Probability in the Engineering and Informational Sciences, **11**, 1997, 107–135. Printed in the U.S.A.

STOCHASTIC INVENTORY MODELS WITH LIMITED PRODUCTION CAPACITY AND PERIODICALLY VARYING PARAMETERS

YOSSI AVIV AND AWI FEDERGRUEN
Graduate School of Business
Columbia University
New York, New York 10027

We consider a single-item, periodic-review inventory model with uncertain demands in which each period's production volume is limited by a capacity level. The demand distributions, capacity levels, and cost parameters vary according to a periodic pattern. We prove that modified base-stock policies are optimal for the finite-horizon planning model and for both the infinite-horizon discounted and undiscounted cost criterion. We further show that the optimal base-stock levels can be calculated via a simple but efficient value-iteration method. Finally, we have conducted a comprehensive numerical study to ascertain the efficiency of this solution method as well as various qualitative properties of the performance of capacitated production/inventory systems under periodically varying demand and cost patterns.

1. INTRODUCTION AND SUMMARY

We consider capacitated, periodic review production/inventory systems in which demands fluctuate from period to period, partially because of systematic periodically varying factors and partially because of intrinsic uncertainties. The combination of stochastic seasonal demand and capacitated production is common in many industries, as demonstrated by Krane and Braun [16], Fair [6], and

 0269-9648/97 $11.00 + .10

Bush and Cooper [4]. It is further complicated when cost structures and capacities vary according to a periodic pattern.

We study in particular the single-item model in which demands are independent across time and in which each period's production volume is limited by a (possibly period-dependent) capacity level. Production costs are proportional to the production volume. All other cost components (in particular, holding and stockout costs) can be expressed as a function of the system's inventory position, equal to inventory on hand plus outstanding orders minus backlogs. The capacity, all cost parameters and cost functions, as well as the demand distributions, follow a periodic pattern, with periodicity K. We assume that unsatisfied demand is (fully) backlogged.

We address both the finite- and infinite-horizon models. In the former, the objective is to minimize expected total (discounted or undiscounted) costs; in the latter, we address both the objective of minimizing total expected discounted and long-run average costs. We prove that a *periodic modified base-stock policy* is optimal for each of these models and associated objectives, under the assumption that the one-step expected costs depend convexly on the inventory position and a few additional regularity conditions. Such a base-stock policy specifies a target base-stock level for each of the K period types; it prescribes in each period a production order to increase the inventory position to a level as close as possible to the period's base-stock level. In the infinite-horizon models, the base-stock levels only depend on the period type; in the finite-horizon models, they depend in addition on the number of periods until the end of the horizon. We also prove that a simple successive approximation method can be used to efficiently compute the optimal base-stock levels. Finally, we have conducted a comprehensive numerical study to ascertain the efficiency of this solution method as well as various qualitative properties of the performance of capacitated production/inventory systems under periodically varying demand and cost patterns. We have focused in particular on the following:

(a) the impact of demand variability on the system-wide costs and optimal base-stock levels,
(b) the impact of the relative cost of carrying inventories versus the cost of backlogs, on the optimal base-stock levels, and
(c) an assessment of the trade-off between capacity and inventory investments.

One of the biggest challenges in managing production/inventory systems is the efficient matching of capacities and demands when their patterns fail to be synchronized (e.g., when demands vary greatly due to seasonal and promotional factors). (In many industries, 30–50% of annual sales are concentrated in a 1-2-month period.) There are two mechanisms to alleviate the problems and costs incurred by this synchronization problem: (a) *smoothing* of the demand pattern by, for example, the elimination of price promotions or by the implementation of off-peak-load pricing schemes (consider, e.g., Wal-Mart's "everyday low prices" sales philosophy and price discounts offered during off-seasons), and (b) the adoption of flexible capacity, that is, adjusting the capacity on a periodic, seasonal basis in (partial) synchronization with the demand pattern.

To date, few if any analytical tools are available to quantify the benefits of the preceding synchronization programs. In fact, if flexible capacity can be adopted, it is not even clear how these capacity levels are most effectively varied over time, that is, how a total capacity budget is best allocated over the individual periods. The last part of our numerical study focuses on the benefits of the preceding pair of synchronization efforts.

Our structural results generalize those obtained for simpler models. In particular, Federgruen and Zipkin [8,9] have proven the optimality of base-stock policies in the *stationary* version of our model where demands are identically distributed and all parameters and cost functions are constant across time. Karlin [14,15] and Zipkin [27] have proven this optimality result for the *uncapacitated* version of the model, that is, where no capacity limits prevail. (Karlin [14] addressed the discounted cost objective, while confining himself to the case of stationary cost parameters; Zipkin extended Karlin's results to the general uncapacitated model.)

As far as computational methods are concerned, Tayur [26] developed for the stationary version of this model an exact solution method, Glasserman and Tayur [11] gave a simulation-based heuristic, and Glasserman [10] provided approximations and bounds for the optimal base-stock levels. The papers by Karlin [15] and Zipkin [27] proposed exact solution methods for the uncapacitated special case of our model, however, with no numerical or theoretical characterization of their efficiency. All of these methods are tailored to specific model assumptions; the algorithm proposed in this paper is an application of a general-purpose method that is very easy to code and easy to adjust to variants of the model.

Other related work includes that by Song and Zipkin [25] and Sethi and Cheng [24] dealing with the uncapacitated special case of our model, in which the parameters are Markov-modulated; that is, they fluctuate as a function of an underlying Markov chain. (The periodic structure treated in this paper can clearly be modeled within this framework, the state of the modulating Markov chain representing the period type.) Metters [18] presented heuristics for a *multi-item* version of our model. (After completing the initial version of this paper, we became aware of Kapuscinski and Tayur [13], who independently established the optimality of a modified base-stock policy in the special case of our model where costs are linear and only the demand varies periodically, while all other model parameters are stationary. These authors also proposed a simulation-based method for identifying optimal base-stock levels.)

The remainder of this paper is organized as follows. In Section 2, we specify the model and introduce notation. Section 3 addresses the finite-horizon problem models and Section 4 the infinite-horizon discounted cost problem. Section 5 characterizes the asymptotic behavior of the discounted cost model for large discount factors; the results here are of interest by themselves, while providing the foundation for the analysis of the infinite-horizon long-run average cost objective, the subject of Section 6. In Section 7, we develop the successive approximation method and prove its convergence to an optimal policy. Finally, Section 8 reports on the numerical study described earlier.

2. MODEL DESCRIPTION AND NOTATION

In this section, we specify the model and introduce the basic notation. At the beginning of each period, a decision is made about whether or not to place an order and, if so, of what size. Demands in consecutive periods are independent. Those arising in the jth period of any cycle of K periods are identically distributed as the random variable D_j, $j = 1, \ldots, K$. We use the following notation:

$x =$ the inventory position at the beginning of a period, before ordering,

$y =$ the inventory position at the beginning of a period, after ordering,

$b_j =$ the capacity in periods of type j,

$c_j =$ the variable production cost rate in periods of type j.

We assume that the expected value of all other cost components that are charged to a period of type j can be expressed as a function $G_j(y)$. Assume, for example, that production orders become available after a lead time of $L \geq 0$ periods and that the carrying (backlogging) cost incurred for an inventory (backlog) of x^+ (x^-) units at the *end* of a period of type l is given by a function $h_l(x^+)\,[p_l(x^-)]$. Using a standard accounting device, we charge to each period the expected holding and backlogging costs incurred one leadtime later; that is,

$$G_j(y) = \mathsf{E}\{h_{j+L}([y - D_j - D_{j+1} \cdots - D_{j+L}]^+) + p_{j+L}([D_j + \cdots + D_{j+L} - y]^+)\} \tag{1}$$

where all subscripts in Eq. (1) are taken mod K.

We make the following assumptions regarding the growth rate of the functions G_j and the finiteness of moments of the demand distributions: we write $\phi(x) = O(\psi(x))$ for any pair of functions $\phi(\cdot)$, $\psi(\cdot)$ if a constant C exists such that $\phi(x) \leq C[\psi(x) + 1]$ for all x. For any pair of *sequences* of functions $\{\phi_l(\cdot)\}_{l=1}^{\infty}$, we write $\phi_l(\cdot) = O(\psi_l(\cdot))$ if the *same* bounding constant C can be used for *all* $l = 1, 2, \ldots$.

Assumption 1: G_j is convex and $\lim_{|y|\to\infty} G_j(y) = \lim_{|y|\to\infty}[c_j y + G_j(y)] = \infty$ for all $j = 1, \ldots, K$.

Assumption 2: $G_j(y) = O(|y|^{\rho})$ for some positive integer ρ $(J = 1, \ldots, K)$.

Assumption 3: $\mathsf{E}[D_j^{\rho}] < \infty$ $(j = 1, \ldots, K)$.

Convexity of the one-step expected cost functions $\{G_j : j = 1, \ldots, k\}$ is satisfied under most commonly used cost structures; for example, in Eq. (1) it holds whenever the functions $\{h_j, p_j\}$ are linear or more generally convex. The second part of Assumption 1 is satisfied whenever the asymptotic marginal backlogging cost is in excess of the period's variable production cost rate; it precludes the trivial and unrealistic case where it is *never* beneficial to carry stock in anticipation of demands. Assumption 2 is similarly general; if the G_j functions are of the form given by Eq. (1), it is satisfied with $\rho = 1$ when $h_j(\cdot)$ and $p_j(\cdot)$ are linear or piecewise linear and with $\rho \geq 2$ when these functions are bounded by polynomials. Assumption 3 is necessary to guarantee that the expected cost over a single-period or multiperiod horizon remains finite; it is often required to ensure that the functions $G_j(\cdot)$ themselves are finite; see Eq. (1).

Our model can be formulated as a Markov decision problem (MDP) with countable state space $S = \{(x,j): x \text{ integer}; j = 1,\ldots,K\}$ and (finite) action sets $Y(x,j) = \{y: x \leq y \leq x + b_j\}$. The state of the system at the beginning of any period is given by the prevailing inventory level x and period type j. Finally, we write $j^+ = (j \bmod K) + 1$. Also, $\mathbf{1} = (1,1,\ldots)$ denotes the infinite vector of ones.

3. THE FINITE-HORIZON PROBLEM

In this section, we prove that a (modified) base-stock policy is optimal for the finite-horizon problem. Let $v_n^*(x,j)$ be the minimal expected discounted cost, over a horizon of n periods when starting in state $(x;j) \in S$. The functions v_n^* satisfy $v_0^* \equiv 0$ and

$$v_n^*(x,j) = -c_j x + \min_{x \leq y \leq x+b_j} J_n(y,j), \qquad n \geq 1, \tag{2}$$

where $J_n(y,j) = c_j y + G_j(y) + \alpha \cdot \mathsf{E}[v_{n-1}^*(y - D_j, j^+)]$.

Theorem 1:

(i) *The functions $J_n(y,j)$ are convex in y and $O(|y|^\rho)$ and have a finite minimum for all $n \geq 1$ and $j = 1,\ldots,K$. Let $\beta_{n,j}^*$ denote the smallest minimizer of $J_n(\cdot,j)$.*

(ii) *The base-stock policy with base-stock levels $\{\beta_{n,1}^*,\ldots,\beta_{n,K}^*\}$ achieves the minima in Eq. (2); $v_n^*(x,j)$ is convex in x and is $O(|x|^\rho)$.*

Proof: Note that $\mathsf{E}[(y - D_j)^\rho] \leq \mathsf{E}[(|y| + D_j)^\rho] = \sum_{r=0}^{\rho}\binom{\rho}{r}\mathsf{E}[D_j^{\rho-r}]|y|^r = O(|y|^\rho)$ by Assumption 3. The remainder of the proof is by standard induction, employing Assumptions 1-3 and verifying that $\lim_{|y|\to\infty} J_n(y,j) = \lim_{|x|\to\infty} v_n^*(x,j) = \infty$. ■

A standard induction proof also establishes the following corollary.

