

Reprinted with permission. No further reproduction is authorized without written permission from the INFORMS. This version of the document is not the version of record. Figures and/or pictures may be missing from this format of the document.

(s, S) Policies for a Dynamic Inventory Model with Stochastic Lead Times

RICHARD EHRHARDT

University of North Carolina, Greensboro, North Carolina

(Received August 1980; accepted October 1983)

This study analyzes a stochastic lead time inventory model under the assumptions that (a) replenishment orders do not cross in time and (b) the lead time distribution for a given order is independent of the number and sizes of outstanding orders. The study extends the existing literature on the finite-horizon version of the model and yields an intuitively appealing dynamic program that is nearly identical to one that would apply in a transformed model with all lead times fixed at zero. Hence, many results that have been derived for fixed lead time models generalize easily. Conditions for the optimality of myopic base-stock policies, and for the optimality of (s, S) policies are established for both finite and infinite planning horizons. The infinite-horizon model analysis is extended by adapting the fixed lead time results for the efficient computation of optimal and approximately optimal (s, S) policies.

WE CONSIDER a periodic review, single-item inventory system where unfilled demand is backlogged, demands during review periods are independent, and the lead time between placement and delivery of an order may vary randomly. We require the joint distribution of lead times to have the properties that (a) replenishment orders do not cross in time and (b) the marginal lead time distribution of each order is independent of the number and size of outstanding orders. These assumptions could be appropriate in practice when, for example, only a single supplier is used and when the stocking organization places orders that are small and infrequent from the supplier's point of view. Replenishment costs are comprised of a setup cost and a cost that is linear in the amount ordered; holding and shortage costs are incurred in each period as a function of period-end inventory. Our optimality criterion is minimization of the expected discounted cost incurred during a planning horizon which may be finite or infinite. Alternatively, when we consider an undiscounted, infinite-horizon model, our criterion is minimization of the expected cost per period.

Kaplan [1970] analyzed a finite-horizon model of this system. The principal results were (a) that optimal policies can be computed using a

Subject classification: 354 (s, S) policies for stochastic lead time, 362 stochastic inventory lead times.

dynamic program having only a scalar state variable, representing inventory on hand plus items on order before ordering and (b) sufficient conditions can be found for the optimality of base stock policies and (s, S) policies. Although the findings in Kaplan represent a breakthrough in the study of stochastic lead time systems, the results had two complicating features that are not present in fixed lead time models. First, the parameters of the dynamic program were not simply related to the marginal lead time distribution. Second, his results did not specify sufficient conditions for the optimality of myopic ordering policies.

In this paper we make two small adjustments to the model of Kaplan, allowing an intuitively appealing analogy with a zero lead-time model. Then we present conditions for the optimality of myopic base stock policies and establish conditions for the optimality of (s, S) policies. We also extend the model to encompass infinite planning horizons and show that optimal (s, S) policies exist under standard conditions on the cost functions. Finally, we present efficient algorithms for computing optimal and approximately optimal (s, S) policies in the infinite-horizon setting.

1. MODEL SPECIFICATION

We initially consider a finite planning horizon of N periods, numbered backward from the end of the horizon; that is, the final period is given number 1, and the initial period is given number N . Demands in successive periods are independent, but not necessarily identically distributed. Specifically, let the demand in period n be represented by the random variable D_n with mean μ_n and cumulative distribution function Φ_n . Also, let $\Phi_{i,j}$ for $j \geq i$ be the convolution of Φ_i, \dots, Φ_j . We assume complete backlogging of unsatisfied demand, and so permit negative inventory levels. Also, there are no losses from the system other than through satisfying demand.

Costs in different periods are related by the single period discount factor α . Let $c_n(z)$ be the cost of ordering z units in period n , with

$$c_n(z) = K_n H(z) + c_n z,$$

where

$$H(z) = \begin{cases} 0, & z = 0 \\ 1, & z > 0. \end{cases}$$

We assume that both the setup cost K_n and the linear portion of the replenishment cost are paid upon delivery of the order. This assumption does not entail a loss of generality, since payment of either portion at the time of ordering can be described via scaling K_n or c_n by the expected value of α^L , where L is the random lead time.

