Computational Issues in an Infinite-Horizon, Multiechelon Inventory Model

AWI FEDERGRUEN and PAUL ZIPKIN

Columbia University, New York, New York

(Received March 1982; accepted May 1983)

Clark and Scarf [1960] characterize optimal policies in a two-echelon, two-location inventory model. We extend their result to the infinite-horizon case (for both discounted and average costs). The computations required are far easier than for the finite horizon problem. Further simplification is achieved for normal demands. We also consider the more interesting case of multiple locations at the lower echelon. We show that, under certain conditions, this problem can be closely approximated by a model with one such location. A rather simple computation thus yields both a near-optimal policy and a good approximation of the cost of the system.

THIS PAPER treats a two-echelon inventory system. The higher echelon is a single location referred to as the *depot*, which places *orders* for exogenous supply of a single commodity. The lower echelon also consists of one point, called the *retail outlet*, which is supplied by *shipments* from the depot, and at which random demands for the item occur. (The case of several retail outlets is discussed below.) Stocks are reviewed and decisions made periodically. Instantaneous, perfect information about inventory at both levels is assumed. Orders and/or shipments may each require a (fixed) leadtime before reaching their respective destinations.

Unfilled demand is backordered at the outlet, incurring a linear penalty cost. Linear holding costs are assessed on "echelon inventories." That is, a cost is incurred for all stock in the system, whether at the depot, at the outlet, or in transit between them; an additional holding cost is charged on inventory at the retail outlet, it being more expensive to keep stock there. The cost of a shipment is assumed linear, while an order incurs a fixed cost as well as a proportional cost. All costs and demands are assumed stationary.

Clark and Scarf [1960], assuming a finite planning horizon, show that an optimal policy for this system can, in principle, be computed by decomposing the problem into two separate single-location problems:

Subject classification: 362 stochastic models, multiechelon; 570 multiechelon inventory systems.

Operations Research
Vol. 32, No. 4, July–August 1984

0030-364X/84/3204-0818 $01.25

The problem for the retail outlet includes only its "own" costs, ignoring all others; a critical-number policy solves this problem. The optimal policy and expected cost function for each period are then used to define a convex "induced penalty cost" function for each period. This function is added to the depot's holding costs and the ordering costs to form the second problem. An (s, S) policy solves this problem and constitutes an optimal order policy for the system as a whole. An optimal shipment policy results from a slight modification of the outlet problem's policy: In each period ship up to the critical number, if the depot has that much stock; if not, ship as much as possible.

While this characterization greatly simplifies the original problem, actual computation of an optimal policy still encounters substantial obstacles. First, two sets of recursive functional equations must be solved numerically. Second, each evaluation of the induced penalty cost function itself entails a numerical integration over the optimal cost function for the outlet's problem; indeed, with an order leadtime of several periods, the computation requires a double integration (see Section 4 below).

Regarding the first issue, we show that the qualitative result of Clark and Scarf extends to the infinite horizon case, first under the criterion of discounted cost (Section 2), and then for long-term average cost (Section 3). The resulting two single-location problems are much easier to solve than their respective finite-horizon versions (as is often the case in dynamic problems).

Turning to the second issue, we show that in the infinite horizon problems the induced penalty cost functions are stationary, do not involve optimal cost functions, and require at most one numerical integration (none for zero order leadtime). The case of normally distributed demands (Section 4) requires *no* explicit numerical integration; the penalty cost can be expressed in terms of functions available in standard packages (specifically, univariate and bivariate normal cumulative distribution functions).

Section 5 presents a method to approximate a problem with several retail outlets by a single-outlet model. Assuming demand at each outlet is normal, with identical coefficients of variation, costs and leadtimes, we argue (partly analytically and partly based on empirical results of Federgruen and Zipkin [1984a, 1984c]) that the approximation should be very accurate. The methods of prior sections, therefore, permit a rather simple computation to find both a near-optimal policy and a good approximation of the cost of the system. The latter can be used in design studies, such as those described by Gross et al. [1981].

Although the model assumes completely centralized control, our proposed policy can be interpreted as prescribing a largely decentralized ("pull") system, where each outlet "orders" up to its own critical number.

Though it may not be obvious, our results can be viewed as justifying and considerably streamlining the original approach of Clark [1958] to this problem (reported in Clark and Scarf, and Gross et al.). That approach assumes that inventories at the outlets are always "in balance," a concept central to the development in Section 5. The analysis begins (where ours ends) with an independent optimal critical-number policy for each outlet. An elaborate parametric optimization is required, however, to construct an approximate cost function for the depot in each period, so the approach becomes awkward for more than a few outlets and time periods. Our method, by contrast, is virtually no harder for many outlets than for one. Incidentally, our results suggest that the balance assumption, and hence the Clark approach, is inappropriate when coefficients of variation are seriously unequal. Section 5 discusses extensions to handle this case.

Section 5 draws heavily on recent work by Eppen and Schrage [1981] and Federgruen and Zipkin [1984a, 1984c], who develop effective, computationally simple methods for the problem where the depot cannot hold stock. Zipkin [1984] provides a formal treatment of inventory balance. Results for the general case under somewhat different assumptions include those of Ignall and Veinott [1969] and Veinott [1971]. Schmidt and Nahmias [1981] have recently studied a problem with two components assembled to form a final product using the framework of Clark and Scarf.

In sum, the results here permit computation of optimal or near-optimal policies for multiechelon systems of considerable complexity and realism, using only readily available and not-too-taxing software.

1. NOTATION AND PRELIMINARIES

The sequence of events represented by the model studied in this paper differs somewhat from that of Clark and Scarf, in accord with what has become common usage. Holding as well as penalty costs are incurred *after* demand in each period. Also, orders and shipments arrive, following their respective leadtimes, at the beginning of a period, that is, *after* costs are assessed in the prior period, and *before* the current decisions. (This change does not affect the analysis in any essential way, and slightly simplifies computations.)

The cost data of the problem are

K = fixed cost to place an order.
c^d = cost rate for proportional order costs.
c^r = cost rate for proportional shipment costs.
h^d = holding cost rate for total system inventory.

h^r = additional holding cost rate for inventory at the retail outlet.
p^r = penalty cost rate for backorders at the retail outlet.

We assume that these cost factors are positive, and that they are related in certain ways, depending on other factors, that preclude it being optimal never to order.

Other parameters are

α = discount rate, $0 \leq \alpha \leq 1$.
l = leadtime for shipments, a nonnegative integer.
L = leadtime for orders, a nonnegative integer.
u = one-period demand, a nonnegative random variable.
$u^{(i)} = i$ − period demand, $i = 1, 2, \cdots$.
$\mu = E(u) < \infty$, $\mu^{(i)} = i\,\mu = E(u^{(i)})$.

For convenience we shall assume u is continuous, though this assumption is not strictly necessary. Demands in different periods are assumed independent.

The evolution of the system is described by several state and action variables. (Our formulation already includes the standard transformation, as in Clark and Scarf, that eliminates state variables representing individual outstanding shipments.)

The variables are

x^r = inventory at the retail outlet, plus shipments currently in transit.
v^d = echelon inventory at the depot
= inventory at the depot, plus x^r.
z = shipment from the depot to the retail outlet.
y = order size.
y^i = outstanding order placed i periods ago, $i = 1, \cdots, L$.
$\tilde{y} = (y^1, \cdots, y^L)$.

The index n will denote the number of periods *remaining* until the end of planning horizon, that is, we number time periods backward. This index is suppressed where possible, as above; when necessary we shall write $x_n{}^r$, y_n, and so forth. The actions are constrained by the inequalities:

$$y \geq 0,\ z \geq 0,\ x^r + z \leq v^d + y^L.$$

The following equations specify the dynamics of the system:

$$x^r_{n-1} = x_n{}^r + z_n - u_n,\ v^d_{n-1} = v_n{}^d + y_n{}^L - u_n,$$

$$\tilde{y}_{n-1} = (y_n, y_n{}^1, \cdots, y_n^{L-1}).$$

To formulate system-wide inventory costs, we (temporarily) need more detailed descriptors of the system: Let w^d = inventory at the depot, w^r = inventory at the retail outlet, and s^i = outstanding shipment placed i

periods ago for $i = 1, \cdots l$. At the end of the current period, system-wide inventory equals

$$w^d + y^L + \sum_{i=1}^{l-1} s^i + [w^r + s^l - u]^+$$
$$= v^d + y^L + \{-(w^r + s^l) + [w^r + s^l - u]^+\}.$$

Our current decisions affect the expression in brackets only l periods later. We shall thus treat the corresponding expected costs through the next $l - 1$ periods as constants, and in the current period count costs occurring l periods later. Under this accounting scheme, the relevant one-period expected holding costs for system-wide inventory are

$$h^d E\{v^d + y^L - \alpha^l(x^r + z - u^{(l)}) + \alpha^l[x^r + z - u^{(l+1)}]^+\}.$$

Similarly, we count now expected penalty and holding costs at the outlet occurring l periods later,

$$\alpha^l E(h^r[x^r + z - u^{(l+1)}]^+ + p^r[u^{(l+1)} - x^r - z]^+),$$

treating prior such costs as constant.

Let A^r denote all the constant costs mentioned above. The one-period holding and penalty costs may thus be represented by $D(v^d + y^L) + R(x^r + z)$, where

$$D(v) = h^d v, \quad \text{and}$$
$$R(x) = \alpha^l\{-h^d(x - \mu^{(l)}) + p^r E[u^{(l+1)} - x]^+ + (h^d + h^r)E[x - u^{(l+1)}]^+\}.$$

Finally, let

$$c^d(y) = \text{order cost function}$$
$$= 0 \quad \text{if} \quad y = 0, K + c^d y \quad \text{if} \quad y > 0.$$

We may now state the finite-horizon dynamic program:

$\hat{g}_n(\tilde{y}, v^d, x^r)$ = minimum total discounted expected costs with n periods remaining, if the system begins in state $(\tilde{y}, v^d, x^r)$, excluding costs represented by A^r.

Then,

$$\hat{g}_0 = 0$$

$$\hat{g}_n(\hat{y}, v^d, x^r) = \min_{y,z}\{c^d(y) + D(v^d + y^L) + c^r z + R(x^r + z) + \alpha E\hat{g}_{n-1}[(y, y^1, \cdots, y^{L-1}), v^d + y^L - u, x^r + z - u]: y \geq 0, z \geq 0, x^r + z \leq v^d + y^L\}, \qquad n \geq 1. \tag{1}$$

(Because of the transformation, the effective planning horizon is shifted from $n = 0$ to $n = l$, so (1) represents a different numbering of time periods.)

The main result of Clark and Scarf is that program (1) can be decomposed into a pair of simpler programs. The first of these programs involves the outlet alone:

$$g_0{}^r = 0,$$

$$g_n{}^r(x^r) = \min_{z \geq 0}\{c^r z + R(x^r + z) + \alpha E g^r_{n-1}(x^r + z - u)\}, \qquad n \geq 1. \tag{2}$$

Since R is convex and $c^r z$ is linear, a critical-number policy solves this problem. Let x^{r*}_n denote the critical number for period n. Now define the induced penalty cost functions:

$$\hat{P}_n(x) = \begin{cases} 0, & x \geq x^{r*}_n, \\ c^r(x - x^{r*}_n) + [R(x) - R(x^{r*}_n)] + \alpha E[g^r_{n-1}(x - u) \\ \quad - g^r_{n-1}(x^{r*}_n - u)], & x < x^{r*}_n. \end{cases}$$

(These functions are easily shown to be nonnegative and convex.)

The second dynamic program is the following:

$$\hat{g}_0{}^d = 0$$

$$\hat{g}_n{}^d(\tilde{y}, v^d) = \min_{y \geq 0}\{c^d(y) + D(v^d + y^L) + \hat{P}_n(v^d + y^L) + \alpha E \hat{g}^d_{n-1}[(y, y^1, \cdots, y^{L-1}), v^d + y^L - u]\}, \quad n \geq 1. \tag{3}$$

Because of the form of $c^d(\cdot)$ and the convexity of D and $\hat{P}_n$, an (s, S) policy solves this problem (cf., Scarf [1960]).

In particular, Clark and Scarf show

$$\hat{g}_n(\tilde{y}, v^d, x^r) = \hat{g}_n{}^d(\tilde{y}, v^d) + g_n{}^r(x^r), \tag{4}$$

and an optimal policy for the system consists of an optimal policy for (3) as the order policy, and a modification (as described in the Introduction) of the critical-number policy optimal for (2) as the shipment policy.

Our goal is to demonstrate a similar separation for both the discounted- and average-cost infinite-horizon problems. To state these problems, we employ a construction that goes back at least to Veinott [1965]. Let $\prod$ be the class of (infinite-horizon) measurable policies and $\tilde{u}$ any realization of the entire sequence of demands. Then, for a given starting state $(\tilde{y}, v^d, x^r)$, a policy $\pi \in \prod$ and $\tilde{u}$ determine the sequences of states and actions, and we may define

$$B_n(\tilde{y}, v^d, x^r \mid \pi, \tilde{u}) = \sum_{i=1}^n \alpha^{n-i}[c^d(y_i) + D(v_i^d + y_i^L) + c^r z_i + R(x_i^r + z_i)]$$

for all n, $(\tilde{y}, v^d, x^r)$, $\pi \in \prod$ and $\tilde{u}$, where $(\tilde{y}_n, v_n^d, x_n^r) = (\tilde{y}, v^d, x^r)$.

Assuming discounted costs ($\alpha < 1$), we let

$$B^\alpha(\tilde{y}, v^d, x^r \mid \pi, \tilde{u}) = \lim_{n\to\infty} B_n(\tilde{y}, v^d, x^r \mid \pi, \tilde{u}).$$

(The limit exists, though it may be $+\infty$, since $\{B_n\}$ is nondecreasing.) Then, $B^\alpha(\tilde{y}, v^d, x^r \mid \pi) \equiv EB^\alpha(\tilde{y}, v^d, x^r \mid \pi, \tilde{u})$ is the expected discounted cost under policy π. Set $B^\alpha(\tilde{y}, v^d, x^r) = \inf\{B^\alpha(\tilde{y}, v^d, x^r \mid \pi): \pi \in \prod\}$. An optimal policy (if it exists) achieves the infimum for every state. We refer to this problem as IH_α.

Under the average-expected-cost criterion (with $\alpha = 1$), let $B_n(\tilde{y}, v^d, x^r \mid \pi) = EB_n(\tilde{y}, v^d, x^r \mid \pi, \tilde{u})$. $B(y, v^d, x^r \mid \pi) = \limsup_{n\to\infty}(1/n)B_n(\tilde{y}, v^d, x^r \mid \pi)$ is taken as the average cost under π, and the problem is to find $B(\tilde{y}, v^d, x^r) = \inf\{B(\tilde{y}, v^d, x^r \mid \pi): \pi \in \prod\}$; again, an optimal policy achieves every infimum. Call this problem IH.

Consider now the analogous problems for the retail outlet alone. That is, let $\prod^r$ be the class of shipment policies (constrained only by $z \geq 0$), and $B_n^r(x^r \mid \pi^r, u) = \sum_{i=1}^n \alpha^{n-i}[c^r z_i + R(x_i^r + z_i)]$, where $x_n^r = x^r$, for all x^r, $\pi^r \in \prod^r$ and $\tilde{u}$. Using limits and expectations as above, define $B^{r\alpha}(x^r \mid \pi^r)$, $B^{r\alpha}(x^r)$, $B^r(x^r \mid \pi^r)$ and $B^r(x^r)$. Denote by IH_α^r and IH^r the discounted and undiscounted problems, respectively. The following are well-known facts about these problems (cf. e.g., Chapter 3 in Heyman and Sobel [1984]) that are true for both discounted and average costs unless stated otherwise:

(a) A stationary critical-number policy is optimal. Let x^{r*} be the critical number.
(b) x^{r*} is the global minimum of $(1 - \alpha)c^r x + R(x)$, and hence solves the equation $(1 - \alpha)c^r + R'(x) = 0$.
(c) The sequence $\{x_n^{r*}\}$ is nondecreasing and converges to x^{r*}.
(d) For all n and $x \leq x_n^{r*}$, $g_n^r(x) = c^r(x_n^{r*} - x) + g_n^r(x_n^{r*})$.
(e) For $\alpha < 1$, the sequence $\{g_n^r\}$ converges to a function g^r, and $g^r = B^{r\alpha}$. Also, $g^r(x) = c^r(x^{r*} - x) + g^r(x^{r*})$, $x \leq x^{r*}$.
(f) For $\alpha = 1$, let a^r denote the minimal average cost. Then $\lim_{n\to\infty}(1/n)g_n^r(x) = B^r(x) = a^r = c^r\mu + R(x^{r*})$, for all x.

(Other properties will be introduced as needed.)

We may now define the appropriate induced penalty cost for the infinite-horizon case:

$$P(x) = \begin{cases} 0, & x \geq x^{r*} \\ (1 - \alpha)c^r(x - x^{r*}) + [R(x) - R(x^{r*})], & x < x^{r*}. \end{cases}$$

(It is straightforward to show that P is nonnegative and convex. Observe that, in contrast to the $\hat{P}_n$, P does not involve optimal cost functions and requires no integration, except perhaps to evaluate R.) Define a program similar to (3), but with P replacing the $\hat{P}_n$:

$$g_0{}^d = 0,$$
$$g_n{}^d(\tilde{y}, v^d) = \min_{y \geq 0}\{c^d(y) + D(v^d + y^L) + P(v^d + y^L) + \alpha E g_{n-1}^d[(y, y^1, \cdots, y^{L-1}), v^d + y^L - u]\}, \quad n \geq 1, \tag{5}$$

and, by analogy with (4), define

$$g_n(\tilde{y}, v^d, x^r) = g_n{}^d(\tilde{y}, v^d) + g_n{}^r(x^r). \tag{6}$$

Finally, define the infinite-horizon problems for the depot $IH_\alpha{}^d$ and IH^d, with $\prod^d$ the class of order policies, setting

$$B_n{}^d(\tilde{y}, v^d \mid \pi^d, \tilde{u}) = \sum_{i=1}^{n} \alpha^{n-i}[c^d(y_i) + D(v_i^d + y_i^L) + P(v_i^d + y_i^L)]$$

for $\pi^d \in \prod^d$ and $(\tilde{y}_n, v_n{}^d) = (y, v^d)$, and defining $B^{d\alpha}(\tilde{y}, v^d \mid \pi^d)$, $B^{d\alpha}(\tilde{y}, v^d)$, $B^d(\tilde{y}, v^d \mid \pi^d)$ and $B^d(\tilde{y}, v^d)$ as above. For both $IH_\alpha{}^d$ (cf. Iglehart [1963b]) and IH^d (cf. Iglehart [1963a]), a stationary (s, S) policy is optimal.

Now let $\pi_\alpha{}^*$ (resp., π^*) denote the stationary policy that specifies orders according to the policy solving $IH_\alpha{}^d$ (resp. IH^d), and shipments according to the (modified) critical-number policy solving $IH_\alpha{}^r$ (resp. IH^r). As we shall see, $\pi_\alpha{}^*$ (resp. π^*) is optimal for IH_α (resp. IH).

To study these problems, we shall examine the limiting behavior of the $\hat{g}_n = \hat{g}_n{}^d + g_n{}^r$. The difficulty is that $\hat{g}_n{}^d$ arises from problem (3) which has nonstationary one-period costs, so we would like to replace $\hat{g}_n{}^d$ by $g_n{}^d$, hence $\hat{g}_n$ by g_n. As shown below, $\{\hat{P}_n\} \to P$; intuitively, as n becomes large, significant differences between the programs (3) and (5) recede infinitely far into the future. The next two sections provide a rigorous basis for this intuition. Indeed, for the average-cost case (Section 3) the difficulty essentially disappears. (A similar problem of convergent costs is studied by Federgruen and Schweitzer [1981] for general, finite-state dynamic programs, and in broad outline our arguments are similar to theirs. See also Evans [1981] for some related results.)

2. THE INFINITE HORIZON CASE: DISCOUNTED COSTS

Assume $\alpha < 1$. Iglehart [1963b] shows that the differences $(g_n{}^r - g^r)$ are bounded and converge uniformly to the zero function on the interval $(-\infty, x^{r*})$. That is, letting $\beta_n = \sup\{\,|g_n{}^r(x) - g^r(x)|: -\infty < x \leq x^{r*}\}$, we have $\beta_n < \infty$ and $\{\beta_n\} \to 0$.

LEMMA 1. *The differences* $(\hat{P}_n - P)$ *are bounded and converge uniformly to the zero function on the entire real line.*

Proof. We shall refer to several of the properties (a)–(f) in Section 1. First, for $x \geq x^{r*}$ and all n, $\hat{P}_n(x) = P(x) = 0$. Defining

$$\gamma_n = \sup\{|\hat{P}_n(x) - P(x)|: -\infty < x < \infty\},$$
$$\gamma_n^- = \sup\{|\hat{P}_n(x) - P(x)|: -\infty < x \leq x_n^{r*}\}, \tag{7}$$
$$\gamma_n^+ = \sup\{|\hat{P}_n(x) - P(x)|: x_n^{r*} \leq x \leq x^{r*}\},$$

we have $\gamma_n = \max\{\gamma_n^-, \gamma_n^+\}$. We wish to show that $\gamma_n < \infty$ and $\{\gamma_n\} \to 0$, and we shall demonstrate these properties for both the γ_n^- and the γ_n^+.

For $x \leq x^{r*}$, using $u \geq 0$ and (e), $E[g^r(x - u) - g^r(x^{r*} - u)] = c^r(x^{r*} - x)$. Thus, for $x \leq x_n^{r*}$, using (c),

$$\begin{aligned}
\hat{P}_n(x) - P(x) &= \{c^r(x - x_n^{r*}) + [R(x) - R(x_n^{r*})] \\
&\quad + \alpha E[g_{n-1}^r(x - u) - g_{n-1}^r(x_n^{r*} - u)]\} \\
&\quad - \{c^r(x - x^{r*}) + [R(x) - R(x^{r*})] \\
&\quad + \alpha E[g^r(x - u) - g^r(x^{r*} - u)]\} \\
&= -c^r(x_n^{r*} - x^{r*}) - [R(x_n^{r*}) - R(x^{r*})] \\
&\quad + \alpha E[g_{n-1}^r(x - u) - g^r(x - u)] \\
&\quad - \alpha E[g_{n-1}^r(x_n^{r*} - u) - g^r(x_n^{r*} - u)] \\
&\quad - \alpha E[g^r(x_n^{r*} - u) - g^r(x^{r*} - u)].
\end{aligned}$$

The third expectation here equals $-c^r(x_n^{r*} - x^{r*})$. Considering the first two expectations, we have

$$E\{|g_{n-1}^r(x - u) - g^r(x - u)|\} \leq \beta_{n-1},$$

and $$E\{|g_{n-1}^r(x_n^{r*} - u) - g^r(x_n^{r*} - u)|\} \leq \beta_{n-1},$$

so $$\gamma_n^- \leq (1 - \alpha)c^r(x^{r*} - x_n^{r*}) + [R(x_n^{r*}) - R(x^{r*})] + 2\alpha\beta_{n-1} < \infty,$$

and $\{\gamma_n^-\} \to 0$.

Clearly, each $\gamma_n^+ = P(x_n^{r*}) < \infty$, and $\{\gamma_n^+\} \to 0$.

LEMMA 2. *The differences* $(\hat{g}_n^d - g_n^d)$ *are bounded and converge uniformly to the zero function.*

Proof. Let $\hat{y}_n^* = \hat{y}_n^*(\tilde{y}, v^d)$ achieve the minimum in (3) and $y_n^* = y_n^*(\tilde{y}, v^d)$ in (5), and define $\delta_0 = 0$ and

$$\delta_n = \sup\{|\hat{g}_n^d(\tilde{y}, v^d) - g_n^d(\tilde{y}, v^d)|: \tilde{y} \geq 0, v^d \in \mathbb{R}\}, \quad n \geq 1.$$

Then, using Lemma 1 gives

$$\hat{g}_n{}^d(\tilde{y}, v^d) - g_n{}^d(\tilde{y}, v^d) \leq [\hat{P}_n(v^d + y^L) - P_n(v^d + y^L)]$$
$$+ \alpha E\{\hat{g}^d_{n-1}[y_n{}^*, y^1, \cdots, y^{L-1}), v^d + y^L - u]$$
$$- g^d_{n-1}[(y_n{}^*, y^1, \cdots, y^{L-1}), v^d + y^L - u]\} \leq \gamma_n + \alpha\delta_{n-1},$$
$$\hat{g}_n{}^d(\tilde{y}, v^d) - g_n{}^d(\tilde{y}, v^d) \geq [\hat{P}_n(v^d + y^L) - P_n(v^d + y^L)]$$
$$+ \alpha E\{\hat{g}^d_{n-1}[(\hat{y}_n{}^*, y^1, \cdots, y^{L-1}), v^d + y^L - u]$$
$$- g^d_{n-1}[(\hat{y}_n{}^*, y^1, \cdots, y^{L-1}), v^d + y^L - u]\} \geq -\gamma_n - \alpha\delta_{n-1},$$

so

$$|\hat{g}_n{}^d(\tilde{y}, v^d) - g_n{}^d(\tilde{y}, v^d)| \leq \gamma_n + \alpha\delta_{n-1}.$$

Letting $\delta_n' = \sum_{i=1}^{n} \alpha^{n-i}\gamma_i$ and using straightforward induction yields $\delta_n \leq \delta_n'$, for all n, and a direct analytical argument shows $\{\delta_n'\} \to 0$.

From Iglehart [1963b] we know $\{g_n{}^d\}$ converges to a function g^d. Defining $g = g^d + g^r$, we have the following result:

LEMMA 3. *The sequence* $\{\hat{g}_n\}$ *converges to* g.

Now consider the infinite-horizon functional equation analogous to (1):

$$f(\tilde{y}, v^d, x^r) = \inf_{y,z}\{c^d(y) + D(v^d + y^L) + c^r z + R(x^r + z)$$
$$+ \alpha E f[(y, y^1, \cdots, y^{L-1}), v^d + y^L - u, x^r + z - u]: \qquad (8)$$
$$y \geq 0, z \geq 0, x^r + z \leq v^d + y^L\}.$$

LEMMA 4. *The function g satisfies* (8), *and the infimum is achieved by the policy* $\pi_\alpha{}^*$.

Proof. Substitute $g = g^d + g^r$ into the right hand side of (8), and use the same argument as in the finite-horizon case.

We are now prepared for the main result:

THEOREM 1. *The policy* $\pi_\alpha{}^*$ *is optimal for problem* IH_α.

Proof. By Lemmas 3 and 4 and Proposition 9.16 in Bertsekas and Shreve [1978], $g = B^\alpha$ (the true infimal cost function). The result now follows from Proposition 9.12 in the same reference. (Note, these propositions require the one-period costs to be bounded below. This condition is easy to establish, provided $\alpha^l p^r \geq (1 - \alpha^l)h^d$; the latter condition *must* hold, otherwise it is optimal never to order.)

3. THE INFINITE HORIZON CASE: AVERAGE COSTS

With $\alpha = 1$, by standard arguments, the average shipping cost is $c^r\mu$ and the average proportional order cost is $c^d\mu$, under all interesting policies. Thus, we may assume $c^d = c^r = 0$ without loss of generality. A straightforward induction demonstrates that $c^r = 0$ implies $x_n^{r*} = x^{r*}$ for all $n \geq 1$, so $g_n^r(x) = nR(x^{r*})$ whenever $x \leq x^{r*}$. But this result implies that $\hat{P}_n = P$, so $\hat{g}_n^d = g_n^d$ and $\hat{g}_n = g_n$ for all $n \geq 1$. Thus, many of the difficulties encountered in the discounted case do not arise here.

We know (Iglehart [1963a]) that $\{g_n^d/n\}$ converges to $a^d = B^d(\tilde{y}, v^d)$, the true minimum average cost for the depot's problem. (See also property (f) in Section 1.)

Letting $a = a^d + a^r$, we have

$$\{\hat{g}_n(\tilde{y}, v^d, x^r)/n\} \to a, \quad \text{for all} \quad (\tilde{y}, v^d, x^r).$$

LEMMA 5. *The policy π^* has average cost a.*

Proof. In a finite number of periods under π^* we will have $x_i^r \leq x^{r*}$ with probability 1, so we may assume $x^r \leq x^{r*}$. Letting z_i^r denote the decision under the unconstrained critical-number policy, we have $x_i^r + z_i^r = x^{r*}$, $i = 1, \cdots, n$ for all n. Thus,

$$R(x_i^r + z_i) = \begin{cases} R(x^{r*}), & v_i^d + y_i^L \geq x^{r*} \\ R(v_i^d + y_i^L), & v_i^d + y_i^L < x^{r*} \end{cases}$$

$$= R(x^{r*}) + P(v_i^d + y_i^L).$$

Thus, for all n, $\tilde{u}$, and $(\tilde{y}, v^d, x^r)$

$$B_n(\tilde{y}\ v^d, x^r \mid \pi^*, \tilde{u}) = B_n^d(\tilde{y}, v^d \mid \pi^*, \tilde{u}) + nR(x^{r*}),$$

so

$$B(\tilde{y}, v^d, x^r \mid \pi^*) = a^d + a^r = a.$$

THEOREM 2. *The policy π^* is optimal for problem IH.*

Proof. We must show $B(\tilde{y}, v^d, x^r) = B(\tilde{y}, v^d, x^r \mid \pi^*) = a$. Clearly, $B(\tilde{y}, v^d, x^r) \leq a$, since π^* is a feasible policy. Now, fix $\epsilon > 0$. There is some policy π'' such that

$$B(\tilde{y}, v^d, x^r \mid \pi'') \leq B(\tilde{y}, v^d, x^r) + \epsilon.$$

By result III.1 on p. 86 of Dynkin and Yushkevich [1979], for $n \geq 1$ there is a (finite-horizon) Markovian policy π_n' such that

$$B_n(\tilde{y}, v^d, x^r \mid \pi_n') = B_n(\tilde{y}, v^d, x^r \mid \pi'').$$

But, by the definition of $\hat{g}_n$,

$$\hat{g}_n(\tilde{y}, v^d, x^r) \leq B_n(\tilde{y}, v^d, x^r \mid \pi_n') = B_n(\tilde{y}, v^d, x^r \mid \pi'').$$

Dividing by n and taking limits yields $a \leq B(\tilde{y}, v^d, x^r) + \epsilon$. Since this is true for all $\epsilon > 0$, we have $a \leq B(\tilde{y}, v^d, x^r)$, completing the proof.

4. COMPUTATION OF OPTIMAL POLICIES: NORMAL DEMANDS

Before turning to normal demands, we summarize what is required to compute an optimal policy in the general case. Here and in the next section, we restrict attention to average costs; the discounted case is similar.

As in property (b) of Section 1, computation of x^{r*} requires the solution of a single equation. Solution of IH^d is less simple. First, we can define an equivalent problem with no order lag by the transformation mentioned in Section 1. Define x^d = all inventory in the system, plus outstanding orders = $v^d + \sum_{i=1}^{L} y^i$. The new problem has the single state variable x^d, action y, dynamics $x_{n-1}^d = x_n{}^d + y_n - u_n$, order cost function $c^d(y)$, and one-period expected costs $D^L(x^d + y) + P^L(x^d + y)$, where

$$D^L(x) = ED[x - u^{(L)}],\ P^L(x) = EP[x - u^{(L)}]. \qquad (10)$$

(As in Section 1, a constant A^d representing short-term costs outside our control is deducted from the problem; this plays no role at all in IH^d.)

From here, the current state of the art requires discretization. Several algorithms are available for solving the discrete version of IH^d, e.g., Veinott and Wagner [1965] and Federgruen and Zipkin [1984b]. Both require multiple evaluations of the functions defined in (10). These computations are the concern of the rest of this section.

Evaluation of $D^L(x) = h^d(x - \mu^{(L)})$ is no problem. Turning to P^L, we let $f^{(i)}$ denote the density of $u^{(i)}$ and let $F^{(i)}$ be its cumulative distribution function (cdf). It is possible to show that

$$R(x) = p^s[\mu^{(l+1)} - x] + (p^s + h^r) \int_{-\infty}^{x} F^{(l+1)}(t)dt - h^d\mu. \qquad (11)$$

where $p^s = h^d + p^r$, and we have

$$P^L(x) = \int_{x-x^{r*}}^{\infty} [R(x - t) - R(x^{r*})]f^{(L)}(t)dt. \qquad (12)$$

In general, (12) evidently requires numerical integration. (The computation is still easier than the finite-horizon analogue, which requires an additional integration.)

Now assume u is normally distributed with mean μ and standard deviation σ. (This distribution violates the assumption that $u \geq 0$, so the previous results may hold only approximately. Since $\Pr\{u < 0\}$ is usually very small for practical μ and σ, we ignore this objection.) After defining some notation, we shall give an expression for P^L and then justify it. Let $\Phi(\cdot)$ = standard normal cdf; $\phi(\cdot)$ = standard normal density. If ξ_1 and ξ_2 have a bivariate normal distribution, with marginal densities ϕ and correlation ρ, let $\Phi(\xi_1, \xi_2; \rho)$ = cdf of (ξ_1, ξ_2). Let $\sigma^{(i)} = i^{1/2}\sigma$ be the standard deviation of $u^{(i)}$. Define

$$\tau_1(x) = -[x - (x^{r*} + \mu^{(L)})]/\sigma^{(L)},$$

$$\tau_2(x) = [x - \mu^{(L+l+1)}]/\sigma^{(L+l+1)},$$

$$\tau_3(x) = -\{x - (x^{r*} + \mu^{(L)}) - [\sigma^{(L)}/\sigma^{(l+1)}]^2[x^{r*} - \mu^{(l+1)}]\}$$
$$/[\sigma^{(L)}\sigma^{(L+l+1)}/\sigma^{(l+1)}],$$

$$\nu^{r*} = [x^{r*} - \mu^{(l+1)}]/\sigma^{(l+1)},$$

$$\epsilon_1(x) = \Phi[\tau_3(x)]\phi[\tau_2(x)]/\sigma^{(L+l+1)},$$

$$\epsilon_2(x) = \Phi(\nu^{r*})\phi[\tau_1(x)]/\sigma^{(L)},$$

$$\Theta(x) = x\Phi(x) + \phi(x).$$

(Note, $\Theta'(x) = \Phi(x)$. Also, $\tau_3(x)$ simplifies to $-\{x - [(L + l + 1)/(l + 1)]x^{r*}\}/[(L + l + 1)/(l + 1)]^{1/2}L^{1/2}\sigma$. The notation above, however, will facilitate the discussion in Section 5.) Then,

$$P^L(x) = p^s\iota(x) - (p^s + h^r)\kappa(x), \tag{13}$$

where

$$\iota(x) = \sigma^{(L)}\Theta[\tau_1(x)],$$

$$\kappa(x) = \sigma^{(l+1)}\Theta(\nu^{r*})\Phi[\tau_1(x)] - \{[\sigma^{(L+l+1)}]^2\epsilon_1(x)$$
$$- [\sigma^{(L)}]^2\epsilon_2(x)\} - [x - \mu^{(L+l+1)}]$$
$$\cdot\Phi[\tau_1(x), \tau_2(x); -\sigma^{(L)}/\sigma^{(L+l+1)}].$$

Note that all the functions required to evaluate P^L are available in standard packages, e.g., IMSL [1977]; while not quite simple, the computation is easier than numerical integration.

In principle, the derivation of formula (13) is an elementary integration problem, but it is sufficiently involved to warrant an outline. We first compute the derivative of P^L:

$$P^{L'}(x) = \int_{x-x^{r*}}^{\infty} R'(x - t)f^{(L)}(t)dt$$
$$= -p^s\Phi[\tau_1(x)] + (p^s + h^r)\int_{x-x^{r*}}^{\infty} F^{(l+1)}(x - t)f^{(L)}(t)dt, \tag{14}$$

using (11). The integral in (14) is precisely $\Pr\{u^{(L)} \geq x - x^{r*}, u^{(L)} + u^{(l+1)} \leq x\}$. This expression may be written in standardized form: Define $\zeta_1 = -[u^{(L)} - \mu^{(L)}]/\sigma^{(L)}$ and $\zeta_2 = [u^{(L)} + u^{(l+1)} - \mu^{(L+l+1)}]/\sigma^{(L+l+1)}$. The correlation between ζ_1 and ζ_2 can be shown to be $-\sigma^{(L)}/\sigma^{(L+l+1)}$, so (14) becomes

$$P^{L'}(x) = -p^s\Phi[\tau_1(x)] + (p^s + h^r)\Phi[\tau_1(x), \tau_2(x); -\sigma^{(L)}/\sigma^{(L+l+1)}]. \quad (15)$$

It is easy to show that $\lim_{x\to\infty}P^L(x) = 0$. Therefore,

$$P^L(x) = -\int_x^\infty P^{L'}(t)dt.$$

Integrating the two terms of (15) separately will lead to ι and κ. The first term is straightforward, since $\lim_{x\to\infty}\iota(x) = 0$, and $\iota'(x) = -\Phi[\tau_1(x)]$. For the second term of (15), $\lim_{x\to\infty}\kappa(x) = 0$ as well (examining κ term by term), so we must show $\kappa'(x) = -\Phi[\tau_1(x), \tau_2(x); -\sigma^{(L)}/\sigma^{(L+l+1)}]$.

For any bivariate normal (ξ_1, ξ_2), $(\xi_1 | \xi_2) \sim N[\rho\xi_2, (1-\rho^2)^{1/2}]$ and $(\xi_2 | \xi_1) \sim N[\rho\xi_1, (1-\rho^2)^{1/2}]$. This fact and tedious algebra can be used to show

$$P\{\zeta_1 \le \tau_1(x) | \zeta_2 = \tau_2(x)\} = \Phi[\tau_3(x)],$$

and

$$P\{\zeta_2 \le \tau_2(x) | \zeta_1 = \tau_1(x)\} = \Phi(\nu^{r*}),$$

and, therefore,

$$(d/dx)\Phi[\tau_1(x), \tau_2(x); -\sigma^{(L)}/\sigma^{(L+l+1)}] = \epsilon_1(x) - \epsilon_2(x).$$

More algebra demonstrates

$$\epsilon_1'(x) = -[\sigma^{(L+l+1)}]^{-2}\{[\sigma^{(l+1)}\phi(\nu^{r*})][\phi(\tau_1(x))/\sigma^{(L)}] + [x - \mu^{(L+l+1)}]\epsilon_1(x)\},$$

and

$$\epsilon_2'(x) = [\sigma^{(L)}]^{-2}\{[\sigma^{(l+1)}\nu^{r*}\Phi(\nu^{r*})][\phi(\tau_1(x))/\sigma^{(L)}] - [x - \mu^{(L+l+1)}]\epsilon_2(x)\}.$$

The last three equations yield the desired result.

We remark that a closed form expression for P^L can be derived from (12) for Erlang-distributed demands as well.

5. SEVERAL RETAIL OUTLETS

Now consider the case of J retail outlets. A subscript j will index the outlets, and an overbar will indicate a vector over j. Thus, $\bar{x}^r = (x_j^r)_{j=1}^J$ is the vector of inventories plus outstanding shipments, and $\bar{z} = (z_j)_{j=1}^J$ is the vector of current shipments. Each outlet experiences a normal demand, so $\bar{u} = (u_j)_{j=1}^J$ is a normal random vector. While demands in different periods remain independent, correlations between components of $\bar{u}$ are allowed. Let μ_j and σ_j denote the parameters of u_j. Assume the u_j have equal coefficients of variation σ_j/μ_j. The symbols x^r, z and u will represent sums over the index j; in particular, z is the total amount withdrawn from the depot, and system-wide demand u itself is normal. We assume the outlets all have the same cost factors, c^r, p^r and h^r. We shall retain z explicitly in the formulation and describe $\bar{z}$ as the *allocation* of z among outlets. We continue assuming $\alpha = 1$.

The system has state $(\tilde{y}, v^d, \bar{x}^r)$ and action $(y, z, \bar{z})$. The constraints and dynamics are given by the following equalities and inequalities:

$$y \geq 0,\ z \geq 0,\ x^r + z \leq v^d + y^L, \qquad \textstyle\sum_{j=1}^{J} z_j = z,\ \bar{z} \geq 0,$$

$$v_{n-1}^d = v_n{}^d + y_n{}^L - u_n,\ \bar{x}_{n-1}^d = \bar{x}_n{}^d + \bar{z}_n - \bar{u}_n,\ \tilde{y}_{n-1} = (y_n, y_n{}^1, \cdots, y_n^{L-1}).$$

Total one-period costs are given by

$$c^d(y) + D(v^d + y^L) + c^r z + \textstyle\sum_{j=1}^{J} R_j(x_j^r + z_j),$$

where

$$R_j(x) = -\ h^d(x - \mu_j^{(l)}) + p^r E[u_j^{(l+1)} - x]^+ + (h^d + h^r)E[x - u_j^{(l+1)}]^+.$$

We now approximate this problem by one having a single outlet, like those of prior sections. We use an approach introduced in Federgruen and Zipkin [1984a, 1984c] related to that of Eppen and Shrage, which performs extremely well when the depot cannot hold stock. After describing the approach we shall argue that, if anything, it should work better in the present case.

Form a *relaxed problem* by dropping the constraints $\bar{z} \geq 0$. Now, the current choice of $\bar{z}$ in no way affects the achievable future values of $\bar{x}^r + \bar{z}$, and hence future costs. We may thus choose $\bar{z}$ to minimize current costs only, that is, to solve the following problem:

$$\min_{\bar{z}} \textstyle\sum_{j=1}^{J} R_j(x_j^r + z_j), \quad \text{subject to:} \quad \sum_{j=1}^{J} z_j = z. \qquad (16)$$

But the minimal cost of (16) can be written as $R(x^r + z)$, where

$$R(x) = -\ h^d(x - \mu^{(l)}) + p^r E[u^{(l+1)} - x]^+ + (h^d + h^r)E[x - u^{(l+1)}]^+,$$

and $u^{(l+1)}$ denotes, *not* system-wide demand in $l + 1$ periods, but rather a normal random variable with the same mean $\mu^{(l+1)} = (l + 1) \sum_{j=1}^{J} \mu_j$ and a different (larger) standard deviation $\sigma^{(l+1)} = (l + 1)^{1/2} \sum_{j=1}^{J} \sigma_j$. Furthermore, the vector $\bar{x}^r$ now affects the problem only through x^r. The relaxed problem thus reduces precisely to problem *IH*, understanding that $u^{(l+1)}$ and hence R have the new interpretations given here. The results of Sections 1–3 can, therefore, be applied directly to determine an optimal policy. In Section 4 the symbol $\sigma^{(l+1)}$ has the meaning given here, $\sigma^{(L)}$ is the true standard deviation of L-period total-system demand, and $\sigma^{(L+l+1)} = \{[\sigma^{(L)}]^2 + [\sigma^{(l+1)}]^2\}^{1/2}$. Because we arrived at this problem by enlarging the original action sets, the optimal cost of this version of *IH* provides a *lower bound* on the true optimal cost of the original problem.

The optimal policy for this instance of *IH* provides a feasible policy for orders and total withdrawals z. To complete the specification of a policy for the original problem, we require an allocation rule. For this purpose, we propose the *myopic allocation* policy (following Federgruen and Zipkin [1984c]), which sets $\bar{z}$ to minimize the *true* current costs, that is, to solve the problem

$$\min_{\bar{z}} \textstyle\sum_{j=1}^{J} R_j(x_j^r + z_j), \quad \text{subject to:} \quad \textstyle\sum_{j=1}^{J} z_j = z, \; \bar{z} \geq 0. \qquad (17)$$

The true total cost of this feasible policy is, of course, an *upper bound* on the optimal cost of the original problem. We shall argue, partly on the basis of earlier empirical results, that the relaxed problem provides a close approximation to the cost of this policy; in other words, the difference between the upper and lower bounds is small. We can then conclude both that the policy is a good one, and that the relaxed problem provides a good cost approximation.

Clearly, the relaxed problem reproduces exactly the costs represented by the functions $c^d(\cdot)$, $D^L(\cdot)$ and $c^r z$. Any error arises in the measurement of $\sum_j R_j$. Comparison of (16) and (17) shows that the approximation would be exact, if the constraints $\bar{z} \geq 0$ were never binding in (17). The accuracy of the approximation as a whole, therefore, depends on how often these constraints are essential and the resulting effects of relaxing the binding constants on the total cost of (16).

Let a "cycle" mean a sequence of periods starting with the arrival of an order at the depot and ending just before the next order arrives. A cycle consists of some "ample-stock periods," when the depot has enough stock to ship up to the critical number x^{r*}, one period when the depot has positive but less-than-ample stock (the "allocation period"), and finally some "empty periods" when the depot has no stock, so $z = 0$. (Not every cycle, of course, need have all three kind of periods.)

Let $x_j^{r^*}$ denote the critical number for outlet j alone (i.e., $x_j^{r^*}$ minimizes R_j), and $\bar{x}^{r^*} = (x_j^{r^*})_{j=1}^{J}$. A direct calculation shows $\sum_j x_j^{r^*} = x^{r^*}$. Suppose a cycle begins with

$$\bar{x}^r \leq \bar{x}^{r^*}. \qquad (18)$$

If there is ample stock in the first period, we have $z = x^{r^*} - x^r \geq 0$, and it is straightforward to show that the solution to both (16) and (17) is

$$\bar{z} = \bar{x}^{r^*} - \bar{x}^r. \qquad (19)$$

In particular, $\bar{z} > 0$ is inessential in (17). The only way (18) could fail in subsequent periods of the cycle, furthermore, is if some demands were negative, events that we assume to have very low probability, as explained in the previous section. It is possible to show that under the proposed policy, from any initial state, we shall arrive at a cycle where (18) holds in a finite number of periods with probability 1. Subject to the qualification expressed earlier about negative demands, therefore, (18) will hold in *all* periods of *all* subsequent cycles, and hence $\bar{z} \geq 0$ is inessential in all but a finite number of periods with ample stock.

In the allocation period, $\bar{z} \geq 0$ may be essential and it (almost surely) will be essential in the empty periods. Observe, however, that $\bar{z} \geq 0$ is inessential in (17) if and only if the solution to (17) satisfies the following condition:

$$[x_j^r + z_j - \mu_j^{(l+1)}]/\sigma_j^{(l+1)} - [x^r + z - \mu^{(l+1)}]/\sigma^{(l+1)} = 0, \quad j = 1, \cdots, J. \tag{20}$$

For given $(x^r + z)$, roughly speaking, the cost of (17) increases with deviations from (20). (Cf. Zipkin [1984] for a more formal treatment.) The process described by these differences during the allocation and empty periods evolves in just the same way as it does during an entire cycle of a system without central stock operating under an (s, S) policy. In both cases the system starts with (20) nearly satisfied; for the remainder of the cycle the differences are generated by the demand process only. Also, the cost functions in the two cases have the same form (refer to (11)). And, as shown empirically in Federgruen and Zipkin [1984a, 1984c], the cost effects of such "imbalances" are quite small in the case of no central stock, so they must be small here as well.

