{d}_k / N^* \leq \bar{d}$ so that

$$\frac{\bar{d}(N^{**} - 1)}{\theta} \geq \frac{\sum_{k=1}^{N^*} \bar{d}_k}{N^*} \cdot \frac{N^{**} - 1}{\theta}. \tag{34}$$

Because we define $N^{**} = 2 + \lceil \theta N^* \rceil$, $(N^{**} - 1) / (\theta N^*) > 1$, we get

$$\frac{\sum_{k=1}^{N^*} \bar{d}_k}{N^*} \cdot \frac{N^{**} - 1}{\theta} > \sum_{k=1}^{N^*} \bar{d}_k. \tag{35}$$

However, by Lemma 14,

$$\sum_{k=1}^{N^*} \bar{d}_k \geq \bar{S}_1^{N^{**}}. \tag{36}$$

From (32) to (36), we have

$$\sum_{k=1}^{N^{**}-1} \bar{d}_k > \bar{S}_1^{N^{**}},$$

which implies that $N^{**} \geq \bar{N}$. Therefore, N^{**} is a forecast horizon. \square

Table 1. The forecast horizon in days for the first infinite horizon optimal production level for a ratio of maximum to minimum daily demands of at most two.

r	v	u				
		1.2	1.4	1.6	1.8	2
0.2	0.2	4	6	8	10	12
0.2	0.1	6	10	14	18	22
0.2	0.05	10	18	26	34	42
0.1	0.2	4	6	8	10	12
0.1	0.1	6	10	14	18	22
0.1	0.05	10	18	26	34	42
0.05	0.2	4	6	8	10	12
0.05	0.1	6	10	14	18	22
0.05	0.05	10	18	26	34	42

COROLLARY 1. If $\beta = 1 + \theta$, then $N^{**} \leq 2 + \beta N^*$.

PROOF.

$$N^{**} = 2 + \lceil \theta N^* \rceil \leq 2 + (1 + \theta) N^*. \quad \square$$

From Corollary 1, we conclude that the effect of uncertainty in demand is to increase the horizon we need to forecast demand over by a constant plus a factor of at most one plus the ratio of maximum to minimum per-period demands.

Also note that the stopping rule of the Solution Algorithm of §4.2 is always met at any forecast horizon N . We conclude that the Solution Horizon Algorithm will terminate after at most N^{**} iterations. Table 1, derived from Table 1 in Smith and Zhang (1998), gives numerical values for N^{**} for a variety of parameters. In the table, we have assumed the ratio of maximum to minimum daily demands is at most two. The maximum unit production cost is the factor u times the unit production cost in period 1; i.e., $u = \bar{c}/c_1$ while minimum inventory cost is v times the first-period production cost, i.e., $v = \bar{h}/c_1$. The table provides the forecast horizon N^{**} for various inventory charges v per day and interest rates r per year for u , respectively, one and two. The results are shown in Table 1.

Because the Solution Horizon Algorithm terminates within a horizon of at most N^{**} , we conclude from Table 1 that it terminates very rapidly for these parameter values with an infinite horizon optimal production decision after only a few increments of the horizon N . Moreover, forecast horizons are evidently remarkably short for problems with significant inventory costs, sometimes just a few days in length.

Appendix

LEMMA 7. If $\pi_n^N \rightarrow \pi_n$ as $N \rightarrow \infty$ for all n when $\pi^N \in \Pi^N$, $N = 1, 2, \dots$, then $Z^N(\pi^N) \rightarrow Z(\pi)$ as $N \rightarrow \infty$.

PROOF. Let $N_k \geq k$ be large enough that $\pi_n^N = \pi_n$ for all $n \leq k$ when $N \geq N_k$. Such an N_k exists because π_n^N is

integer and uniformly bounded over N for all n . Moreover, $N_k \rightarrow \infty$ as $k \rightarrow \infty$. By the definition of $Z(\pi)$ and $Z^N(\pi^N)$, we have for all N ,

$$|Z(\pi) - Z^N(\pi^N)| = \left| \sum_{n=1}^{\infty} \alpha^{n-1} P^n(\pi) Z_n(\pi) - \sum_{n=1}^{N-1} \alpha^{n-1} P^n(\pi^N) Z_n(\pi^N) - \alpha^{N-1} P^N(\pi^N) \tilde{Z} \right|.$$

Now, for all $N \geq N_k$, because $\pi_n = \pi_n^N$ for all $n \leq k$, we have

$$P^n(\pi) = P^n(\pi^N) \quad \text{and} \quad Z_n(\pi) = Z_n(\pi^N) \quad \text{for all } n \leq k.$$

Thus, for all k and $N \geq N_k$,

$$\begin{aligned} |Z(\pi) - Z^N(\pi^N)| &= \left| \sum_{n=1}^{\infty} \alpha^{n-1} P^n(\pi) Z_n(\pi) - \sum_{n=1}^{N-1} \alpha^{n-1} P^n(\pi^N) Z_n(\pi^N) - \alpha^{N-1} P^N(\pi^N) \tilde{Z} \right| \\ &\leq \alpha^{k-1} \sum_{n=k}^{\infty} \alpha^{n-k} P^n(\pi) |Z_n(\pi)| \\ &\quad + \alpha^{k-1} \sum_{n=k}^{N-1} \alpha^{n-k} P^n(\pi^N) |Z_n(\pi^N)| \\ &\quad + \alpha^{N-1} P^N(\pi^N) |\tilde{Z}|. \end{aligned}$$