Let $C_n(x)$ represent the holding and shortage costs in period n , where x is the ending inventory level in that period. Also, we define the function

$$\begin{aligned} g_n(i, y) &\equiv EC_{n-i}[y - (D_n + \dots + D_{n-i})] \\ &= \int_0^\infty C_{n-i}(y - u)d\Phi_{n,n-i}(u). \end{aligned} \quad (1)$$

Expression 1 can be interpreted from the perspective of a model with a fixed lead time of i periods. When viewed in this way, $g_n(i, y)$ is the expected holding and shortage costs in period $n - i$ when y is the value of inventory on hand plus items on order after ordering in period n .

We specify replenishment lead times as identically distributed random variables that can take on values from zero up to a fixed maximum m . Let a given lead time be represented by the random variable L having the probability distribution $l_i = P\{L = i\}$, $i = 0, \dots, m$. The joint distribution of lead times is characterized by our assumptions (identical to those in Kaplan) that (a) replenishment orders do not cross in time and (b) the lead time of an order is independent of the number and size of outstanding orders.

In developing the lead-time model, we follow the elegant explanation given by Nahmias [1979, pp. 911-913], who used the same lead time assumptions in analyzing a lost-sales inventory system. Nahmias observes that the replenishment delivery process arises from a sequence of independent and identically distributed random variables $\{A_1, A_2, \dots, A_N\}$. Each of the random variables characterizes the delivery of orders in a given period. One can interpret A_n as the age of the youngest order arriving in period n . That is, the event $\{A_n = k\}$ means that in period n all orders that have been outstanding for at least k periods are delivered. It is evident that this mechanism for specifying the delivery process guarantees that orders do not cross in time. Of course, if in period n there are no orders that have been outstanding for k or more periods, then the event $\{A_n = k\}$ means that no delivery occurs in period n .

Let the probability distribution of A_n be specified by the masses $\{p_0, p_1, \dots, p_m\}$, where $p_j = P\{A_n = j\}$ for $j = 0, 1, 2, \dots, m$. This distribution can be related to the distribution of an arbitrary lead time by focusing on an order placed in period n and its lead time L . The event $\{L = i\}$ is equivalent to the event $\{A_n > 0, A_{n+1} > 1, \dots, A_{n+i-1} > i - 1, A_{n+i} \leq i\}$. Observing that $\{A_1, A_2, \dots, A_N\}$ are independent and identically distributed, we have

$$P\{L = i\} = P\{A_n > 0\}P\{A_{n+1} > 1\} \dots P\{A_{n+i-1} > i - 1\}P\{A_{n+i} \leq i\}.$$

Now, since

$$P\{A_n > k\} = 1 - \sum_{j=0}^k p_j,$$

the value of $l_i = P\{L = i\}$ is given by

$$l_i = \begin{cases} p_0, & i = 0 \\ (1 - p_0)(1 - p_0 - p_1) \cdots (1 - \sum_{j=0}^{i-1} p_j) \sum_{j=0}^i p_j, & i = 1, 2, \dots, m. \end{cases} \quad (2)$$

Expression 2 is similar to one on p. 495 of Kaplan, which contains a minor error. The substantive results of Kaplan are entirely correct, though, since they are stated exclusively in terms of the p_j 's and do not depend on the relationship between the l_i 's and p_j 's. We will show, however, that (2) simplifies the interpretation of the dynamic program in Kaplan.

The final aspect of model specification concerns the costs which must be included in computing optimal policies. We include all costs that are incurred during periods N through 1, plus those that occur in the following m periods due to orders placed during the planning horizon. A terminal reward (or salvage value) is also applied to the inventory level at the end of the horizon. This procedure differs from Kaplan, who considers costs incurred only in periods N through 1, and sets the terminal reward arbitrarily at zero for all terminal states. The change in cost accounting allows us to derive conditions for the optimality of myopic base stock policies.

2. FINITE PLANNING HORIZONS

The central finding in Kaplan is that all policy-dependent costs can be included in a dynamic program that has inventory on hand plus items on order as its only state variable. Let $h_n(x)$ be the minimum expected discounted cost when x is the inventory on hand plus items on order immediately before ordering in period n . Theorem 4.1 of Kaplan gives a recursion for $h_n(x)$. The recursion remains valid if the model parameters vary from period to period, as Kaplan states on p. 493. It is also true that the recursion holds if we relax Kaplan's assumption of continuously distributed demand. With these adjustments, we have

$$\left. \begin{aligned} h_n(x) &= \min_{y \geq x} \left\{ \sum_{i=0}^m \alpha^i l_i [K_n H(y - x) + (y - x)c_n + g_n(i, y)] \right. \\ &\quad \left. + \alpha E h_{n-1}(y - D_n) \right\} \\ h_0(x) &= -c_0 x \sum_{i=0}^m \alpha^i l_i. \end{aligned} \right\} \quad (3)$$

Notice in (3) that l_i appears in place of Kaplan's q_{ii} , which was specified as a function of the p_j 's and equal to our Expression 2. Recursion (3) differs from Kaplan's also in that our cost accounting includes additional terms at the end of the horizon, as described above at the end of Section 1. The terminal reward function $h_0(\cdot)$ is derived in a manner similar to

the fixed lead time case of Veinott [1966b]. Cost accounting continues for one lead time beyond the last decision period, since these periods are influenced by the last decision. At that time we postulate that all the remaining inventory (or backlog) is salvaged (or purchased) at unit price c_0 . If the last order has lead time i , then the salvage value (in period $-i$) is given by

$$v_i(x) = \begin{cases} c_0x, & i = 0 \\ c_0(x - \sum_{j=0}^{i-1} D_j), & i = 1, \dots, m. \end{cases}$$

Hence, the expected discounted cost of the salvage operation is

$$-\sum_{i=0}^m l_i \alpha^i v_i(x) = -c_0x \sum_{i=0}^m \alpha^i l_i + c_0 \sum_{i=1}^m l_i \alpha^i \sum_{j=0}^{i-1} D_j.$$

The terminal reward function $h_0(\cdot)$ is taken as the first term on the right hand side, since the second term is independent of the replenishment decisions and cannot influence the choice of an optimal policy.

Now let

$$\beta \equiv E\alpha^L = \sum_{i=0}^m \alpha^i l_i,$$

$$k_n(y) \equiv \sum_{i=0}^m \alpha^i l_i g_n(i, y)/\beta,$$

and

$$F_n(x) \equiv h_n(x)/\beta.$$

Recursion (3) can be rewritten in terms of this notation as

$$\left. \begin{aligned} F_n(x) &= \min_{y \geq x} \{K_n H(y - x) + (y - x)c_n + k_n(y) \\ &\quad + \alpha E F_{n-1}(y - D_n)\} \\ F_0(x) &= -c_0x. \end{aligned} \right\} \quad (4)$$

Notice that (4) is of the same form as a recursion for a zero lead-time system with $k_n(\cdot)$ representing the single-period expected holding and shortage costs.

Following the approach of Veinott [1966b], we establish sufficient conditions for the optimality of myopic base stock policies. Consider the zero lead-time analogy for recursion (4), and let $F_n(x_n | Y)$ be the expected discounted cost in periods n through 0 when following a particular ordering policy Y and x_n is the starting inventory in period j . Also let $\{y_i, i = 1, \dots, N\}$ be the sequence of inventory on hand after ordering and before demand. Then one can show that

$$F_n(x_n | Y) = \sum_{i=1}^n \alpha^{n-i} E[K_i H(y_i - x_i) + G_i(y_i)] \\ + [\sum_{i=0}^{n-1} \alpha^{n-i} c_i \mu_{i+1} - c_n x_n]$$

where $G_n(y) = (c_n - \alpha c_{n-1})y + k_n(y)$. The functions $G_n(\cdot)$ are composites of expected holding and shortage costs and the linear purchase costs.

They can be interpreted (Veinott and Wagner [1965]) as the conditional expected holding and shortage cost functions of an equivalent model with unit purchase costs c_n set equal to zero. We shall use this interpretation, and hereafter refer to $G_n(\cdot)$ as a conditional expected holding and shortage cost function.

Consider the case of $K_n = 0$ for all n . It follows that if $-G_n(y)$ is unimodal with a minimum at \bar{y}_n , and if $\bar{y}_n - D_n \leq \bar{y}_{n-1}$ for $n = 2, \dots, N$ with probability one, then it is optimal to order $\max(\bar{y}_n - x_n, 0)$ in period n . For alternative conditions that ensure the optimality of myopic base stock policies, see Veinott [1966b].

An interesting parallel to the fixed lead time problem arises in the solution for the base-stock values $\{\bar{y}_i, i = 1, \dots, N\}$ when demand has a density and the single-period holding and shortage costs are given by

$$L_n(x) = h \max(x, 0) + p \max(-x, 0), \quad n = -m + 1, \dots, N. \quad (5)$$

Then one can show that \bar{y}_n is a solution to

$$\theta_n(\bar{y}_n) = (p - c_n + \alpha c_{n-1}) / (h + p),$$

where

$$\theta_n(y) = \sum_{i=0}^m \alpha^i l_i \Phi_{n,n-i}(y) / \beta.$$

Notice that the functions $\theta_n(y)$ are linear combinations of convoluted demand distributions and are legitimate distribution functions in their own right.

When the model does not possess an optimal myopic base-stock policy, we consider the function

$$\begin{aligned} f_n(x_n | Y) &= F_n(x_n | Y) - [\sum_{i=0}^{n-1} \alpha^{n-i} c_i \mu_{i+1} - c_n x_n] \\ &= \sum_{i=1}^n \alpha^{n-i} E[K_i H(y_i - x_i) + G_i(y_i)]. \end{aligned}$$

Notice that all policy-dependent costs are included in $f_n(x_n | Y)$. Therefore, an optimal policy can be found by computing

$$f_n(x) = \min_Y f_n(x | Y)$$

using the dynamic programming recursion

$$\left. \begin{aligned} f_n(x) &= \min_{y \geq x} \{K_n H(y - x) + G_n(y) + \alpha E f_{n-1}(y - D_n)\}, \\ & \qquad \qquad \qquad n = 1, \dots, N \end{aligned} \right\} \quad (6)$$

$$f_0(x) = 0.$$

Expression 6 is easily recognized as a standard form in inventory theory. Therefore, conclusions about the structure of optimal policies are immediate. For example, if $K_n = 0$ for all n , one can show (Veinott

[1966b]) that a base-stock policy is optimal when $G_n(y)$ is convex for all n . The base stock levels are given by the values $\{y_n^*, n = 1, \dots, N\}$ that minimize the expression in braces on the right-hand side of (6).

For models having $K_n > 0$ for at least one value of n , there are several sets of conditions that establish the optimality of (s, S) policies. The function $G_n(\cdot)$ is typically required to be convex or unimodal, and additional conditions are imposed on the variability of parameters over time. See, for example, Denardo [1982], Schäl [1976] and Veinott [1966a].

3. INFINITE PLANNING HORIZONS

We consider an infinite horizon version of our model in which all data are stationary. Our notation is simplified in this setting by suppressing subscripts that denote period numbers whenever the quantity of interest does not vary with time. Hence, recursion (6) becomes

$$\left. \begin{aligned} f_n(x) &= \min_{y \geq x} \{KH(y-x) + G(y) + \alpha E f_{n-1}(y-D)\}, \quad n \geq 1 \\ f_0(x) &= 0, \end{aligned} \right\} \quad (7)$$

where

$$G(y) = (1 - \alpha)cy + k(y),$$

and

$$k(y) = \sum_{i=0}^m \alpha^i l_i g(i, y) / \beta.$$

Recursion (7) is just like one for a fixed lead-time model, with $G(y)$ representing the conditional expected holding and shortage costs. Hence, we know that if $G(y)$ is convex, a stationary (s, S) policy is optimal in (7), as n approaches infinity. The conclusion is supported by the argument in Iglehart [1963], which also establishes the existence of $f(x) = \lim_{n \rightarrow \infty} f_n(x)$. Also from Iglehart, we know that $f(x)$ satisfies the functional equation

$$f(x) = \min_{y \geq x} \{KH(y-x) + G(y) + \alpha E f(y-D)\}. \quad (8)$$

The only difference between (8) and a fixed lead-time model is in the function $G(\cdot)$. In fact, $G(\cdot)$ can be expressed in the same form as the conditional expected holding and shortage cost function of a fixed lead-time model with a transformed demand distribution. We have

$$\begin{aligned} G(y) &= (1 - \alpha)cy + k(y) \\ &= (1 - \alpha)cy + \sum_{i=0}^m \alpha^i l_i g(i, y) / \beta \\ &= (1 - \alpha)cy + \sum_{i=0}^m \alpha^i l_i \int_0^\infty L(y-u) d\Phi^{*(i+1)}(u) / \beta, \end{aligned}$$

where Φ^{*j} is the j -fold convolution of the demand distribution Φ . Hence, $G(\cdot)$ can be expressed in the form

$$G(y) = (1 - \alpha)cy + \int_0^\infty L(y - u)d\Psi_\alpha(u), \quad (9)$$

where

$$\Psi_\alpha(x) = \sum_{i=0}^n \alpha^i l_i \Phi^{*(i+1)}(x)/\beta. \quad (10)$$

Notice that the function $\Psi_\alpha(\cdot)$ has all the properties of a distribution function. Therefore, we call $\Psi_\alpha(\cdot)$ the discounted lead time demand distribution.

When demand is discrete and $G(\cdot)$ is convex, expressions (9) and (10) allow a simple adaptation of the efficient Veinott-Wagner procedure for computing optimal (s, S) policies. Veinott and Wagner assume a fixed lead time λ , and compute the conditional holding and shortage cost function using the convoluted lead time demand distribution $\Phi^{*(\lambda+1)}$. To adapt Veinott and Wagner for our stochastic lead-time model, we merely substitute Ψ_α in place of $\Phi^{*(\lambda+1)}$ in any computation related to $G(\cdot)$. Specifically, the important expressions in Veinott and Wagner that require modification are (20), (21), (22), (23), (26), (27), and the two unnumbered expressions immediately preceding (21).

We have performed computations using the procedure described above. We summarize the results below, in Section 4.

We conclude this section with a discussion of approximately optimal (s, S) policies for the infinite horizon model. We have just shown how to compute optimal (s, S) policies by modifying a fixed lead-time procedure. Basically, the same kind of modification can be used to compute approximately optimal (s, S) policies as well.

For example, consider the common assumptions of $\alpha = 1$ and linear holding and shortage costs (as given by (5)). Roberts [1962] analyzed this model using asymptotic renewal theory to characterize the limiting behavior of an optimal policy (s^*, S^*) as the parameters K and p grow large. He obtained the following expressions for optimal policy parameters s^* and $D^* \equiv S^* - s^*$, as D^* becomes large:

$$D^* = \sqrt{2K\mu/h} + o(D^*),$$

$$\int_{s^*}^\infty (u - s^*)d\Phi^{*(\lambda+1)}(u) = D^*/(1 + p/h) + o(D^*),$$

where λ is the fixed lead-time, μ is the single period demand mean, and $o(D^*)/D^*$ converges to zero as D^* becomes infinite. Ehrhardt [1979] used these expressions to construct an approximately optimal policy (the Power Approximation) that is easy to compute and requires for demand

information only the mean and variance of demand. Specifically, the Power Approximation requires the single-period demand mean and variance, μ and σ^2 , as well as the mean and variance of $\Phi^{*(\lambda+1)}$, μ_λ and σ_λ^2 .