Corollary 1: *For all $n \geq 1$ and $(x,j) \in S$, $v_n^*(x,j)$ is nonincreasing and jointly convex in the capacity vector b.*

This corollary has important implications for *static resource-allocation problems*, such as where a common capacity pool (of B_j units in period $j = 1,\ldots,K$) needs to be shared by multiple items; it also has important implications when allocating capacities over the cycle, as in flexible capacity schemes; see Sections 1 and 8.

4. THE INFINITE-HORIZON DISCOUNTED COST PROBLEM

In this section, we show that a (modified) base-stock policy is optimal, when minimizing the expected discounted costs over an *infinite* horizon and that this policy satisfies the infinite-horizon optimality equation:

$$v(x,j) = (Tv)(x,j) \quad \text{for all } (x,j) \in S, \tag{3}$$

where

$$(Tv)(x,j) = \min_{x \le y \le x+b_j} \{c_j(y-x) + G_j(y) + \alpha \mathsf{E}[v(y-D_j, j^+)]\} \qquad \text{for all } v : S \to \mathbf{R}. \tag{4}$$

(A policy is said to satisfy the optimality equation for some solution v if in every state $(x,j) \in S$ it prescribes an action achieving the minimum in Eq. (4).) These results are obtained by showing in addition that $\{v_n^*\}$ converges to a solution of Eq. (3), when $\alpha < 1$. For MDPs with bounded one-step expected costs, the latter convergence result is usually obtained by showing that the T-operator is a contraction mapping with respect to the regular l_∞-norm. Because the one-step expected costs in our model are unbounded, we follow the approach of Lippman [17] and show that T is a contraction mapping with respect to the norm $\|\cdot\|_w$, defined as follows: for any function $u : S \to \mathbf{R}$,

$$\|u\|_w \doteq \sup_{(x,j) \in S} \frac{|u(x,j)|}{w(x)^\rho}, \quad \text{where } w(x) = \max(|x|,1).$$

We first need the following lemma: let $\eta = \max\{\sqrt[i]{\mathsf{E}[(D_j + b_j)^i]} : i = 1,\ldots,\rho;\ j = 1,\ldots,K\} < \infty$, by Assumption 3.

LEMMA 1: *For all $m = 1,\ldots,\rho$ and all $(x,j) \in S$,*

$$\sup_{x \le y \le x+b_j} \mathsf{E}[\max(|y-D_j|,1)^m] \le (\max(|x|,1) + \eta + 1)^m.$$

PROOF:

$$\begin{aligned}
\sup_{x \le y \le x+b_j} \mathsf{E}[\max(|y-D_j|,1)^m] &\le 1 + \sup_{x \le y \le x+b_j} \mathsf{E}[|y-D_j|^m] \le 1 \\
&\quad + \sup_{x \le y \le x+b_j} \mathsf{E}[(|y| + D_j)^m]. \\
&\le 1 + \mathsf{E}[(|x| + b_j + D_j)^m] \\
&= 1 + \sum_{i=0}^{m} \binom{m}{i} \mathsf{E}[(b_j + D_j)^i]|x|^{m-i} \\
&\le 1 + \sum_{i=0}^{m} \binom{m}{i} \eta^i |x|^{m-i} \\
&= 1 + (|x| + \eta)^m \le (|x| + \eta + 1)^m \\
&\le (\max(|x|,1) + \eta + 1)^m. \qquad \blacksquare
\end{aligned}$$

Let B be the Banach space of all real-valued functions $v : S \to \mathbf{R}$ with $\|v\|_w < \infty$ (endowed with the corresponding metric $d(u_1, u_2) = \|u_1 - u_2\|_w$ for all $u_1, u_2 \in B$). Also, let $v^*(x,j)$ denote the minimum expected total discounted costs over an infinite horizon when starting in state $(x,j) \in S$.

THEOREM 2:

(a) $v^* = \lim_{n\to\infty} v_n^*$ *(pointwise).*

(b) v^* *is the unique nonnegative solution in* B *of optimality Eq.* (3).

(c) *There exists a policy* f^* *that satisfies optimality Eq.* (3) *for* $v = v^*$, *which is a stationary modified base-stock policy (say, with* β^* *as the vector of base-stock levels) and which minimizes the expected total discounted costs over an infinite planning horizon.*

(d) *The sequence of finite-horizon optimal base-stock vectors* $\{\beta_n^*\}$ *is bounded, and every one of its limit points is an optimal base-stock vector in the infinite-horizon model.*

(e) v^* *is nonincreasing and convex in* b.

PROOF: Note that the set of feasible actions in any state $(x,j) \in S$ is finite. $\alpha < 1$ and Lemma 1 show that Assumptions 1 and 3 in Lippman [17] are satisfied. Also, by Assumption 2 in this paper we have for any $(x,j) \in S$ that $\sup_{x\le y\le x+b_j}[c_j(y-x) + G_j(y)] \le c_j b_j + \max_{l=0,\ldots,b_j} G_j(x+l) = O(|x|^p)$ so that $\|\sup_{x\le y\le x+b_j}[c_j(y-x) + G_j(y)]\|_w < \infty$, verifying Assumption 2 in Lippman [17]. Lippman showed that, under his Assumptions 1–3, T is an N-stage contraction mapping with respect to the $\|\ \|_w$-norm for N sufficiently large; hence, $\lim_{n\to\infty} v_n^* = u^*$ (see Denardo [5]). Theorem 1 in Lippman implies part (b) and the fact that any policy satisfying the optimality equation for $v = v^*$ is optimal. Because v_n^* is convex in x and $\lim_{|x|\to\infty} v_n^*(x,j) = \infty$ for all $n \ge 1$ and $j = 1, \ldots, K$ (see Theorem 1 and its proof), the same properties apply to v^*. It follows that a modified base-stock policy is optimal, completing the proof of part (c).

To prove part (d), assume to the contrary that $\{\beta_n^*\}$ fails to be bounded; then, in view of Assumption 1, there exists $n \ge 1$ and $j \in \{1, \ldots, K\}$ with $G_j(\beta_{n,j}^*) \ge 2v^*(0,j)$, which implies that

$$\infty > v^*(0,j) \ge v_n^*(0,j) \ge v_n^*(\beta_{n,j}^*, j) \ge G_j(\beta_{n,j}^*) \ge 2v^*(0,j),$$

which is a contradiction. (The first inequality follows from part (a); the second inequality results from $\{v_n^*\}$ being nondecreasing, which is easily verified by induction; the third inequality follows from the definition of $\beta_{n,j}^*$; and the fourth inequality follows from $v_n^*(\beta_{n,j}^*, j) = G_j(\beta_{n,j}^*) + \alpha \mathrm{E}[v_{n-1}^*(\beta_{n,j}^* - D_j, j^+)]$.) Thus, $\{\beta_n^*\}$ has at least one limit point. If β^* is a limit point, and because β_n^* is integer-valued, there exists a subsequence $\{n_k\}_{k=1}^\infty$ such that $\beta_{n_k}^* = \beta^*$. For any $(x,j) \in S$: $v_{n_k}^*(x,j) \ge v_{n_k}^*(\beta_j^*, j)$ and, taking limits for $k \to \infty$, $v^*(x,j) \ge v^*(\beta_j^*, j)$. This implies that β^* is an optimal base-stock vector in the infinite-horizon model.

Part (e) is immediate from part (a) and Corollary 1. ∎

5. THE ASYMPTOTIC BEHAVIOR OF THE DISCOUNTED COST MODEL FOR LARGE DISCOUNT FACTORS

In this section, we investigate the asymptotic behavior of the optimal base-stock vector β^* and the minimum cost function v^* as the discount factor α converges to 1. (We make the dependency on α explicit by writing v_α^* for v^* and β_α^* for β^*.) To this end, we need the following additional assumption.

Assumption 4: Let $\mu_j = \mathrm{E}[D_j]$, $\mu = \sum_{j=1}^K \mu_j$, and $B = \sum_{j=1}^K b_j$. Then, $0 < \mu < B$.

Assumption 4 is the necessary and sufficient condition for the system to be stable under some (e.g., base-stock) policy (see Glasserman and Tayur [11]). In addition, we require the finiteness of one additional moment of the $\{D_j\}$ variables beyond those required in Assumption 3.

Assumption 3': $\mathrm{E}[D_j^{l+1}] < \infty$ $(j = 1, \ldots, K)$ has finite moments of all orders up to $\rho + 1$.

Assumption 3′ guarantees that the long-run average cost of some (e.g., base-stock) policies exists and is finite; see also Glasserman and Tayur [11], who required finite demand variances to ensure the existence of a finite long-run average cost value under linear holding and backlogging costs.

In Markov decision processes with *finite* state and action sets, it is known that a so-called Blackwell-optimal policy exists, that is, the *same* policy is optimal for *all* sufficiently larger discount factors α, and that the minimum cost function $v_\alpha^* = O((1-\alpha)^{-1})$ as $\alpha \to 1$. (In fact, in finite MDPs, v_α^* can be expanded as a Laurent series in $(1-\alpha)$ of the form $v_\alpha^* = \sum_{l=-1}^{\infty} w_l (1-\alpha)^l$; see Miller and Veinott [19].) In our model, we show that $v_\alpha^* = O((1-\alpha)^{-1})$ continues to apply. Moreover, even though a Blackwell-optimal policy fails to exist in general MDPs with *infinite* state or action sets, we show that the optimal base-stock vector β_α^* is at least bounded in α.

THEOREM 3:

(a) *There exists a function* $w_1 : S \to \mathbf{R}$ *and a constant* $\bar{g}$ *such that*

$$v_\alpha^*(x,j) \le w_1(x,j) + \frac{\bar{g}}{(1-\alpha)}, \quad \text{for all } (x,j) \in S.$$

(b) β_α^* *is bounded in* α.

(c) *For every sequence* $\{\alpha_n\}$ *of discount factors converging to* 1, *there exists a subsequence* $\{\alpha_{n_k}\}_{k=1}^{\infty}$ *and a base-stock vector* β^* *such that* $\beta_{\alpha_{n_k}}^* = \beta^*$ *for all* $k \ge 1$.

To prove this theorem as well as the existence of an optimal base-stock policy under the long-run average-cost criterion, we need to show that under certain policies, certain finite subsets of S can be reached with finite expected cost from any starting state. More specifically, for every modified base-stock policy β, we construct a specific nonstationary modified base-stock policy $\tilde{\pi}(\beta)$ and an associated *finite* set of states such that the expected cost incurred until reaching this set from any state $(x,j) \in S$ is finite and $O(|x|^{\rho+1})$.

Let $r_j = \min\{d \geq 0 : \Pr(D_j = d) > 0\}$ denote the essential infimum of D_j's distribution, $j = 1, \ldots, K$, and let $R = \sum_{j=1}^{K} r_j$. For each period $j = 1, \ldots, K$, we define a finite interval of states $S_j(\beta) = \{(x,j) \in S : t_j - b_j < x \leq t_j\}$, where the upper bounds $\{t_j : j = 1, \ldots, K\}$ are recursively defined by

$$t_K = \beta_K; \qquad t_j = t_{j+1} + r_j, \quad j = 1, \ldots, K-1. \tag{5}$$

$S(\beta) = \bigcup_{j=1}^{K} S_j(\beta)$ constitutes a finite *band* within the state space S. Let $\bar{S}(\beta) = \{(x;i) : 1 \leq i \leq K, x > t_i\}$ and $\underline{S}(\beta) = \{\hat{x} \in S \mid \hat{x} \notin S(\beta) \cup \bar{S}(\beta)\} = \{(x;i) : 1 \leq i \leq K, x \leq t_i - b_i\}$ denote the sets of states "above" and "below" $S(\beta)$, respectively. The band $S(\beta)$ is constructed in such a way as to guarantee the following properties.