Hence, for all k ,

$$\begin{aligned} \lim_{N \rightarrow \infty} |Z(\pi) - Z^N(\pi^N)| &\leq \alpha^{k-1} \sum_{n=k}^{\infty} \alpha^{n-k} P^n(\pi) |Z_n(\pi)| \\ &\quad + \alpha^{k-1} \sum_{n=k}^{\infty} \alpha^{n-k} P^n(\pi^N) |Z_n(\pi^N)| \\ &\quad + \lim_{N \rightarrow \infty} \alpha^{N-1} P^N(\pi^N) |\tilde{Z}| \\ &= \alpha^{k-1} \sum_{n=k}^{\infty} \alpha^{n-k} P^n(\pi) |Z_n(\pi)| \\ &\quad + \alpha^{k-1} \sum_{n=k}^{\infty} \alpha^{n-k} P^n(\pi^N) |Z_n(\pi^N)| \end{aligned}$$

because $\lim_{N \rightarrow \infty} \alpha^{N-1} P^N(\pi^N) |\tilde{Z}| = 0$. Assumptions 1, 2, and 3 imply that expected costs and revenues per period are uniformly bounded. Therefore, $Z_n(\pi)$ and $Z_n(\pi^N)$ are uniformly bounded; i.e., there exist $M < \infty$ such that each component of $Z_n(\pi)$ and $Z_n(\pi^N)$ is bounded by M for all n . Hence, we have

$$\begin{aligned} \sum_{n=k}^{\infty} \alpha^{n-k} P^n(\pi) |Z_n(\pi)| + \sum_{n=k}^{\infty} \alpha^{n-k} P^n(\pi^N) |Z_n(\pi^N)| \\ \leq \sum_{n=k}^{\infty} \alpha^{n-k} (2M) \leq \frac{2M}{1-\alpha}. \end{aligned}$$

Thus, taking the limit as $k \rightarrow \infty$, we get

$$\lim_{N \rightarrow \infty} |Z(\pi) - Z^N(\pi^N)| \leq \lim_{k \rightarrow \infty} \alpha^{k-1} \frac{2M}{1-\alpha} = 0.$$

Hence, $\lim_{N \rightarrow \infty} |Z(\pi) - Z^N(\pi^N)| = 0$ or

$$\lim_{N \rightarrow \infty} Z^N(\pi^N) = Z(\pi). \quad \square$$

Acknowledgments

This work was partially supported by the National Science Foundation under Grants DMI-9713723, 982074, and 9900267.

References

- Alden, J., R. L. Smith. 1992. Rolling horizon procedures in nonhomogeneous Markov decision processes. *Oper. Res.* **40**(Suppl. 2) S183–S194.
- Bean, J., R. L. Smith. 1984. Conditions for the existence of planning horizons. *Math. Oper. Res.* **9**(3) 391–401.
- Denardo, E. V. 1982. *Dynamic Programming, Models and Applications*. Prentice-Hall, Englewood Cliffs, NJ.
- Federgruen, A., M. Tzur. 1996. Detection of minimal forecast horizons in dynamic programs with multiple indicators of the future. *Naval Res. Logist.* **43** 169–189.
- Garcia, S., R. L. Smith. 2000a. Solving nonstationary infinite horizon stochastic production planning problems. *Oper. Res. Lett.* **27** 135–141.
- Garcia, S., R. L. Smith. 2000b. Solving nonstationary infinite horizon dynamic optimization problems. *J. Math. Anal. Appl.* **244**(2) 304–317.
- Hernandez-Lerma, O., J. Lasserre. 1990. Error bounds for rolling horizon policies in discrete-time Markov control processes. *IEEE Trans. Automatic Control* **35**(10).
- Karlin, S. 1960a. Dynamic inventory policy with varying stochastic demand. *Management Sci.* **6** 231–258.
- Karlin, S. 1960b. Optimal policy for dynamic inventory process with stochastic demand. *SIAM* **8** 611–629.
- Kleindorfer, P., H. Kunreuther. 1978. Stochastic horizons for the aggregate planning problem. *Management Sci.* **25** 1020–1031.
- Lovejoy, W. S. 1992. Stopped myopic policies in some inventory models with generalized demand processes. *Management Sci.* **38** 688–707.
- Morton, T. E. 1978. The nonstationary infinite horizon inventory problem. *Management Sci.* **24** 1474–1482.
- Morton, T. E., D. W. Pentico. 1995. The finite horizon nonstationary stochastic inventory problem: Near-myopic bounds, heuristics, testing. *Management Sci.* **41** 334–343.
- Sethi, S., F. Cheng. 1997. Optimality of (s, S) policies in inventory models with Markovian demand. *Oper. Res.* **45** 931–939.
- Smith, R. L., R. Q. Zhang. 1998. Infinite horizon production planning in time varying systems with convex production and inventory costs. *Management Sci.* **44** 1313–1320.
- Sobel, M. J. 1981. Myopic solution of Markov decision processes and stochastic games. *Oper. Res.* **29** 995–1009.
- Topkis, D. M. 1998. *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.
- Veinott, A. F. 1965. Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Sci.* **12** 206–222.
- Zhang, R. Q. 1998. New results in infinite horizon production planning. Technical report, Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI.