OPERATIONS RESEARCH
Vol. 52, No. 5, September–October 2004, pp. 739–755
ISSN 0030-364X | EISSN 1526-5463 | 04 | 5205 | 0739

DOI 10.1287/opre.1040.0131

Optimal Policies for a Capacitated Two-Echelon Inventory System

Rodney P. Parker
Yale School of Management, New Haven, Connecticut 06520-8200, rodney.parker@yale.edu

Roman Kapuscinski
University of Michigan Business School, 701 Tappan Street, Ann Arbor, Michigan 48109-1234, roman.kapuscinski@umich.edu

This paper demonstrates optimal policies for capacitated serial multiechelon production/inventory systems. Extending the Clark and Scarf (1960) model to include installations with production capacity limits, we demonstrate that a modified echelon base-stock policy is optimal in a two-stage system when there is a smaller capacity at the downstream facility. This is shown by decomposing the dynamic programming value function into value functions dependent upon individual echelon stock variables. We show that the optimal structure holds for both stationary and nonstationary stochastic customer demand. Finite-horizon and infinite-horizon results are included under discounted-cost and average-cost criteria.

Subject classifications: inventory/production: multiechelon, stochastic uncertainty, policies, capacity expansion.
Area of review: Manufacturing, Service, and Supply Chain Operations.
History: Received May 2001; revisions received June 2003, September 2003; accepted September 2003.

1. Introduction

The production, storage, and delivery of goods between factories, warehouses, and retailers is a rich area of study with many interesting questions still unsolved. These are issues of substantial importance, because the coordination of goods between members of supply chains constitutes a significant investment in terms of managerial attention, inventory costs, and capital investments. Careful analysis of these problems can increase responsiveness to end-customers' needs without necessarily increasing costs. In this paper, we analyze the basic model that includes two critical elements: multiple stages (two echelons) and production capacity constraints at both stages.

The issue of optimal ordering and inventory policies in multiechelon production and inventory systems without any capacity constraints was the focus of Clark and Scarf (1960). They consider a purely serial supply chain (known as a multiechelon system), with the lowest installation facing stochastic demand from the end customer. In a finite-horizon setting, Clark and Scarf (1960) determine that the optimal ordering policy for the entire multiechelon system can be decomposed into decisions based solely on echelon inventories. Federgruen and Zipkin (1984) extend the multiechelon result to an infinite-time horizon. Rosling (1989) demonstrates that a pure assembly system can be reduced to a serial multiechelon system. The Clark and Scarf (1960) results have been reproven by Chen and Zheng (1994) using lower bounds on the long-run costs, and by Muharremoglu and Tsitsiklis (2003) using an alternative approach based on item-customer decomposition.

The research for systems with limited production capacity (defined as a finite upper limit on the amount that may be processed in a single period) under periodic review is mostly constrained to one-echelon systems. Federgruen and Zipkin (1986a, b) consider a capacitated single installation for infinite-horizon average-cost and discounted-cost criteria, respectively. Following Zipkin's (1989) analysis of the cyclical single-stage system without capacity limits, Kapuscinski and Tayur (1998) and Aviv and Federgruen (1997) study a single installation with limited capacity facing stochastic cyclical demand and find the optimal policy for such systems; Metters (1997) applies heuristics to the same problem with lost sales.

There has been very little research on a multiechelon system with limited capacity at each installation. Speck and Van der Wal (1991a) consider a two-echelon system, present a counterexample justifying why a modified base-stock policy is not optimal, and suggest that numerically "nearby" modified base-stock policies provide a good approximation to the optimal policy. Speck and Van der Wal (1991b) provide an algorithm to determine parameter values for such a base-stock policy. Glasserman and Tayur (1994, 1995) consider this problem, but show analysis for the capacitated supply chain *assuming* that the system operates under a base-stock policy. Roundy and Muckstadt (2000) also assume base-stock policy (for one stage) and propose an efficient approximation.

This paper demonstrates that when the smallest capacity is at the downstream facility in a capacitated serial supply chain, the optimal policy is a simple modification of an echelon base-stock policy. The policy for lower echelon is unchanged—we order up to a specific target (subject to the availability from the higher echelon). The policy for the higher echelon is modified. The higher echelon orders up to a specific echelon target, taking care not to exceed specific installation (on-hand) inventory. Due to the additional constraint on the installation inventory, we label the policy as the modified echelon base-stock (MEBS) policy. This policy can be interpreted as having a structure of generalized kanban production-inventory policies. These are policies with release mechanisms that could encompass constraints on the amounts of inventory in a subset of consecutive installations. Axsäter and Rosling (1999) rank various policies, including generalized kanban policies, in multistage systems with various release mechanisms, according to how general they are. Dallery and Liberopoulos (2000) and Liberopoulos and Dallery (2000) also consider generalized kanban control policies and contain references to other papers in this area. Even though most of this literature deals with Poisson arrivals of demand and exponentially distributed service times in continuous time, some of the observations can be related to our findings. Veatch and Wein (1994) numerically show that kanban policies with unit-sized containers are generally superior to order-up-to policies when the lower installation's capacity is smaller, and this policy relationship is reversed when the capacity condition is reversed. For the case when lower installation has lower or equal capacity in a periodic setting, the policy we show to be optimal can be interpreted as generalized kanban. Although not exactly the same, it is similar to the one considered in Veatch and Wein (1994), which justifies the good performance of the Veatch and Wein heuristics.

Our main result—providing the structure of the optimal policy for a two-echelon system with the constraining capacity closer to the customer—is based on a few observations. We show that nondominated ordering is limited to a "band" of states (formally defined in Definition 1) and that for the states in the "band," the cost function is separable. Specifically, we demonstrate that it will never be optimal for the higher installation (farther from the customer) to hold more inventory than can be processed in a single period by the lowest stage, which is a bottleneck. This immediately implies that the conventional (Clark and Scarf 1960) echelon base-stock policy cannot be optimal. According to the conventional policy, a huge spike of demand would generate the same-size order at the higher installation, which may exceed the capacity constraint, thus generating unnecessary holding costs. Using this "band" dominance, in a two-stage system we substitute the constraints upon production by a constraint on the inventory, and show separability of the cost function. The resulting MEBS policy resembles the conventional multiechelon policy, except for the additional "band" constraint.

We permit general (multiple of period length) lead times leading to the lower installation, but restrict the lead time at the higher installation to one period. Without this limitation, we would not be able to guarantee that the inventory remains in the undominated "band." The limitation to two stages is directly linked to the lead-time issue; in fact, Glasserman and Tayur (1994) describe a technique of inserting dummy installations to act as a surrogate for positive lead times. The final limitation of our model is the requirement that the constraining capacity is at the lower installation. We address this issue with an example in §4.3.4. We show that our results extend to other system configurations in §4.3.3.

This paper is organized as follows. In §2, we describe the model in detail, stating the sequence of events and formulating the model. Section 3 illustrates the deceptive nature of the numerical results where some behavior that is apparently not base-stock in nature is optimal. Section 4 contains the paper's key results for the finite horizon. In this section we prove the optimal policy and extend the model to specific lead-time models and to a model with nonstationary stochastic demands. Section 5 extends the results to discounted-cost and average-cost criteria for the infinite horizon. Concluding remarks are in §6.

2. Model

Consider a serial multiechelon supply chain with N installations (see Figure 1). Each of the installations, $j > 1$, supplies its immediately lower installation, $j-1$, in the supply chain and receives goods from installation $j+1$. Installation N receives goods from an outside uncapacitated supplier. Each installation j, other than N, is limited in its order by the available inventory held at installation $j+1$. Additionally, at installation j there is a capacity limit K_j that serves as an upper limit on the amount that can be processed in each period. Installation 1 supplies goods to the end customer, whose demand, D_n, is stochastic and independent from period to period. $D_n \geqslant 0$ and $\mu_n = \mathbb{E}D_n < \infty$.

The sequence of events is as follows: (1) at the beginning of every period, installation 1 places an order with installation 2, installation 2 places an order with installation 3, and so forth up to installation N; (2) installation 2 then delivers the ordered amount to installation 1, installation 3 delivers to installation 2, and so forth, until the outside supplier delivers to installation N; and (3) end-customer demand is then realized, and installation 1 attempts to satisfy this demand as closely as possible using its available stock. The sequence of orders in Step (1) is listed in this way to demonstrate the dynamics of the system, although these ordering decisions are taken concurrently.[1] Unsatisfied demand is backlogged with a linear penalty cost per period, p. Echelon j incurs an incremental holding cost, $h_j \geqslant 0$, for each unit of inventory in each period, and

Figure 1. Multiechelon system with N installations.

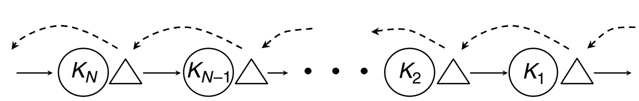

the installation holding cost is $H_j = \sum_{k \geqslant j} h_k$. The objective is to minimize the discounted sum of holding and penalty costs. In the finite-horizon case, time is counted backwards and n represents the number of periods remaining until the end of the horizon. Initially, it is assumed that there are no delivery lead times, other than delays caused by the dynamics of the system. In addition, we assume without loss of generality that the shipping costs are zero.

The variables considered here are: The inventory at installation j at the start of period n is x_n^j; the amount ordered by installation j in period n is $a_n^j \geqslant 0$; and y_n^j is inventory in installation j after shipments are made in period n. The inventory dynamics are described by

$$x_{n-1}^1 = y_n^1 - D_n = x_n^1 + a_n^1 - D_n,$$
$$x_{n-1}^j = y_n^j = x_n^j - a_n^{j-1} + a_n^j \quad \text{for } j > 1.$$

An echelon inventory is defined as the amount in transit to and in stock at an installation plus the amount in transit to and in stock at all lower installations. The corresponding echelon variables $\widetilde{X}_n = (X_n^j)_{j=1}^N \in \Re^N$ and $\widetilde{Y}_n = (Y_n^j)_{j=1}^N \in \Re^N$ are defined as $X_n^j = \sum_{i=1}^j x_n^i$ and $Y_n^j = \sum_{i=1}^j y_n^i$. Clearly,

$$Y_n^j = X_n^j + a_n^j,$$
$$x_{n-1}^j = X_{n-1}^j - X_{n-1}^{j-1} = Y_n^j - Y_n^{j-1}.$$

The equivalence of installation and echelon variables is apparent, but the benefits of using echelon variables may not be obvious. Clark and Scarf (1960) show that the optimal ordering policy for this model could be described by echelon base-stock levels, $z_1, z_2, \ldots, z_N$: If the echelon inventory X_n^j is below z_j, order $z_j - X_n^j$; otherwise, order nothing. (When there is insufficient stock at an upstream installation, a partially filled order is preferable to no deliveries at all.)

Clearly, the echelon inventory variables are nondecreasing in echelon levels and (see Figure 2) each potential value of Y_n^j is bounded below by X_n^j and above by X_n^{j+1}, except for $j = N$. Capacity limits impose an additional constraint on each Y_n^j. When demand is realized, all Xs and Ys are shifted leftwards by the size of the demand.

We define the entire model with the following recursion statements.

Definition (Core Model)

$$V_n(\widetilde{X}_n) = \min_{\widetilde{Y}_n \in \mathscr{A}(\widetilde{X}_n)} J_n(\widetilde{Y}_n), \tag{1}$$

$$J_n(\widetilde{Y}_n) = L(\widetilde{Y}_n) + \beta \mathbb{E}_{D_n} V_{n-1}(\widetilde{Y}_n - D_n), \tag{2}$$

$$V_0(\cdot) = 0, \tag{3}$$

Figure 2. Echelon variables provide natural ordering.

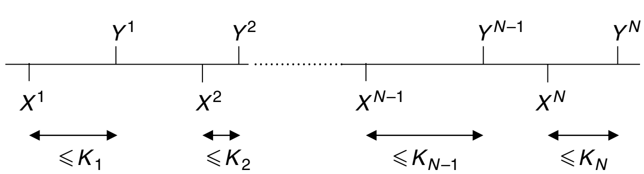

where

$$L(\widetilde{Y}_n) = \mathbb{E}_{D_n}\left\{\left(\sum_{j=1}^N h_j\right)(Y_n^1 - D_n)^+ + p(D_n - Y_n^1)^+\right\} + \sum_{i=2}^N h_i(Y_n^i - Y_n^1)$$
$$= \sum_{i=1}^N h_i(Y_n^i - \mathbb{E}_{D_n} D_n) + \left(p + \sum_{i=1}^N h_i\right)\mathbb{E}_{D_n}(D_n - Y_n^1)^+, \tag{4}$$

$$\mathscr{A}(\widetilde{X}_n) = \left\{ \widetilde{Y}_n \in \Re^N \;\middle|\; \begin{array}{c} X_n^i \leqslant Y_n^i \leqslant X_n^{i+1}, Y_n^i - X_n^i \leqslant K_i, \\ i = 1, \ldots, N-1, \\ X_n^N \leqslant Y_n^N, Y_n^N - X_n^N \leqslant K_N, \end{array} \right\}$$

and

$$\widetilde{X}_n = (X_n^1, X_n^2, \ldots, X_n^N),$$
$$\widetilde{Y}_n = (Y_n^1, Y_n^2, \ldots, Y_n^N),$$
$$\widetilde{Y}_n - D_n = (Y_n^1 - D_n, Y_n^2 - D_n, \ldots, Y_n^N - D_n).$$

V_n represents the expected discounted costs of operating under the optimal inventory policy in this capacity-limited system for a time horizon of n periods. L represents the periodic costs, or the costs (holding and backorder) incurred in a single period. The notation x^+ denotes $\max(x, 0)$ and $(a, b)^+$ denotes $\max(a, b)$. The discount factor β assumes values $\beta \geqslant 0$. The costs of purchasing goods, $c \geqslant 0$ (assumed linear in amount), from the external supplier, and the revenues, $r \geqslant 0$ (assumed linear in amount), collected from the end customer are omitted from this model. It can be easily demonstrated that these amounts can be absorbed into the holding and penalty costs by redefining $p' = p + (r - c)(1 - \beta)$ and $h'_N = h_N + c(1 - \beta)$. When $p + r(1 - \beta) > c(1 - \beta)$, the derivations and the optimal policy structure (similar to Veinott 1966) are not affected.

3. The Disguised Base-Stock Policy

A base-stock policy attempts to bring an echelon inventory up to its base-stock level if the inventory is below this level and orders nothing otherwise. In this section, we demonstrate some behavior of the capacitated model that could be construed as non-base-stock. Table 1 shows optimal behavior of the capacitated model for the following parameters. In this example, $\mathbb{E}D = 9.6 < 10 = K_1 = K_2$.

$N = 2, \quad c = 0, \quad h_2 = 0.05, \quad h_1 + h_2 = 1, \quad p = 10,$

$r = 0, \quad \beta = 0.9, \quad K_1 = 10, \quad K_2 = 10, \quad \Pr(D = 7) = 0.1,$

$\Pr(D = 8) = 0.2, \quad \Pr(D = 9) = 0.25, \quad \Pr(D = 10) = 0.1,$

$\Pr(D = 11) = 0.2, \quad \Pr(D = 12) = 0.1, \quad \Pr(D = 13) = 0.05.$

As can be seen in the last two columns of Table 1, for $x^1 \leqslant 15$ and $x^2 = 8$, the model is ordering up to echelon

Table 1. Example of the "push-ahead" effect.

Initial installation inventory		Initial echelon inventory		Installation orders		Ending echelon inventory	
x^1	x^2	X^1	X^2	a^1	a^2	Y^1	Y^2
5	8	5	13	8	10	13	23
6	8	6	14	8	10	14	24
7	8	7	15	8	10	**15**	**25**
8	8	8	16	7	9	**15**	**25**
.	.	.	.	.	.	.	.
15	8	15	23	0	2	**15**	**25**
16	8	16	24	0	2	16	26
17	8	17	25	0	2	17	**27**
18	8	18	26	0	1	18	**27**
19	8	19	27	0	0	19	**27**
20	8	20	28	0	0	20	28

levels of 15 and 25 for echelons 1 and 2, respectively, but the policy then varies from this for higher initial stocks at installation 1. The last four rows of Table 1 show the target level for echelon 2 increasing above a level of 25, apparently approaching another "base-stock" level of 27. This type of behavior, labeled the "push-ahead" effect by Speck and Van der Wal (1991a), appears to be an exception to the base-stock policy. The model optimally orders more than a typical base-stock would suggest. However, this can be shown to be a disguised version of a base-stock policy shown to be optimal in §4.

4. Finite-Horizon Results

4.1. Main Case

Consider a case when the capacity of installation 1 is the smallest in the supply chain. We demonstrate that then a modified echelon base-stock policy is optimal for the capacitated serial supply chain with $N = 2$.

Lemma 1. *Let $K_1 \leqslant K_j$. For any $\widetilde{X}_n$, all optimal $\widetilde{Y}_n$ satisfy $y_n^j \leqslant \max(K_1, x_n^j - a_n^{j-1})$ for $j > 1$.*

Proof. See the appendix.

Corollary 1. *Assume that $K_1 \leqslant K_j$ and $X_n^j - X_n^{j-1} \leqslant K_1$ for all $j \in \{2, \dots, N\}$. Then,*

(a) *The optimal Y_n^j satisfy $Y_n^j - Y_n^{j-1} \leqslant K_1$ for all $j > 1$.*

(b) *If the optimal policy is followed, then the inventory positions satisfy $X_m^j - X_m^{j-1} \leqslant K_1$ for all $j > 1$ and $m < n$.*

(c) *All capacities may be replaced with capacities equal to K_1 without affecting costs.*

Corollary 1(c) is true because none of the capacity in excess of a level equal to K_1 is used. This is apparent for all installations $j < N$ because all $a_n^j \leqslant x_n^{j+1} \leqslant K_1$ (from Corollary 1(b)). It is true for installation N also, because the resulting inventory $x_{n-1}^N \leqslant K_1$, and thus $a_n^N \leqslant K_1$ also.

Remark 1. If the bottleneck capacity is at echelon i_0, i.e., $K_{i_0} \leqslant K_i$ for all i, then Lemma 1 and Corollary 1 hold for all echelons $j > i_0$.

Lemma 1 can be easily modified to include lead times:

Remark 2. If there are positive lead times λ_j for delivery of goods from installation $j+1$ to j, then the installation inventories x_j and y_j additionally include goods shipped from installation $j+1$ that have not yet arrived at installation j. Under this redefinition,

- For any $\widetilde{X}_n$, all optimal $\widetilde{Y}_n$ satisfy $y_n^j \leqslant \max((\lambda_i+1)K_{i_0}, x_n^j - a_n^{j-1})$ for $j > i_0$, where $K_{i_0} = \min_j K_j$ (i_0 is the bottleneck).

Consequently, upstream of the bottleneck, Corollaries 1(a) and 1(b) hold, but the limit of K_1 (or "band," see below) needs to be replaced by a limit of $(\lambda_j + 1)K_1$. Corollary 1(c) holds unchanged. For further discussion of the effect of lead times, see §4.3.1.

It may be opportune to define the following region we commonly refer to as the "band."

Definition 1. Feasibility band, or simply band, is defined as $\mathcal{S} := \{\widetilde{X} \subset \Re^N \mid X^j \leqslant X^{j+1} \leqslant X^j + K_1, j \geqslant 1\}$.