LEMMA 2:

(a) *$\bar{S}(\beta)$ cannot be reached from $\underline{S}(\beta)$ without entering $S(\beta)$, under any policy.*

(b) *Under the base-stock policy with base-stock vector $t = t(\beta)$ (defined as in Eq. (5) for any vector β), $S_K(\beta)$ can be reached from any state $(x,j) \in S(\beta) \backslash S_K(\beta)$.*

PROOF:

(a) Assume the system is in state $(x,j) \in \underline{S}(\beta)$ and the next period in state $(y,j^+) \notin \underline{S}(\beta)$. Then, $t_{j^+} - b_{j^+} < y \leq x + b_j - r_j \leq t_j - b_j + b_j - r_j \leq t_{j^+}$, that is, $(y,j^+) \in S(\beta)$. (The last inequality holds as an equality for $j = 1, \ldots, K-1$ and it holds for $j = K$ because the t_j-values are nonincreasing in j; see Eq. (5).)

(b) We show for every state $(x,j) \in S(\beta) \backslash S_K(\beta)$ that the state (β_K, K) can be reached in $(K-j)$ periods, for example, by following the path $(t_{j+1}, j^{+1}), \ldots, (t_K, K)$ the likelihood of which path equals to $\prod_{l=j}^{K-1} \Pr\{D_l = r_l\} > 0$, because r_l is the essential infimum of D_l, $l = 1, \ldots, K$, and using Eq. (5). ∎

The (nonstationary) policy $\tilde{\pi}(\beta_\alpha^*)$ starts out prescribing actions according to the (stationary) modified base-stock policy $t = t(\beta_\alpha^*)$ (defined as in Eq. (5) for $\beta = \beta_\alpha^*$), until the first entrance to the set $S_K(\beta_\alpha^*)$ and switches to the policy β_α^* thereafter. (If the starting state is in $S_K(\beta_\alpha^*)$, β_α^* is used throughout.)

For any set of states $U \subset S$, and any policy π, let $C_{x,j}^{\pi}(U)$ be the expected *undiscounted* cost incurred under policy π until the first visit to the set U, when starting in state $(x,j) \in S$. We now show the following proposition.

PROPOSITION 1: *For all $(x,j) \in S \backslash S_K(\beta)$,*

(a) $C_{x,j}^{\tilde{\pi}(\beta)}(S(\beta)) = O(\max\{|x|^{\rho+1}, |\beta_K|^{\rho+1}\})$,

(b) $C_{x,j}^{\tilde{\pi}(\beta)}(S_K(\beta)) = O(\max\{|x|^{\rho+1}, |\beta_K|^{\rho+1}\})$.

PROOF:

(a) We distinguish among three cases:

(i) $(x,j) \in \underset{\sim}{S}(\beta)$: Let f_b denote the policy that places full capacity orders throughout. Note that

$$C_{x,j}^{\tilde{\pi}(\beta)}(S(\beta)) = C_{x,j}^{f_b}(S(\beta)) \leq C_{x,j}^{f_b}(\{(y,j): y > t_j - b_j\}). \qquad (6)$$

(The equality follows from $\tilde{\pi}(\beta)$ prescribing the same actions as f_b until the first exit from $\underset{\sim}{S}(\beta)$, which by Lemma 2 coincides with the first entry into $\underset{\sim}{S}(\beta)$. The set $\{(y,j): y > t_j - b_j\}$ is entirely outside $\underset{\sim}{S}(\beta)$ so that (again by Lemma 2) it cannot be reached without passing through $S(\beta)$, justifying the inequality in Eq. (6).) The upper bound in Eq. (6) represents the expected cost on a random walk, which makes transitions after a full cycle (of K periods) and whose increments are distributed as $(B - D)$, where $D = \sum_{j=1}^{K} D_j$. (This random walk has positive drift by Assumption 4.) The path traveled by the random walk from level x to $[t_j - b_j + 1, \infty)$ can be decomposed into consecutive ladder epochs (where a ladder epoch T is defined as the time required to advance by at least one unit; note that the distribution of T is independent of the starting state). Let $\bar{C}(x)$ denote the expected cost incurred until the first ladder epoch in the random walk when starting at level x. Note that $\bar{C}(x) = \mathsf{E} \sum_{l=0}^{T-1} \sum_{j=1}^{K} [c_j b_j + G_j(x_{lK+j})]$, where x_{lK+j} is the inventory level at the beginning of the $(lK+j)$th period in the ladder epoch. Clearly, $x - TB \leq x_{lK+j} \leq x + B$ for all $l = 0, \ldots, T-1$ and $j = 1, \ldots, K$. Also, by Assumption 2, there exists a constant $A > 0$ such that $G_j(x) \leq A|x|^{l+1} + A$ for all $j = 1, \ldots, K$. Thus, because the function $|x|^{l+1}$ is convex,

$$\begin{aligned}
\bar{C}(x) &\leq \left(\sum_{j=1}^{K} c_j b_j + AK \right) \mathsf{E}T + A\mathsf{E} \sum_{l=0}^{T-1} \sum_{j=1}^{K} |x_{lK+j}|^{l+1} \\
&\leq \left(\sum_{j=1}^{K} (c_j b_j + AK) \right) \mathsf{E}T \\
&\quad + AK\mathsf{E} \sum_{l=0}^{T-1} \max(|x - TB|^{l+1}, |x + B|^{l+1}) \\
&= O(|x|^{l+1}),
\end{aligned}$$

because $\mathsf{E}T^{l+1} < \infty$. (The latter follows from (1.6) in Janson [12] because, for the random walk's increment, $\mathsf{E}(|B - D|^{l+1}) \leq \mathsf{E}\{(B + D)^{l+1}\} < \infty$, by Assumption 3′.) Thus, there are constants $K_1, K_2 > 0$ such that

$$C_{x,j}^{\tilde{\pi}(\beta)}(S(\beta)) \leq C_{x,j}^{f_b}(\{(y,j) : y > t_j - b_j\}) \leq \sum_{y=x}^{t_j-b_j} \bar{C}(y)$$

$$\leq \sum_{y=x}^{\beta_K+R} (K_1|y|^\rho + K_2) = O(\max\{|x|^{\rho+1}, |\beta_K|^{\rho+1}\}).$$

(ii) $(x,j) \in S(\beta)\backslash S_K(\beta)$: Under $\tilde{\pi}(\beta)$ we increase the inventory level to t_j both in state (x,j) and in state $(t_j - b_j, j) \in \underline{S}(\beta)$; hence, $C_{x,j}^{\tilde{\pi}(\beta)}(S(\beta)) = C_{t_j-b_j,j}^{\tilde{\pi}(\beta)}(S(\beta)) - (x - t_j + b_j)c_j = O(|\beta_K|^{\rho+1})$ by part (i).

We conclude from cases (i) and (ii) that constants $K_3, K_4 > 0$ exist such that

$$C_{x,j}^{\tilde{\pi}(\beta)}(S(\beta)) \leq K_3 \max\{|x|^{\rho+1}, |\beta_K|^{\rho+1}\} + K_4,$$
$$\text{for all } x \leq t_j \text{ and } j = 1, \ldots, K. \quad (7)$$

(iii) $(x,j) \in \bar{S}(\beta)$: Let $C_{x,j}^1(S(\beta))$ denote the expected (undiscounted) cost incurred under policy $\tilde{\pi}(\beta)$, when starting in state (x,j), until the first exit from the set $\bar{S}(\beta)$ and $C_{x,j}^2(S(\beta)) = C_{x,j}^{\tilde{\pi}(\beta)}(S(\beta)) - C_{x,j}^1(S(\beta))$. ($C_{x,j}^2(S(\beta)) > 0$ may occur because after exiting $\bar{S}(\beta)$ the next state may be in $\underline{S}(\beta)$.) Observe first that $C_{x,j}^1(S(\beta))$ represents the expected costs incurred on a trajectory θ during which $\tilde{\pi}(\beta)$ does not prescribe any orders; that is, the inventory level is nonincreasing. Let $p_0 \doteq \Pr\{D_1 + \cdots + D_K = 0\} < 1$, by Assumption 4. While the system state is in $\bar{S}(\beta)$, the number of *full* cycles during which a given inventory level $y \leq x$ is maintained is stochastically smaller than a geometric random variable with mean $(1 - p_0)^{-1}$ (it is stochastically smaller because for $\beta_K \leq y \leq t_1$ one exists $\bar{S}(\beta)$ during a cycle even in the absence of any demand). This implies that the expected cost incurred while maintaining a given inventory level $y \leq x$ during the trajectory θ is bounded by $[(1 - p_0)^{-1} + 1]\sum_{j=1}^K G_j(y)$. Because the set of inventory levels encountered during θ is a subset of $\{\beta_K + 1, \ldots, x\}$, we have for appropriate constants $K_5, K_6 > 0$

$$C_{x,j}^1(S(\beta)) \leq [(1 - p_0)^{-1} + 1]$$
$$\times \sum_{y=\beta K}^{x} \sum_{j=1}^{K} G_j(y) \leq K_3 \max\{|\beta_K|^{\rho+1}, |x|^{\rho+1}\} + K_4. \quad (8)$$

To obtain a bound for $C_{x,j}^2(S(\beta))$, let $\{(x_n, j_n)\}$ denote the states visited under $\tilde{\pi}(\beta)$, where $(x_0, j_0) = (x,j)$. Define the function $\hat{C}: S \to \mathbf{R}$ by $\hat{C}(y,j) = C_{y,j}^{\tilde{\pi}(\beta)}(S(\beta))$ if $(y,j) \in \underline{S}(\beta)$ and 0 otherwise. Thus,

$$C^2_{x,j}(S(\beta)) \le \max_{n\ge 0} \mathsf{E}[\hat{C}(x_n - D_{j_n}, j_{n+1})\mathbf{1} \times \{t_{j_n} < x_n \le x \text{ and } x_n - D_{j_n} \le t_{j_{n+1}}\}]$$
$$\le \max_{n\ge 0} \mathsf{E}[\hat{C}(x_n - D_{j_n}, j_{n+1})\mathbf{1}\{\beta_K \le x_n \le x\}]$$
$$\le \max_{j=1,\dots,K} \max_{\beta_K \le y \le x} \mathsf{E}[\hat{C}(y - D_j, j^+)] \times \max_{j=1,\dots,K} \max_{\beta_K \le y \le x} [K_3 h_j(y) + K_4],$$

by Eq. (7), where for all $j = 1, \dots, K$, $h_j(y) \doteq \mathsf{E}[\max\{|y - D_j|^{\rho+1}, |\beta_K - D_j|^{\rho+1}\}] \le \mathsf{E}[|y - D_j|^{\rho+1}] + \mathsf{E}[|\beta_K - D_j|^{\rho+1}] \le K_7|y|^{\rho+1} + K_7|\beta_K|^{\rho+1} + K_8 \le 2K_7 \max\{|y|^{\rho+1}, |\beta_K|^{\rho+1}\} + K_8$, for appropriate constants $K_7, K_8 > 0$. (The second upper bound for $h_j(\cdot)$ is obtained using Binomial expansions of $(|y| + D_j)^{\rho+1}$ and $(|\beta_K| + D_j)^{\rho+1}$ and invoking Assumption 3′.) Thus,

$$C^2_{x,j}(S(\beta)) \le \max_{j=1,\dots,K} \max_{\beta_K \le y \le x} [2K_3K_7 \max\{|y|^{\rho+1}, |\beta_K|^{\rho+1}\} + K_4 + K_3K_8]$$
$$\le 2K_3K_7 \max\{|x|^{\rho+1}, |\beta_K|^{\rho+1}\} + K_4 + K_3K_8, \quad (9)$$

where the second inequality follows from the fact that the function $\max\{|y|^{\rho+1}, |\beta_K|^{\rho+1}\}$ is convex in y as the maximum of a convex function and a constant. Case (iii) thus follows from Eqs. (7) and (9).

(b) In view of part (a), it suffices to show that the total expected costs from the first entry into $S(\beta)$ until the first visit to $S_K(\beta)$ is uniformly bounded in the starting state $(x,j) \in S$. Note that until the first visit to $S_K(\beta)$, $\tilde{\pi}(\beta)$ prescribes the exact same actions as the stationary base-stock policy $t(\beta)$.

Consider now the process embedded on consecutive visits to $S(\beta)$. This is a semi-Markov process on the finite state space $S(\beta)$, with finite one-step expected costs and transition times, in view of part (a). (To show that the transition times have finite expectations, repeat part (a) with $G_j(\cdot)$ replaced by $\bar{G}(y) \equiv 1$ for all $(y,j) \in S$.) By Lemma 2(b), we know that the set $S_K(\beta)$ can be reached from any state in $S(\beta)\backslash S_K(\beta)$. Thus, let F denote the maximum first passage time to $S_K(\beta)$ (in this embedded semi-Markov process) for any state in $S(\beta)\backslash S_K(\beta)$, and let $\bar{C} = \max_{(y,j)\in S(\beta)\backslash S_K(\beta)} C^{\tilde{\pi}(\beta)}_{y,j}(S(\beta))$ denote the maximum expected cost incurred between consecutive visits to $S(\beta)$ (before reaching $S_K(\beta)$). We conclude that for all $(x,j) \in S\backslash S_K(\beta) : C^{\tilde{\pi}(\beta)}_{x,j}(S_K(\beta)) \le C^{\tilde{\pi}(\beta)}_{x,j}(S(\beta)) + F\bar{C}$. Part (b) thus follows from part (a). ∎

We are now ready to prove Theorem 3.

PROOF OF THEOREM 3:

(a) Consider an arbitrary base-stock vector β^0, and let $\beta^1 = t(\beta^0)$. Note that $t(\beta^1) = \beta^1$ (see Eq. (5)) so that the policy $\tilde{\pi}(\beta^1)$ is the stationary base-stock policy β^1. Clearly, for all $(x,j) \in S$,

$$v_\alpha^*(x,j) \le v_\alpha^{\tilde{\pi}(\beta^1)}(x,j) \le C_{x,j}^{\tilde{\pi}(\beta^1)}(S_K(\beta^1)) + \alpha \max_{y \in S_K(\beta^1)} \{v_\alpha^{\tilde{\pi}(\beta^1)}(y,K)\}$$
$$\le C_{x,j}^{\tilde{\pi}(\beta^1)}(S_K(\beta^1)) + \alpha v_\alpha^{\tilde{\pi}(\beta^1)}(\beta_K^1 - b_K, K). \qquad \textbf{(10)}$$

The last inequality follows from $v_\alpha^{\tilde{\pi}(\beta^1)}(y,K) = v_\alpha^{\tilde{\pi}(\beta^1)}(\beta_K^1,K) + c_K(\beta_K^1 - y)$ for *all* $\beta_K^1 - b_K \le y \le \beta_K^1$ because $\tilde{\pi}(\beta^1)$ is the base-stock policy β^1. To verify the second inequality, note that the total expected discounted cost under policy $\tilde{\pi}(\beta^1)$ can be decomposed into two terms: (i) the *discounted* cost over the first $(T-1)$ periods with T the period in which the set $S_K(\beta^1)$ is visited first (which is bounded by the *undiscounted* cost on this trajectory) and (ii) the infinite-horizon discounted cost from the first entry state in $S_K(\beta^1)$ multiplied by $\alpha^T \le \alpha$. Proposition 1 shows that the upper bounds in Eq. (10) are finite for $(x,j) \in S \backslash S_K(\beta^1)$; moreover, because $\tilde{\pi}(\beta^1)$ is the stationary base-stock policy β^1 we have $C_{x,j}^{\tilde{\pi}(\beta^1)}(S_K(\beta^1)) < \infty$ for $(x,j) \in S_K(\beta^1)$ as well, following the proof of part (b) of Proposition 1. Substituting $(x,j) = (\beta_K^1 - b_K, K)$ into Eq. (10), we obtain $v_\alpha^{\tilde{\pi}(\beta^1)}(\beta_K^1 - b_K, K) \le (1-\alpha)^{-1} C_{\beta_K^1 - b_K, K}^{\tilde{\pi}(\beta^1)}(S_K(\beta^1))$ so that part (a) follows with $w_1(x,j) = C_{x,j}^{\tilde{\pi}(\beta^1)}(S_K(\beta^1))$ and $\bar{g} = C_{\beta_K^1 - b_K, K}^{\tilde{\pi}(\beta^1)}(S_K(\beta^1))$.

(b) and (c)
Fix $j \in \{1, \ldots, K\}$. Because for all $l = 1, \ldots, K$, $\beta_{\alpha,l}^*$ is a minimum of the function $v_\alpha^*(\cdot, l)$, we have $v_\alpha^*(0,j) \ge v_\alpha^*(\beta_{\alpha,j}^*) \ge G_j(\beta_{\alpha,j}^*) + \alpha v_\alpha^*(\beta_{\alpha,j^+}^*, j^+)$. Iterating this inequality, we obtain

$$v_\alpha^*(0,j) \ge \left(\sum_{i=0}^{\infty} \alpha^{iK}\right) G_j(\beta_{\alpha,j}^*) + \sum_{l<j} \left(\sum_{i=1}^{\infty} \alpha^{iK+l-j}\right) G_l(\beta_{\alpha,l}^*)$$
$$+ \sum_{l>j} \left(\sum_{i=0}^{\infty} \alpha^{iK+l-j}\right) G_l(\beta_{\alpha,l}^*)$$
$$\ge \frac{G_j(\beta_{\alpha,j}^*)}{1-\alpha^K} \ge \frac{G_j(\beta_{\alpha,j}^*)}{K(1-\alpha)},$$

where the last inequality follows from $K - \alpha K \ge 1 - \alpha^K$ because $K \ge \sum_{l=0}^{K-1} \alpha^l = (1-\alpha^K)/(1-\alpha)$. Combining this lower bound for $v_\alpha^*(0,j)$ with the upper bound in Theorem 3(a), we get $[G_j(\beta_{\alpha,j}^*)]/[K(1-\alpha)] \le v_\alpha^*(0,j) \le w_1(0,j) + \bar{g}/(1-\alpha)$ so that for all $\alpha < 1$

$$G_j(\beta^*_{\alpha,j}) \le K(1-\alpha)\left[w_1(0,j) + \frac{\bar{g}}{1-\alpha}\right] \le K w_1(0,j) + K\bar{g}. \quad (11)$$

In view of Assumption 1, this implies that $\{\beta^*_{\alpha,j}\}$ is confined to a finite interval, proving part (b) and hence part (c). ■

6. THE AVERAGE COST CRITERION

In this section, we show that a base-stock policy continues to be optimal under the long-run average cost criterion. In fact, we show that any of the base-stock policies that arise in part (c) of Theorem 3, that is, any base-stock policy that is optimal under the discounted cost criterion for a sequence of discount factors approaching one, is optimal under the average cost criterion as well. To prove these results, we need to verify one of several sets of sufficient conditions for the existence of an optimal stationary policy and the existence of a solution to the average cost optimality equation:

$$h(x,j) = \min_{x \le y \le x+b_j} \left\{ c_j \cdot (y-x) + G_j(y) - g^* + \sum_{\xi=0}^{\infty} \Pr\{D_j = \xi\} \cdot h(y-\xi, j^+) \right\}, \quad (12)$$

where g^* denotes the minimum long-run average cost value.

We verify the existence conditions in Sennott [22]. In addition to the minimum total discounted costs being finite, that is, $v^*_\alpha(x,j) < \infty$ for all $\alpha < 1$ and all $(x,j) \in S$ (which we showed in Theorem 2(b)), the remaining conditions (Assumptions 2 and 3*) in Sennott consist of a characterization of the *relative* discounted cost difference with respect to a given reference state. More precisely, we choose $(0,K)$ as the reference state and derive upper and lower bounds for the values of the cost functions $\{v^*_\alpha(\cdot,\cdot) - v^*_\alpha(0,K)\}$.

PROPOSITION 2: *For all* $(x,j) \in S$ *and all* $0 \le \alpha < 1$

(a) (*see Assumption* 2 *in Sennott* [22]), *there exists a constant* $N \ge 0$ *such that* $v^*_\alpha(x,j) - v^*_\alpha(0,K) \ge -N$,

(b) (*see Assumption* 3* *in Sennott* [22]), *there exists a nonnegative function* $M: S \to \mathbf{R}$ *with* $M(x,j) = O(|x|^{\rho+1})$ *such that* $v^*_\alpha(x,j) - v^*_\alpha(0,K) \le M(x,j)$. *Moreover, for all* $x \le y \le x + b_j$,

$$\sum_{\xi=0}^{\infty} \Pr\{D_j = \xi\} M(y-\xi, j^+) < \infty. \quad (13)$$

PROOF: Fix $\alpha < 1$. Hitherto, all of our results apply regardless of which of the periods is chosen as period 1. Therefore, number the periods such that

$$v^*_\alpha(\beta^*_{\alpha,K}, K) = \min_{(x,j)\in S} \{v^*_\alpha(x,j)\}. \quad (14)$$

(Recall that the minimum is achieved in one of the states $(\beta^*_{\alpha,1},1),\ldots,(\beta^*_{\alpha,K},K)$.) With this numbering, let $(0,K')$ denote the selected reference state, previously referred to as $(0,K)$.

Following the proof of Theorem 3(a) (see Eq. (10)) one verifies that for all $(x,j) \in S$

$$\begin{aligned} v^*_\alpha(x,j) &\le \mathbf{1}\{(x,j) \notin S_K(\beta^*_\alpha)\} \cdot C^{\bar{\pi}(\beta^*_\alpha)}_{x,j}(S_K(\beta^*_\alpha)) + \max_{(y,K)\in S_K(\beta^*_\alpha)} [v^{\bar{\pi}(\beta^*_\alpha)}_\alpha(y,K)] \\ &= \mathbf{1}\{(x,j) \notin S_K(\beta^*_\alpha)\} \cdot C^{\bar{\pi}(\beta^*_\alpha)}_{x,j}(S_K(\beta^*_\alpha)) + \max_{(y,K)\in S_K(\beta^*_\alpha)} [v^*_\alpha(y,K)] \\ &= \mathbf{1}\{(x,j) \notin S_K(\beta^*_\alpha)\} \cdot C^{\bar{\pi}(\beta^*_\alpha)}_{x,j}(S_K(\beta^*_\alpha)) + v^*_\alpha(\beta^*_{\alpha,K},K) + c_K(b_K - 1). \end{aligned} \tag{15}$$

(The first equality follows from the policy $\bar{\pi}(\beta^*_\alpha)$ switching to the policy β^*_α upon entering the set $S_K(\beta^*_\alpha)$ and the definition of β^*_α; this base-stock policy places, in all states in $S_K(\beta^*_\alpha)$, an immediate order to increase the inventory level to $\beta^*_{\alpha,K}$, thus justifying the second equality.)

(a) Apply Eq. (15) to the state $(0,K')$. In view of Proposition 1(b) and Theorem 3(b), we have that $[\mathbf{1}\{(0,K') \notin S_K(\beta^*_\alpha)\} \cdot C^{\bar{\pi}(\beta^*_\alpha)}_{0,K'}(S_K(\beta^*_\alpha)) + c_K(b_K - 1)]$ is uniformly bounded in α by some constant N. Hence, $v^*_\alpha(0,K') \le N + v^*_\alpha(\beta^*_{\alpha,K},K) \le N + v^*_\alpha(x,j)$ for all $(x,j) \in S$ by Eq. (14).

(b) Again, by Eq. (14), $v^*_\alpha(\beta^*_{\alpha,K},K) \le v^*_\alpha(0,K')$. This, together with Eq. (15), implies $v^*_\alpha(x,j) - v^*_\alpha(0,K') \le \mathbf{1}\{(x,j) \notin S_K(\beta^*_\alpha)\} \cdot C^{\bar{\pi}(\beta^*_\alpha)}_{x,j}(S_K(\beta^*_\alpha)) + c_K(b_K - 1) \le M(x,j) \doteq M_1|x|^{\rho+1} + M_2$ for appropriate constants M_1, M_2 (independent of α), again in view of Proposition 1(b) and Theorem 3(b). Finally, Eq. (13) follows immediately from Assumption 3′ (see, e.g., the proof of Theorem 1). ∎

We are ready for the main result for the average cost model. Let β^* denote any base-stock vector such that $\beta^* = \beta^*_{\alpha_n}$ for a sequence of discount factors $\{\alpha_n\}$ converging to 1, which exists in view of Theorem 3(c).

THEOREM 4:

(a) $\lim_{\alpha\to 1}(1-\alpha)v^*_\alpha(x,j) = g^*$ *for all* $(x,j) \in S$.

(b) *Let* $\beta^* = \beta^*_{\alpha_n}$ *for a sequence of discount factors* $\{\alpha_n\}$ *converging to* 1. *The sequence of relative cost functions* $\{v^*_{\alpha_n}(x,j) - v^*_{\alpha_n}(0,K)\}$ *has at least one limit point and each limit point is a function* $h^*(\cdot,\cdot): S \to \mathbf{R}$, *with* $-N \le h^*(x,j) \le M_1|x|^{\rho+1} + M_2$ *for appropriate constants* $M_1, M_2 > 0$, *which satisfies the optimality Eq.* (12); *moreover,* β^* *is average cost optimal and satisfies the optimality equation for any such solution* $h = h^*$.

(c) g^* *is nonincreasing and convex in b.*

PROOF: Theorem 2(b) shows that $v_\alpha^*(x,j) < \infty$ for all $(x,j) \in S$ and all $\alpha < 1$, that is, Assumption 1 in Sennott [22]. Propositions 2(a) and 2(b) verify Assumptions 2 and 3* in Sennott, respectively. Parts (a) and (b) of Theorem 4 now follow from the theorem in Sennott [22]. Part (c) is immediate from part (e) in Theorem 2 and part (a) of this theorem. ∎

7. AN EFFICIENT SOLUTION METHOD FOR THE AVERAGE COST MODEL

In this section, we show that a slight modification of recursive scheme (2), employed with $\alpha = 1$, can be used to identify an optimal base-stock policy β^* for the average cost model. More specifically, we propose a modified value-iteration scheme, generating the functions $\tilde{v}_n$ as follows:

$$\tilde{v}_{n+1}(x,j) = \min_{x \le y \le x+b_j} \left\{ \tau c_j(y-x) + \tau G_j(y) + (1-\tau)\tilde{v}_n(x,j) + \tau \sum_{\xi=0}^{\infty} \Pr\{D_j = \xi\}\tilde{v}_n(y-\xi, j^+) \right\}, \tag{16}$$

with τ an arbitrary constant such that $0 < \tau < 1$. We show in particular that $\{\tilde{v}_n - n\tau g^*\}_{n=1}^{\infty}$ converges to a solution of optimality Eq. (12).

Recursive scheme (16) is a modification of the standard value-iteration method in which the functions $\{v_n^*\}$ are generated via Eq. (2). The modified scheme transforms the transition probability matrix P of any of the Markov chains, induced by a stationary policy, to $\tilde{P} = (1-\tau)I + \tau P$, with I the identity matrix. Note that $\tilde{P}_{(x,j),(x,j)} \ge (1-\tau) > 0$ for all $(x,j) \in S$; that is, all matrices $\tilde{P}$ are aperiodic. Indeed, the *modified* scheme can be viewed as the *standard* value-iteration method applied to the transformed MDP in which all one-step expected costs and one-step transition probabilities are multiplied by $0 < \tau < 1$, with the residual probability mass of $(1-\tau)$ added to the probabilities of transitions from any state to itself. The transformed MDP has the same set of solutions to the optimality equation as the original MDP has, with a long-run average cost value $\tilde{g} = \tau g^*$ (see Schweitzer [21], Federgruen and Schweitzer [7], and Remark 2 in Aviv and Federgruen [2]). (In contrast, the sequence $\{v_n^* - ng^*\}$ derived from standard value-iteration scheme (2) fails to converge because in the original MDP all stationary policies induce a Markov chain that has a periodicity of K or larger.)

The value-iteration method (in its basic form or with the preceding data transformation) is used as a general-purpose solution method for general MDPs with *finite* state and action sets. As shown in Sennott [23] and Aviv and Federgruen [2], the method is also applicable in many MDPs with a *countable* state space in which the existence of a solution to the average cost optimality equation, bounded by an appropriate order function, can be verified.

To prove that value-iteration scheme (16) converges to a solution v^* of optimality Eq. (12), it is in fact necessary to apply a second modification, this time with respect to the set of feasible actions in some of the states. In particular, we *restrict* orders in states with a sufficiently high (low) inventory level to be of zero (full-capacity) size only. This second modification is only needed to prove theoretical convergence. In practice, we have observed that the sequence $\{\tilde{v}_n - n\tau g^*\}$ converges even when the preceding restriction of the feasible action sets is not implemented. More specifically, for $j = 1,\ldots,K$, let

$$u_j = \sup\{x : G_j(x) \le K[C_{0,j}^{\tilde{\pi}(t(0))}(S_K(t(0))) + C_{-b_K,K}^{\tilde{\pi}(t(0))}(S_K(t(0)))]\} + 1,$$

$$l_j = \inf\{x : G_j(x) \le K[C_{0,j}^{\tilde{\pi}(t(0))}(S_K(t(0))) + C_{-b_K,K}^{\tilde{\pi}(t(0))}(S_K(t(0)))]\} - b_j.$$

Recall that $t(0)$ denotes the base-stock vector $(\sum_{i=1}^{K-1} r_i, \sum_{i=2}^{K-1} r_i, \ldots, r_{K-1}, 0)$ (see Eq. (5)) and that $C_{x,j}^{\tilde{\pi}(t(0))}(S_K(t(0)))$, which denotes the expected total cost incurred under this base-stock policy until the first entry into the interval $\{(\xi,K) : -b_K < \xi \le 0\}$ when starting in a state (x,j) outside this interval, is finite by Proposition 1. It follows that $-\infty < l_j < u_j < \infty$ by Assumption 1. We restrict the feasible action set in states (x,j) with $x \ge u_j$ to the singleton $\hat{Y}(x,j) = \{y = x\}$ and in states (x,j) with $x < l_j$ to the singleton $\hat{Y}(x,j) = \{y = x + b_j\}$. In all other states (x,j), we maintain the action set $\hat{Y}(x,j) = \{x \le y \le x + b_j\}$. We refer to the resulting model as the *restricted* model.

This restriction is without loss of optimality.

LEMMA 3: *Let $h^* : S \to \mathbf{R}$ be any function and β^* be any base-stock policy as defined in Theorem* 4. *Then,*

(a) *h^* is a solution of optimality Eq.* (12) *both in the original and in the restricted models,*

(b) *β^* satisfies the optimality equation both in the restricted and in the original model, for the solution $h = h^*$ and is average cost optimal in both models. In particular, any base-stock policy that is average cost optimal in the restricted model is average cost optimal in the original model.*

PROOF: Any base-stock policy β^* that arises in Theorem 4 is the limit of a sequence $\{\beta^*_{\alpha_n}\}$ with $\{\alpha_n\}$ converging to 1, where $\beta^*_{\alpha_n}$ is an optimal base-stock policy in the original model, discounting future costs with a factor α_n. Apply Eq. (11) with $\beta^1 = t(0)$ to obtain for all $j = 1,\ldots,K$,

$$G_j(\beta^*_{\alpha_n,j}) \le K[C_{0,j}^{\tilde{\pi}(t(0))}(S_K(t(0))) + C_{-b_K,K}^{\tilde{\pi}(t(0))}(S_K(t(0)))]. \tag{17}$$

Taking limits as $n \to \infty$ and because G_j is convex and hence continuous (see Assumption 1), we conclude that $G_j(\beta_j^*)$ is bounded by the right-hand side of Eq. (17) as well. Thus, by the definition of l_j and u_j, we have $l_j + b_j \le \beta_j^* < u_j$. This implies that the base-stock policy β^* is feasible in the restricted model and, hence, optimal for that model as well. It remains to be shown that h^* is a solution to the optimality equation of the restricted model and that β^* satisfies this optimality equation for the solution $h = h^*$. Because h^* satisfies Eq. (12) in the original model,

$$h^*(x,j) \le c_j(y-x) + G_j(y) - g^* + \sum_{\xi=0}^{\infty} \Pr\{D_j = \xi\} h^*(y-\xi, j^+); \qquad (x,j) \in S, \quad y \in \hat{Y}(x,j). \tag{18}$$

Because β^* satisfies Eq. (12) in the original model for $h = h^*$, and because $l_j + b_j \le \beta_j^* \le u_j$ for all $j = 1, \ldots, K$, we have

$$h^*(x,j) = G_j(x) - g^* + \sum_{\xi=0}^{\infty} \Pr\{D_j = \xi\} h^*(x-\xi, j^+), \qquad \beta_j^* \le x, \tag{19}$$

$$h^*(x,j) = c_j(\min\{\beta_j^*, x+b_j\} - x) + G_j(\min\{\beta_j^*, x+b_j\}) - g^* + \sum_{\xi=0}^{\infty} \Pr\{D_j = \xi\} h^*(\min\{\beta_j^*, x+b_j\} - \xi, j^+), \quad l_j \le x < \beta_j^*, \tag{20}$$

$$h^*(x,j) = c_j b_j + G_j(x+b_j) - g^* + \sum_{\xi=0}^{\infty} \Pr\{D_j = \xi\} h^*(x+b_j-\xi, j^+), \qquad x \le l_j. \tag{21}$$

Thus, Eqs. (19)-(21) establish that, for all $(x,j) \in S$, Eq. (12) holds as an equality for some $y \in \hat{Y}(x,j)$, namely, the action y prescribed by the policy β^*. This implies that h^* is a solution to optimality Eq. (12) in the restricted model and that β^* satisfies this equation for $h = h^*$. ∎

We conclude in particular that in the restricted model there exists a solution to average cost optimality Eq. (12) such that h^* is bounded by a simple order function that is known in closed form. In particular, $h^*(x,j) \le r(x) \doteq M_1|x|^{\rho+1} + M_2$ for all $(x,j) \in S$. Aviv and Federgruen [2], building on earlier results by Sennott [23], showed that for such MDPs convergence of the value-iteration method can be established by verifying a single, relatively simple additional condition: for any policy π, let (X_n, J_n) denote the sequence of states visited.

Condition (C): $\mathsf{E}[r(X_n) \mid (X_0, J_0) = (x,j)] = O(|x|^{\rho+1})$ for all $(x,j) \in S$ and $n \ge 1$; that is, there exists a constant C such that for all $n \ge 1$ $\mathsf{E}[r(X_n) \mid (X_0, J_0) = (x,j)] \le C(|x|^{\rho+1} + 1)$.

We now verify Condition (C) assuming that all demand distributions $\{D_j\}$ have finite moments of order $\rho + 2$, a slight refinement of Assumptions 3 and 3′.

LEMMA 4: *Assume* $\mathsf{E}[D_j^{\rho+2}] < \infty$ *for all* $j = 1, \ldots, K$. *Condition* (C) *applies.*

PROOF: Let $\bar{u} = \max\{0, u_1, \ldots, u_K\}$; $\underline{l} = \min\{0, l_1, \ldots, l_K\}$ and $\bar{b} = \max_j b_j$. Fix a policy π in the restricted model and a starting state $(X_0, J_0) = (x,j)$. Observe that

$$\min\{X_n, \underline{l}\} \le X_n \le \max\{x, \bar{u} + \bar{b}\} \quad \text{for all } n \ge 1. \tag{22}$$

(Let $n_0 = \min\{n \geq 0 : X_n \leq \bar{u}\}$; X_n is nondecreasing during the first n_0 periods and thereafter the inventory level after ordering is never in excess of $\bar{u} + \bar{b}$, thus verifying the upper bound in Eq. (22).) By the convexity of the function $|x|^{\rho+1}$, we obtain

$$|X_n|^{\rho+1} \leq \max\{(-\min\{X_n, \underline{l}\})^{\rho+1}, (\max\{x, \bar{u} + \bar{b}\})^{\rho+1}\}$$
$$\leq (-\min\{X_n, \underline{l}\})^{\rho+1} + \max\{|x|^{\rho+1}, (\bar{u} + \bar{b})^{\rho+1}\}. \tag{23}$$

Thus, to show that $\mathsf{E}[r(X_n) \mid (X_0, J_0) = (x, j)] = O(r(x))$, it suffices to show that $\mathsf{E}[(-\min\{X_n, \underline{l}\})^{\rho+1} \mid (X_0, J_0) = (x, j)] = O(|x|^{\rho+1})$. To do so, define a nonstationary *reflected random walk* as follows: $R_0 = \min\{x, \underline{l}\}$; $R_{n+1} = \min\{R_n + b_{j_n} - D_{j_n}, \underline{l}\}$. We show by induction that

$$R_n \leq \min\{X_n, \underline{l}\} + \bar{b}. \tag{24}$$

Observe first that $R_0 = \min\{x, \underline{l}\} \leq \min\{X_0, \underline{l}\} + \bar{b}$. Now assume that Eq. (24) holds for some $n \geq 0$. Consider the following two cases:

$$X_n > \underline{l} \geq R_n : R_{n+1} = \min\{R_n + b_{j_n} - D_{j_n}, \underline{l}\} \leq \min\{X_n + b_{j_n} - D_{j_n}, \underline{l}\}$$
$$\leq \min\{X_n - D_{j_n}, \underline{l}\} + \bar{b} \leq \min\{X_{n+1}, \underline{l}\} + \bar{b},$$
$$X_n \leq \underline{l} : R_{n+1} = \min\{R_n + b_{j_n} - D_{j_n}, \underline{l}\}$$
$$\leq \min\{\min\{X_n, \underline{l}\} + \bar{b} + b_{j_n} - D_{j_n}, \underline{l}\}$$
$$= \min\{X_n + b_{j_n} - D_{j_n} + \bar{b}, \underline{l}\} = \min\{X_{n+1} + \bar{b}, \underline{l}\}$$
$$\leq \min\{X_{n+1}, \underline{l}\} + \bar{b},$$

thus concluding that $(-\min\{X_n, \underline{l}\})^{\rho+1} \leq (|R_n| + \bar{b})^{\rho+1}$.

To show that $\mathsf{E}[(-\min\{X_n, \underline{l}\})^{\rho+1} \mid (X_0, J_0) = (x, j)] = O(|x|^{\rho+1})$, it thus suffices to verify that $\mathsf{E}[|R_n|^{\rho+1} \mid (X_0, J_0) = (x, j)] = O(|x|^{\rho+1})$. To do so, we compare the random walk R with one in which reflection occurs only once every K periods, that is, once per cycle, and at the possibly lower level $R_0 = \min\{x, \underline{l}\}$. More specifically, let

$$\tilde{R}_{nK} = \min\{\tilde{R}_{nK-1} + b_{j_{K-1}} - D_{j_{K-1}}, R_0\}, \qquad n \geq 0,$$
$$\tilde{R}_{nK+m} = \tilde{R}_{nK+m-1} + b_{j_{m-1}} - D_{j_{m-1}}, \qquad n \geq 0, m = 1, \ldots, K-1,$$

with $\tilde{R}_0 = R_0 = \min\{x, \underline{l}\}$. Note, however, that the embedded process $\{\tilde{R}_{nK}\}$ is a translation of a *standard* reflective random walk with positive drift and increments distributed as $I \doteq \left(B - \sum_{i=1}^{K} D_i\right)$. Moreover, the random walks R and $\tilde{R}$ "move" in "close" proximity. More specifically, we shall show that

$$R_{nK+m} \geq \tilde{R}_{(n+1)K} - 2B \quad \text{(a.s.) for all } n \geq 0 \text{ and } m = 1, \ldots, K-1, \tag{25}$$

which implies that $\mathsf{E}[|R_{nK+m}|^{\rho+1}] \leq \mathsf{E}[(|\tilde{R}_{(n+1)K}| + 2B)^{\rho+1}]$. Now, for all $n \geq 1$, $\tilde{R}_{nK} =_d \min\{x, \underline{l}\} + \min\{S_n, S_n - S_1, S_n - S_2, \ldots, S_n - S_{n-1}, 0\}$, where S_n denotes the sum of n independent increments all distributed as I (see Asmussen

[1, Proposition 1.1 on p. 181]). The sequence $\{\tilde{R}_{nK}\}$ is thus stochastically decreasing with a limiting distribution equal to that of $[\min\{x,\underline{l}\} + R^*]$, where $\mathsf{E}[|R^*|^{\rho+1}] < \infty$ because $\mathsf{E}[|I|^{\rho+2}] < \infty$ (see Asmussen [1, Theorem 2.1 on p. 184]). We conclude that a constant C exists such that for all $n \geq 1$ $\mathsf{E}[|\tilde{R}_{nK}|^{\rho+1}] \leq C(|x|^{\rho+1} + 1)$.

It remains to be shown that Eq. (25) holds. By the construction of the process $\tilde{R}$, it is clear that, for all $1 \leq m < K$, $\tilde{R}_{(n+1)K} = \min\{\tilde{R}_{nK+m} + \sum_{i=m+1}^{K} (b_{j_i} - D_{j_i}),\ R_0\} \leq \tilde{R}_{nK+m} + B$. Hence, to verify Eq. (25) it suffices to show that $\tilde{R}_{nK+m} \leq R_{nK+m} + B$ (a.s.) for all $n \geq 0$ and $m = 0, \ldots, K-1$. The proof is by induction.

For $n, m = 0$ the inequality holds because $\tilde{R}_0 = R_0$; assume now that it is true for some $n \geq 0$ and $0 \leq m \leq K-1$, and note that

$$\begin{aligned}
\tilde{R}_{nK+m+1} - R_{nK+m+1} &\leq \tilde{R}_{nK+m} + b_{j_m} - D_{j_m} - \min\{R_{nK+m} + b_{j_m} - D_{j_m}, \underline{l}\} \\
&\leq \max\{\tilde{R}_{nK+m} - R_{nK+m}, \tilde{R}_{nK+m} + b_{j_m} - \underline{l}\} \\
&\leq \max\left\{B, \sum_{i=0}^{m} b_{j_i}\right\} = B. \qquad \blacksquare
\end{aligned}$$

Thus, let $\{\hat{v}_n\}$ denote the sequence of functions generated by recursive scheme (16) in which the action set $Y(x,j) = [x, x + b_j]$ is restricted to $\hat{Y}(x,j)$ for all $(x,j) \in S$.

THEOREM 5: $\{\hat{v}_n - n\tau g^* \mathbf{1}\}$ *converges to a solution* h^* *of average cost optimality Eq.* (12), *with* $h^*(x,j) = O(|x|^{\rho+1})$.

PROOF: Lemma 4 establishes Condition (C). Also, the data transformation in Eq. (16) ensures that the Markov chain induced by each of the stationary policies is aperiodic (as already explained). The theorem now follows from Theorem 1 in Aviv and Federgruen [2]. ∎

We now describe how the sequence $\{\hat{v}_n\}$ can be generated in an efficient manner. First, to apply the restriction of the action sets to the sets $\{\hat{Y}(x,j) : (x,j) \in S\}$ we need to compute the values $\{u_j, l_j; j = 1, \ldots, K\}$. This can be achieved by the following value-iteration scheme: let $z(x,j) = C_{x,j}^{\bar{\pi}(t(0))}(S_K(t(0)))$. Then, $z = \lim_{n\to\infty} z_n$ by Corollary 9.17.1 in Bertsekas and Shreve [3], where for all (x,j) with $x \leq \sum_{l=j}^{K-1} r_l = t(0)_j$

$$z_n(x,j) = c_j[y(x,j) - x] + G_j(y(x,j)) + \sum_{\xi=0}^{\infty} \Pr\{D_j = \xi\} z_{n-1}(y(x,j) - \xi, j^+), \qquad j \neq K-1, \tag{26}$$

$$z_n(x,K-1) = c_{K-1}[y(x,K-1) - x] + G_{K-1}(y(x,K-1)) + \sum_{\xi=y(x,K-1)+b_K}^{\infty} \Pr\{D_{K-1} = \xi\} z_{n-1}(y(x,K-1) - \xi, K), \tag{27}$$

with $y(x,j) = \min\{x + b_j, \sum_{l=j}^{K-1} r_l\}$ and $z_0 \equiv 0$. (Note that the set $S_K(t(0))$ can only be reached from states (x,j) with $j = K - 1$. Also, by a slight modification of the proof of Lemma 2(a), one easily verifies that the set $\bar{S}(t(0))$ is never reached, when starting outside this set. This explains why the summation in Eq. (27) should be started at $\xi = y(x,K-1) + b_K$.)

Also, instead of the functions $\hat{v}_n$, which grow linearly in n, one should generate the normalized value functions w_n defined by $w_n(x,j) \doteq \hat{v}_n(x,j) - \hat{v}_n(x_0,j_0)$ for some reference state $(x_0,j_0) \in S$ (e.g., $(x_0,j_0) = (0,K)$ as in Section 6). Note that $\{w_n\}$ can be generated from the recursion

$$\begin{aligned} w_{n+1}(x,j) = \min_{y \in \hat{Y}(x,j)} \Big\{ &\tau c_j(y-x) + \tau G_j(y) + (1-\tau) w_n(x,j) \\ &+ \tau \sum_{\xi=0}^{\infty} \Pr\{D_j = \xi\} w_n(y-\xi, j^+) \Big\} \\ - \min_{y \in \hat{Y}(0,K)} \Big\{ &\tau c_K y + \tau G_K(y) + (1-\tau) w_n(0,K) \\ &+ \tau \sum_{\xi=0}^{\infty} \Pr\{D_K = \xi\} w_n(y-\xi, 1) \Big\}, \end{aligned} \tag{28}$$

that these functions remain bounded in n in the course of the algorithm and that the same sequence of base-stock policies is generated by Eq. (28) as by the original scheme. Finally, $\{w_n\}$ converges to a solution of optimality Eq. (2) and $\lim_{n\to\infty}[w_n - w_{n-1}] = \lim_{n\to\infty}[\hat{v}_n - \hat{v}_{n-1}] = \tau g^* \mathbf{1}$.

Remark: We have observed that, in practice, it is neither necessary to restrict the action sets upfront to the sets $\{\hat{Y}(x,j)\}$ nor to apply data transformation (16). In other words, the sequence $\{v_n^* - ng^*\mathbf{1}\}$ itself converges where v_n^* is defined in Eq. (2) with $\alpha = 1$. Because in practice convergence occurs significantly fast, it is not necessary to generate the normalized functions w_n to avoid numerical instability.

The computational effort associated with the value-iteration scheme, low as it is for general MDPs, can be reduced significantly by exploiting all of the properties stated in Theorem 1, in particular, the convexity of the functions $J_n(y,j)$ in y and the fact that their minimizers $\{y_{n,j}^* : j = 1, \ldots, K\}$ can be used as optimal base-stock levels in the nth iteration of the algorithm. Thus, with v_{n-1}^* known at the beginning of the nth iteration, it suffices to compute the values $y_{n,j}^*$—for example, via a simple *bisection method*. Thereafter, one evaluates $v_n^*(x,j)$ for all required inventory levels x, via $v_n^*(x,j) = -c_j x + J_n(\min\{y_{n,j}^*, x + b_j\}, j)$, thus completing the nth iteration.

Finally, as for all numerical methods used to solve infinite state space models, it is of course necessary to *truncate* the state space at certain sufficiently high and low inventory levels U_j and L_j, $j = 1, \ldots, K$. A basic implementation would simply truncate the summation in Eq. (16) at the level $\xi = y - L_j$. This can, however, result in significant evaluation errors. Instead, we propose replacing the uncomputed values $\tilde{v}_n(x,j^+)$ for $x \le L_j$ (i.e., for states outside the truncated state space) by an appropriate *extrapolation* of the values $\{\tilde{v}_n(x,j)\}$ inside this truncated space. In Theorem 5, we have shown that the limit function of the sequence $\{\hat{v}_n(x,j) - n\tau g^* \mathbf{1}\}$ is $O(|x|^{\rho+1})$. It is therefore appropriate to extrapolate the functions $\hat{v}_n$ by a polynomial of degree $\rho + 1$; for example, when holding and backlogging costs are linear functions, $\rho = 1$ and extrapolation by a quadratic function is called for.

8. NUMERICAL STUDY

In this section, we report on the numerical study described in Section 1. We have evaluated a total of 57 different problem instances, all with $K = 6$ periods, linear holding and backlogging costs, and zero leadtimes (see Eq. (1)). In addition, all variable production cost rates $\{c_j\}$, holding cost rates $\{h_j\}$, and backlogging cost rates $\{p_j\}$ are taken to be stationary, that is, $c_j = c$, $h_j = h$, $p_j = p$ for all $j = 1, \ldots, K$.

First, in terms of computational times, we have observed that the sequence of optimal base-stock levels $\{y_n^*\}$ converges quite rapidly; that is, the same base-stock policy is generated after at most 20 iterations, or several scores of CPU seconds on a 486-based PC. Obtaining convergence of the sequence $\{\hat{v}_n - n\tau g^* \mathbf{1}\}$ within a precision of 0.1% can require up to several minutes of CPU time on 486-based PCs; the computational times increase with the utilization rate. We have validated the accuracy of the estimate of g^*, obtained via the sequence $\{\hat{v}_n - \hat{v}_{n-1}\}$, with that achieved via high precision simulations.

In our basic set of problem instances, the vector of mean demands μ is given by $\mu = (30, 35, 50, 60, 40, 25)$; this represents a pattern with a single mode and with sales within the peak month 2.4 times the volume in the slowest month. Within this set, we have evaluated all 35 combinations of seven capacity levels, assumed to be stationary, that is, $b_1 = b_l = \cdots = b_K = b$, ranging from 45 to 100 (corresponding with a utilization rate ranging from 40 to 88.89%) and five types of demand distribution: a Geometric and Negative Binomial distribution representing the number of "failures" in a series of Bernoulli trials until the first and fifth "success," respectively; a Poisson distribution and two compound Poisson distributions where the compounding order size has a Binomial distribution with mean 1 and $n = 2$ and $n = 5$ trials, respectively. Table 1 reports the maximum and minimum coefficient of variation of one-period demands over all period types, as well as the optimal base-stock levels $\beta_1^*, \ldots, \beta_K^*$.

We conclude that in all periods optimal base-stock levels are nonincreasing in the capacity of the system. The sensitivity of the optimal base-stock levels with respect to the system's capacity is relatively small when the demand variability is relatively low (e.g., the instances with Poisson and compound Poisson

Table 1. Optimal Base-Stock Levels: $K = 6$, $\mu = (30, 35, 50, 60, 40, 25)$, $h = 0.5$, and $p = 1.0$

Distribution Type	b	Utilization Rate	Optimal Base-Stock Levels 1	2	3	4	5	6
Geometric	45	88.89%	138	162	181	164	124	109
min CV = 0.98, max CV = 0.99	50	80.00%	91	116	139	124	84	63
	60	66.67%	54	75	102	92	57	37
	70	57.14%	42	58	86	80	49	31
	80	50.00%	37	50	76	74	46	29
	90	44.44%	35	45	70	70	44	28
	100	40.00%	34	43	66	68	43	27
Negative Binomial ($r = 5$)	45	88.89%	70	89	104	96	59	47
min CV = 0.44, max CV = 0.44	50	80.00%	45	63	82	80	48	31
	60	66.67%	35	47	69	73	45	28
	70	57.14%	34	42	63	70	45	28
	80	50.00%	33	40	60	69	44	28
	90	44.44%	33	39	58	68	44	28
	100	40.00%	33	39	57	68	44	28
Compound Poisson [Bin($n = 5, \theta = 2$)]	45	88.89%	40	55	68	66	43	28
min CV = 0.17, max CV = 0.27	50	80.00%	34	46	62	65	43	27
	60	66.67%	33	39	57	64	43	27
	70	57.14%	33	38	54	64	43	27
	80	50.00%	33	38	54	64	43	27
	90	44.44%	33	38	54	64	43	27
	100	40.00%	33	38	54	64	43	27
Compound Poisson [Bin($n = 2, \theta = 0.5$)]	45	88.89%	39	54	67	65	43	28
min CV = 0.16, max CV = 0.24	50	80.00%	33	45	62	64	43	27
	60	66.67%	33	39	56	64	43	27
	70	57.14%	33	38	54	64	43	27
	80	50.00%	33	38	53	64	43	27
	90	44.44%	33	38	53	64	43	27
	100	40.00%	33	38	53	64	43	27
Poisson	45	88.89%	37	52	66	64	42	27
min CV = 0.13, max CV = 0.20	50	80.00%	32	44	60	63	42	27
	60	66.67%	32	38	55	63	42	27
	70	57.14%	32	37	53	63	42	27
	80	50.00%	32	37	53	63	42	27
	90	44.44%	32	37	53	63	42	27
	100	40.00%	32	37	53	63	42	27

Abbreviations: min CV = minimum coefficient of variation; max CV = maximum coefficient of variation.

distributions). On the other hand, for instances with a relatively high demand variability, the optimal base-stock levels may be reduced by a factor of 3–4 when the capacity is changed from $b = 100$ to $b = 45$. The results also suggest that the optimal base-stock levels vary convexly with the system's capacity b. Similarly, we observe that the base-stock levels increase as the demand variability increases; as can be explained, the increase is particularly large when the utilization rate is high. These observations indicate the importance of appropriate and accurate forecasting systems, in particular under high utilization rates.

Considering the base-stock levels for a given instance, one observes that the periods with the largest and most variable demands (i.e., those with the highest standard deviations) may fail to have the largest base-stock level, even though the production lead time is zero. This occurs because of the need to build up inventories in advance of peak demand periods. Consider, for example, the instances with Negative Binomial demand distributions. Under low-capacity volumes (e.g., $b = 45$), the base-stock level for period 3 is higher than that of the peak period, period 4; when the capacity b is increased to $b = 60$ the relative *order* of the base-stock levels is reversed, as the safety stock requirements to cover the larger and more variable demand in the peak period now dominate over the need to build up inventories in response to the capacity constraints.

Figure 1 displays the long-run average system-wide costs as a function of the capacity level for the five demand distributions considered. As proven in

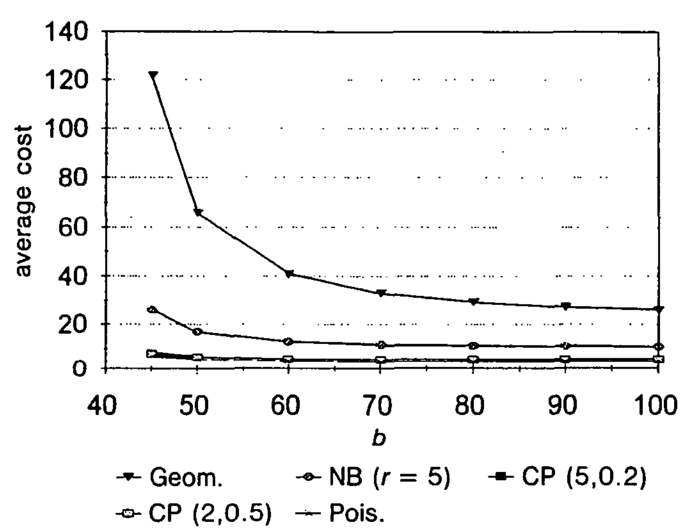

FIGURE 1. Long-run average cost as a function of the capacity limits, for various demand distributions: $h = 0.5$, $p = 1.0$, $c = 0$. Geom. = Geometric; NB = Negative Binomial; CP = Compound Poisson; Pois. = Poisson.

Theorem 4(c), the average cost value, g^*, is nonincreasing and convex in b. Observe that the cost increases dramatically as the demand variability increases, in particular under high utilization rates. These observations suggest that major benefits can be reaped by reducing the variability of demands. This can be achieved, for example, by improved coordination with customers, receiving advanced notice of future orders as in many *quick response* or *vendor-managed replenishment* programs. Figures 2 and 3 display the average inventory and average backlog sizes as a function of the system's capacity, again in five different curves, corresponding with the five types of demand distributions considered. The former represent capacity/inventory investment tradeoff curves. The observations made with respect to Figure 1 apply here as well. Notice in particular the large inventory reductions that can be achieved by capacity expansions, in particular when demands are highly variable.