We suggest modifying the Power Approximation for our stochastic lead-time model by replacing μ_λ and σ_λ^2 with the mean and variance of Ψ_α , μ_l , and σ_l^2 . Specifically, expressions (13)–(16) in Ehrhardt require this change. In computing μ_l and σ_l^2 , we note that $\alpha = 1$ implies that $\beta = 1$ also. Therefore,

$$\mu_l = \int_0^\infty u d\Psi_1(u) = \sum_{i=0}^m l_i \int_0^\infty u d\Phi^{*(i+1)}(u) \quad (11)$$

$$\mu_l = \sum_{i=0}^m l_i(i+1)\mu = (EL+1)\mu,$$

$$\text{and } \sigma_l^2 = \int_0^\infty u^2 d\Psi_1(u) - \mu_l^2 = (EL+1)\sigma^2 + \mu^2 \text{Var}(L). \quad (12)$$

Notice that μ_l and σ_l^2 are merely the mean and variance of demand during

TABLE I
LEAD-TIME DISTRIBUTIONS

Probability Mass $l_i, i = 0, \dots, 4$					EL	Var L	γ
$i: 0$	1	2	3	4			
0	0	1	0	0	2	0	0
0	0.250	0.500	0.250	0	2	0.50	0.354
0.0667	0.2333	0.4000	0.2333	0.0667	2	1.00	0.500
0.2	0.2	0.2	0.2	0.2	2	2.00	0.707

$(L+1)$ periods. The use of (11) and (12) in place of μ_λ and σ_λ^2 is a familiar heuristic approach for modifying a fixed lead-time policy. Until now, however, this approach has not been theoretically justified for periodic review systems.

We assess the effectiveness of the modified Power Approximation in Section 4, below, which compares it with optimal policies for a variety of parameter settings.

4. NUMERICAL RESULTS

We have performed computations using the procedures described above for infinite horizon problems. In this section we consider a set of 12 inventory items under a variety of assumptions about the distribution of lead time. First, we show how optimal expected costs vary with the variance of the lead-time distribution. Then we compare the performance of optimal policies with that of the Power Approximation as modified by (11) and (12).

TABLE II
OPTIMAL TOTAL COST PER PERIOD OF 12 ITEMS FOR 4 DIFFERENT LEAD-TIME DISTRIBUTIONS

Lead-Time		Total Aggregate Cost	Costs Aggregated by Parameter Value						
			Penalty cost		Setup cost		Mean demand		
Mean	Variance		4	9	32	64	2	4	8
2	0	280	129	150	124	156	64	90	126
2	½	293	135	159	131	162	65	93	135
2	1	306	140	166	137	168	66	96	143
2	2	327	149	178	149	179	69	102	156

Consider a system of 12 inventory items, each having a negative binomial demand distribution with a variance-to-mean ratio of 3. Mean demand μ has three values, 2, 4, and 8. Each item has linear holding and shortage costs as given by (5). Since the total cost function is linear in the parameters h , p , and K , the value of the unit holding cost is a redundant parameter which is set at unity. The unit shortage costs are 4 and 9, and the setup cost values are 32 and 64. The computations considered all combinations of these parameter settings, yielding 12 items.

We consider the four lead-time distributions displayed in Table I. Each is a symmetrical triangular distribution over the range $[0, 4]$, with a mean value of 2. The variance of lead time ranges from a minimum of 0 for the deterministic case to a maximum of 2 for the uniform distribution. We also list the coefficient of variation γ of each lead-time distribution, which is defined as the ratio of the standard deviation to the mean.