All of these observations taken together suggest that the overall approximation provided by the relaxed problem should indeed be very accurate.

The assumption of equal σ_j/μ_j enters the argument only in the last step concerning the allocation and empty periods. As shown in Federgruen and Zipkin [1984c], with very unequal σ_j/μ_j, if we start with (20), the imbalance grows very fast and results in much higher cost. Federgruen and Zipkin [1984a] present a special modified allocation rule, which solves this problem satisfactorily when there is no central stock. The method does not apply directly to the current case, unfortunately, but we suspect some related approach can be fashioned to work just as well. Similarly, further research will hopefully show how to adapt the techniques of Federgruen and Zipkin [1984a] and Zipkin [1982] for unequal penalty and holding costs.

Equation 19 expresses a remarkable fact: Under the proposed policy the shipments to individual outlets are the same as if each outlet followed its own critical-number policy, and the depot simply satisfied the resulting orders when there is enough stock to do so (in ample-stock periods). The only period in a cycle requiring an explicit allocation decision is the allocation period (hence the name). Thus, even though the original model describes a system under centralized control, a maximally decentralized policy should perform very well. This situation contrasts markedly with the case of no central stock, where no such interpretation is possible. Rosenfield and Pendrock [1980] among others have noticed before that centralizing some inventory *permits* some decentralization of decision making, but we believe the discussion above provides the first strong evidence that such an approach is economically effective. We emphasize, however that our construction of the cost functions for *IH* to determine the depot's order policy (especially the calculation of $\sigma^{(L+l+1)}$) uses data on the individual outlets in a nontrivial way.

ACKNOWLEDGMENTS

This research was supported in part by the Faculty Research Fund, Graduate School of Business, Columbia University. We thank the referees for many valuable suggestions.

REFERENCES

BERTSEKAS, D., AND S. SHREVE. 1978. *Stochastic Optimal Control: The Discrete Time Case.* Academic Press, New York.

CLARK, A. 1958. A Dynamic, Single-Item, Multi-Echleon Inventory Model, Report, Rand Corporation, Santa Monica.

→ CLARK, A., AND H. SCARF. 1960. Optimal Policies for a Multi-Echelon Inventory Problem. *Mgmt. Sci.* **6,** 475–490.

DYNKIN, E., AND A. YUSHKEVICH. 1979. *Controlled Markov Processes.* Springer, New York.

EPPEN, G., AND L. SCHRAGE. 1981. Centralized Ordering Policies in a Multi-warehouse System with Leadtimes and Random Demand. In *Multi-Level Production/Inventory Control Systems: Theory and Practice*, pp. 51–69, L. Schwarz (ed.). North-Holland, Amsterdam.

→ EVANS, R. 1981. Markov Chain Design Problems. *Opns. Res.* **29,** 959–971.

FEDERGRUEN, A., AND P. SCHWEITZER. 1981. Nonstationary Markov Decision Problems with Converging Parameters. *J. Optim. Theory Appl.* **34,** 207–241.

FEDERGRUEN, A., AND P. ZIPKIN. 1984a. Allocation Policies and Cost Approximations for Multilocation Inventory Systems, Research Working Paper, Graduate School of Business, Columbia University, New York (to appear in *Naval Research Logistics Quarterly*).

FEDERGRUEN A., AND P. ZIPKIN. 1984b. An Efficient Algorithm for Computing Optimal (*s*, *S*) Policies, Research Working Paper, Graduate School of Business, Columbia University, New York (to appear in *Operations Research*).

→ FEDERGRUEN, A., AND P. ZIPKIN. 1984c. Approximations of Dynamic Multilocation Production and Inventory Problems, *Mgmt. Sci.* **30,** 69–85.

GROSS, D., R. SOLAND AND C. PINKUS. 1981. Designing a Multi-Product, Multi-Echelon Inventory System. In *Multi-Level Production/Inventory Control Systems: Theory and Practice*, pp. 11–49, L. Schwarz (ed.). North-Holland, Amsterdam.

HEYMAN, D., AND M. SOBEL. 1984. *Stochastic Models in Operations Research*, Vol. II. McGraw Hill, New York.

IGLEHART, D. 1963a. Dynamic Programming and Stationary Analysis of Inventory Problems. In *Multistage Inventory Models and Techniques*, pp. 1–31, H. Scarf, D. Gilford and M. Shelly (eds.). Stanford University Press, Stanford.

IGLEHART, D. 1963b. Optimality of (s, S) Policies in the Infinite Horizon Dynamic Inventory Problem. *Mgmt. Sci.* **9,** 259–267.

IGNALL, E., AND A. VEINOTT. 1969. Optimality of Myopic Inventory Policies for Several Substitute Products. *Mgmt. Sci.* **15,** 284–304.

INTERNATIONAL MATHEMATICAL AND STATISTICAL LIBRARIES. 1977. *IMSL Reference Manual*, Ed. 6. International Mathematical and Statistical Libraries, Inc., Houston.

ROSENFIELD, D., AND M. PENDROCK. 1980. The Effects of Warehouse Configuration Design on Inventory Levels and Holding Costs. *Sloan Mgmt. Rev.* **21,** (4), 21–33.

SCARF, H. 1960. The Optimality of (s, S) Policies in the Dynamic Inventory Problem. In *Mathematical Methods in the Social Sciences*, pp. 196–202, K. Arrow, S. Karlin and P. Suppes (eds.). Stanford University Press, Stanford.

SCHMIDT, C., AND S. NAHMIAS. 1981. Optimal Policy for a Two-Stage Assembly System under Random Demand, Working Paper, University of Santa Clara, Santa Clara.

VEINOTT, A. 1965. Optimal Policy for a Multi-Product, Dynamic, Nonstationary Inventory Problem. *Mgmt. Sci.* **12,** 206–222.

VEINOTT, A. 1971. Least d-Majorized Network Flows with Inventory and Statistical Applications. *Mgmt. Sci.* **17,** 547–567.

VEINOTT, A., AND H. WAGNER. 1965. Computing Optimal (s, S) Inventory Policies. *Mgmt. Sci.* **11,** 525–552.

ZIPKIN, P. 1984. Approximations of Dynamic, Multilocation Production and Inventory Problems II: Centrally Held Stocks, Research Working Paper, Graduate School of Business, Columbia University, New York (to appear in *Mathematics of Operations Research*).

ZIPKIN, P. 1982. Exact and Approximate Cost Functions for Product Aggregates. *Mgmt. Sci.* **28,** 1002–1012.