In our second set of instances, we investigate the impact of the ratio h/p, that is, the relative magnitude of holding versus backlogging costs. Starting with the four instances in the first set, in which the capacity $b = 50$ and $b = 100$, and the demand distributions Negative Binomial and Poisson, we systematically considered three alternative values of h, in particular $h = 1$, 0.2, and 0.1. Table 2 displays the optimal base-stock levels for the new set of 24 instances. As can be expected, the base-stock levels increase as the ratio h/p decreases. We observe that the base-stock levels decrease convexly with h, throughout. As already mentioned, the optimal base-stock level for period 3 is often larger than that of

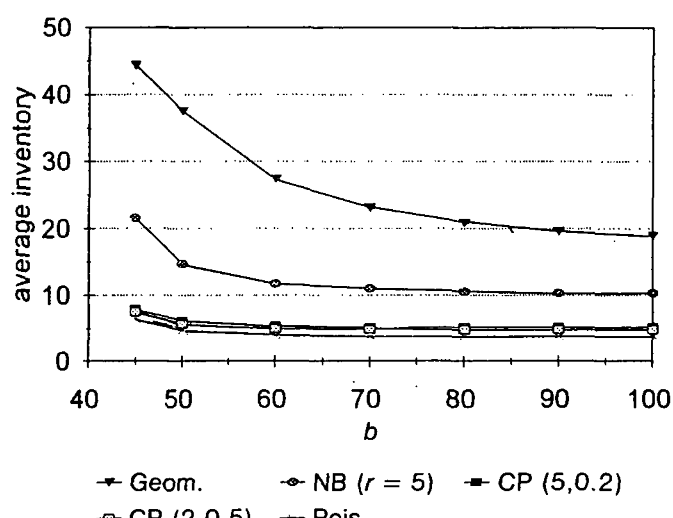

Figure 2. Average inventory on hand as a function of the capacity limits, for various demand distributions: $h = 0.5$, $p = 1.0$, $c = 0$. See caption of Figure 1 for abbreviations.

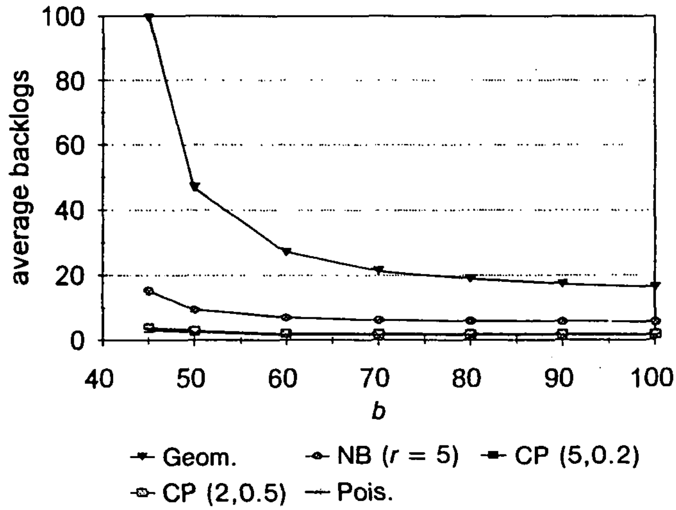

FIGURE 3. Average backlog occurrences as a function of the capacity limits, for various demand distributions: $h = 0.5$, $p = 1.0$, $c = 0$. See caption of Figure 1 for abbreviations.

TABLE 2. Optimal Base-Stock Levels for Different h/p Ratios: $\mu = (30, 35, 50, 60, 40, 25)$, $p = 1.0$, and $h = 0.1, 0.2, 0.5, 1$

			Optimal Base-Stock Levels					
Distribution Type	b	h	1	2	3	4	5	6
Negative Binomial ($r = 5$)	50	1.0	33	47	66	66	39	24
		0.5	45	63	82	80	48	31
		0.2	66	88	107	100	62	44
		0.1	85	107	125	116	74	59
Poisson	50	1.0	30	39	56	60	40	25
		0.5	32	44	60	63	42	27
		0.2	36	51	67	68	46	30
		0.1	39	57	72	71	48	32
Negative Binomial ($r = 5$)	100	1.0	27	32	46	55	36	22
		0.5	33	39	57	68	44	28
		0.2	42	50	72	85	56	35
		0.1	49	58	84	97	65	41
Poisson	100	1.0	30	35	50	60	40	25
		0.5	32	37	53	63	42	27
		0.2	35	41	57	67	46	30
		0.1	37	43	59	70	48	32

period 4 in spite of the latter's demand distribution having a higher mean and standard deviation. The relative ranking $\beta_3^* > \beta_4^*$ occurs primarily when h is small (relative to p) as the incentive to respond to capacity constraints by building inventories in advance, decreases with h. Notice that $\beta_3^* > \beta_4^*$ only occurs when the capacity is relatively low ($b = 50$); in case the demand distributions are Poisson, $\beta_3^* > \beta_4^*$ when $h = 0.1$, but $\beta_3^* \leq \beta_4^*$ for all $h \geq 0.2$.

We conclude this numerical study with an assessment of the benefits that can be achieved when synchronizing the capacity and the demand profile. As already mentioned, this can be done in two ways: in Table 3 we investigate the benefits of demand smoothing, starting from our base vector $\mu = (30, 35, 50, 60, 40, 25)$ and replacing it by a pattern of *constant* mean demands, with the same aggregate $\mu = 240$. In doing so, we maintain the variances of the individual periods' demands ($\sigma^2 = (40, 40, 50, 60, 40, 40)$). As already noted, under seasonal demand fluctuations the base-stock levels as well as the long-run average cost decrease considerably as capacity is expanded; on the other hand, the benefits of capacity expansion are significantly smaller when the demand pattern is smooth. Alternatively, one observes that in the basic instance system-wide costs can be reduced by 43.6% when expanding capacity from $b = 45$ to $b = 80$. However, most of these savings (i.e., 34.6%) can be achieved by introducing a smooth demand pattern, and the latter may in certain settings be implemented with a far smaller investment.

The alternative mechanism to synchronizing the demand and capacity patterns is to adopt flexible, that is, period-dependent, capacity levels. As shown in Theorem 4(c), the long-run average cost is convex in the vector b. Given an aggregate budget constraint with total capacity B used in a cycle, an *optimal* capacity allocated can be determined by solving the convex program

TABLE 3. Optimal Base-Stock Levels and Optimal Long-Run Average Costs, under Fluctuating and Constant Mean-Demand Patterns, for Different Capacity Limits (b): $\sigma^2 = (40, 40, 50, 60, 40, 40)$, $p = 1.0$, and $h = 0.5$; demands follow compound Poisson distributions

		Optimal Base-Stock Levels						
Pattern	b	1	2	3	4	5	6	g^*
Fluctuating demand	45	38	52	66	64	42	28	6.54
	50	33	44	60	63	42	27	4.90
	80	32	37	53	63	42	27	3.69
Smooth demand (constant means over time)	45	44	44	45	45	44	44	4.28
	50	43	43	43	43	43	43	3.76
	80	42	42	43	43	42	42	3.69

$$(\mathrm{P})\min g^*(b_1,\ldots,b_K)$$
$$\text{s.t.} \sum_{j=1}^{K} b_j = B$$
$$b_j \geq 0.$$

We postpone a discussion of effective solution methods for (P) to a future publication. We have experimented with capacity allocations of the type $b_j = \gamma(\mu_j + kr_j)$ for appropriate factors γ and k. Allocations of this type appear to result in modest cost reductions; for example, in the instance of our basic set (see Table 1) in which $b = 50$ and demands have Negative Binomial distributions, we have observed cost savings of approximately 5% by allocating capacities proportional to the periods' mean demands. It remains an open question whether or not and under what circumstances significantly larger cost savings can be achieved under an optimal capacity allocation.

References

1. Asmussen, S. (1987). *Applied Probability and Queues*. New York: Wiley.
2. Aviv, Y. & Federgruen, A. (1995). The value-iteration method for countable state Markov decision processes. Working Paper, Graduate School of Business, Columbia University, New York.
3. Bertsekas, D.P. & Shreve, S.E. (1978). *Stochastic Optimal Control: The Discrete Time Case*. Englewood Cliffs, NJ: Prentice-Hall.
4. Bush, C.M. & Cooper, W.D. (1988). Inventory level decision support. *Production and Inventory Management Journal* First Quarter: 69–73.
5. Denardo, E.V. (1967). Contraction mapping in the theory underlying dynamic programming. *SIAM Review* 9: 165–177.
6. Fair, R.C. (1989). The production-smoothing model is alive and well. *Journal of Monetary Economics* 24: 353–370.
7. Federgruen, A. & Schweitzer, P.J. (1979). Discounted and undiscounted value-iteration in Markov decision problems: A survey. In M.L. Puterman (ed.), *Dynamic Programming and Its Applications*. New York: Academic Press, pp. 23–52.
8. Federgruen, A. & Zipkin, P. (1986). An inventory model with limited production capacity and uncertain demands, I: The average cost criterion. *Mathematics of Operations Research* 11: 193–207.
9. Federgruen, A. & Zipkin, P. (1986). An inventory model with limited production capacity and uncertain demands, II: The discounted cost criterion. *Mathematics of Operations Research* 11: 208–215.
10. Glasserman, P. (1996). Bounds and asymptotics for planning critical safety stocks. *Operations Research* (to appear).
11. Glasserman, P. & Tayur, S. (1995). Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Science* 42: 263–281.
12. Janson, S. (1986). Moments for first-passage and last-exit times, the minimum, and related quantities for random walks with positive drift. *Advances in Applied Probability* 18: 865–879.
13. Kapuscinski, R. & Tayur, S. (1995). A capacitated production-inventory model with periodic demand. Working Paper, GSIA, Carnegie Mellon University, Pittsburgh.
14. Karlin, S. (1960). Dynamic inventory policy with varying stochastic demands. *Management Science* 6: 231–258.
15. Karlin, S. (1960). Optimal policy for dynamic inventory process with stochastic demands subject to seasonal variations. *Journal of SIAM* 8: 611–629.

16. Krane, S.D. & Braun, S.N. (1991). Production smoothing evidence from physical-product data. *Journal of Political Economy* 99: 558–581.
17. Lippman, S. (1975). On dynamic programming with unbounded rewards. *Management Science* 21: 1225–1233.
18. Metters, R. (1994). Production of multiple products with stochastic seasonal demand and capacitated production. Working Paper, Owen Graduate School of Management, Vanderbilt University, Nashville, TN.
19. Miller, B.L. & Veinott, A.F., Jr. (1969). Discrete dynamic programming with a small interest rate. *Annals of Mathematical Statistics* 40: 366–370.
20. Puterman, M.L. (ed.). (1978). *Dynamic Programming and Its Applications.* New York: Academic Press.
21. Schweitzer, P.J. (1971). Iterative solution of the functional equations for undiscounted Markov renewal programming. *Journal of Math Analysis and Applications* 34: 495–501.
22. Sennott, L.I. (1989). Average cost optimal stationary policies in infinite state Markov decision processes with unbounded costs. *Operations Research* 37: 626–633.
23. Sennott, L.I. (1991). Value iteration in countable state average cost Markov decision processes with unbounded costs. *Annals of Operations Research* 28: 261–272.
24. Sethi, S.P. & Cheng, F. (1993). Optimality of (s, S) policies in inventory models with Markovian demand. Working Paper, Faculty of Management, University of Toronto, Toronto.
25. Song, J.S. & Zipkin, P. (1993). Inventory control in a fluctuating demand environment. *Operations Research* 41: 351–370.
26. Tayur, S. (1992). Computing the optimal policy for capacitated inventory models. *Communications in Statistics—Stochastic Models* 9: 585–598.
27. Zipkin, P. (1989). Critical number policies for inventory model with periodic data. *Management Science* 35: 71–80.