Table II shows optimal total cost per period as a function of parameter values for each of the four lead-time distributions given in Table I. Notice that the total aggregate cost of the 12 items increases monotonically with lead-time variance. The largest lead-time variance yields an optimal total cost of 327 for the 12 items, 17% higher than the deterministic lead-time

TABLE III
PERCENTAGES ABOVE OPTIMAL TOTAL COST PER PERIOD FOR 12-ITEM SYSTEMS UNDER APPROXIMATELY OPTIMAL CONTROL

Lead-Time		Total Aggregate Cost	Costs Aggregated by Parameter Value						
			Penalty cost		Setup cost		Mean demand		
Mean	Variance		4	9	32	64	2	4	8
2	0	0.1	0.2	0.1	0.1	0.3	0.2	0.2	0.1
2	½	0.2	0.2	0.1	0.1	0.2	0.3	0.1	0.2
2	1	0.2	0.3	0.1	0.2	0.1	0.2	0.2	0.1
2	2	0.3	0.3	0.2	0.4	0.2	0.0	0.3	0.4

cost of 280. When costs are aggregated by parameter values, we see that the larger lead-time variances produce slightly larger cost increases for items with a penalty cost p of 9, than for those with p equal to 4. Items with a setup cost K of 32 as opposed to those with K equal to 64 also exhibit slightly larger cost increases. The bulk of the cost increase, however, can be attributed to items with the largest value of mean demand. Notice that items with μ equal to 8 shows a 24% increase in total cost from 126 for the deterministic lead-time system to 156 for the high lead-time variance system. The corresponding percentage increase for items with μ equal to 2 is merely 8%. This fact is not surprising, since we have held the demand variance-to-mean ratio constant. Therefore, items with the largest mean demand also have the largest variance of demand, yielding especially large values of σ_t^2 in (12).

Table III lists percentage increases in expected total cost per period when Power Approximation policies are compared with optimal policies. The Power Approximation yields costs within a few tenths of a percent of optimal for all parameter settings. This level of performance is comparable with the data in Ehrhardt, who considered only deterministic lead times.

ACKNOWLEDGMENTS

The author is pleased to acknowledge the helpful comments and suggestions of his colleagues, Professors Alastair MacCormick and Harvey M. Wagner, and of an anonymous referee and the area editor, Daniel Heyman. He also expresses his appreciation to James Lent who performed the computer calculations. This research was supported by the Office of Naval Research under Contract N00014-78-C-0467.

REFERENCES

- DENARDO, E.V. 1982. *Dynamic Programming: Models and Applications*. Prentice-Hall, Englewood Cliffs, N.J.
- EHRHARDT, R. 1979. The Power Approximation for Computing (s, S) Inventory Policies. *Mgmt. Sci.* **25**, 777-786.
- IGLEHART, D. L. 1963. Optimality of (s, S) Policies in the Infinite Horizon Dynamic Inventory Problem. *Mgmt. Sci.* **9**, 259-267.
- KAPLAN, R. S. 1970. A Dynamic Inventory Model with Stochastic Lead Times. *Mgmt. Sci.* **16**, 491-507.
- NAHMIAS, S. 1979. Simple Approximations for a Variety of Dynamic Lead-time Lost-sales Inventory Models. *Opns. Res.* **27**, 904-924.
- ROBERTS, D. M. 1962. Approximations to Optimal Policies in a Dynamic Inventory Model. In *Studies in Applied Probability and Management Science*, K. Arrow, S. Karlin and H. Scarf (eds.). Stanford University Press, Stanford, Calif.

- SCHÄL, M. 1976. On the Optimality of (s, S) -Policies in Dynamic Inventory Models with Finite Horizon. *SIAM J. Appl. Math.* **30**, 528–537.
- VEINOTT, A. F., JR. 1966a. On the Optimality of (s, S) Inventory Policies: New Conditions and a New Proof. *SIAM J. Appl. Math.* **14**, 1067–1083.
- VEINOTT, A. F., JR. 1966b. The Status of Mathematical Inventory Theory. *Mgmt. Sci.* **12**, 745–777.
- VEINOTT, A. F., JR., AND H. M. WAGNER. 1965. Computing Optimal (s, S) Inventory Policies. *Mgmt. Sci.* **11**, 525–552.