This band establishes the region where the inventory at all installations, except installation 1, does not exceed K_1. Lemma 1 determines that when following an optimal policy from an inventory position outside the band, the system will traverse to the band in the most direct manner. Corollary 1(a) implies that the installations upstream (i.e., away from the consumer) of the most constrained installation will not hold more than K_1 units of inventory once the inventory levels are within the band; i.e., once in the band, the system will remain within the band.

Before showing the structure of the optimal policy, we prove some properties of the model.

Property 1. (a) J_n is continuous and convex.

(b) V_n is continuous and convex.

(c) $\{V_n\}$ is nondecreasing in n.

Proof. (a) and (b) are proved by induction. Because $V_0 = 0$, $J_1 = L$ is clearly continuous and convex because each term in L (see (4)) is continuous and convex. Using Proposition B-4 in Heyman and Sobel (1984), convexity of J_n and set convexity of $\mathcal{A} := \bigcup_{\widetilde{X}_n} \mathcal{A}(\widetilde{X}_n)$ implies that V_n is convex. Continuity of V_n follows from convexity of V_n for all internal points of the feasibility set and is guaranteed from the continuity of J_n and compactness and convexity of $\mathcal{A}(\widetilde{X}_n)$ for the border points. Assume that V_n is continuous and convex. Because expectation and linear transformation preserve both continuity and convexity, $\mathbb{E}V_n(\widetilde{Y} - D)$ is continuous and convex. Therefore, J_{n+1}, as a sum of two continuous and convex functions, is continuous and convex. (c) $\{V_n\}$ cannot decrease in n from the nonnegativity of the one-period function, L. □

While the objective function is not generally separable, it does separate within the band.

Theorem 1. *Assume $N = 2$, $X_n^2 - X_n^1 \leqslant K_1$, and $K_1 \leqslant K_2$. The two-installation model can be decomposed into* *programs dependent only upon the echelon inventories, as follows*:

$$V_n(X_n^1, X_n^2) = V_n^1(X_n^1) + V_n^2(X_n^2). \tag{5}$$

In addition, V_n^1 and V_n^2 are convex.

First, we state without a proof, a lemma used to show decoupling of the function value.

Lemma 2 (Karush 1959). (a) *If a function $f(y)$ is convex on $(-\infty, \infty)$ and attains its minimum at y^*, then*

$$\min_{a \leqslant y \leqslant b} f(y) = f^L(a) + f^U(b),$$

where $f^L(a) := \min_{a\leqslant y} f(y) = f(\max(a, y^*))$ is convex nondecreasing in a and $f^U(b) := f(b) - \min_{b\leqslant y} f(y) = f(b) - f(\max(b, y^*))$ is convex nonincreasing in b.

(b) *If a function $f(y)$ is quasi-convex on $(-\infty, \infty)$ and attains its minimum at y^*, then*

$$\min_{a \leqslant y \leqslant b} f(y) = f^L(a) + f^U(b),$$

where $f^L(a) := \min_{a\leqslant y} f(y) = f(\max(a, y^*))$ is nondecreasing in a and $f^U(b) := f(b) - \min_{b\leqslant y} f(y) = f(b) - f(\max(b, y^*))$ is nonincreasing in b.

Remark 3. Note that Lemma 2(a) is taken directly from Karush (1959), adjusted for notation. Lemma 2(b) also may be straightforwardly shown. Porteus (2002) explains that when $g(x) := \mathbb{E} f(x - X)$, where X is a Pólya frequency function random variable and f is quasi-convex, then g is also quasi-convex.[2] However, the decomposition of the optimality results cannot be extended to the context with quasi-convex periodic costs for echelon 1. This is simply because the optimization occurs across both echelon inventory decisions, and the sum of quasi-convex functions is not necessarily quasi-convex. There appear to be no obvious conditions that could prove sufficient to yield structural results in such a context.

Proof of Theorem 1. The critical element of the proof is rephrasing the constraints into simpler, but *equivalent*, conditions. Such rephrasing holds in *all periods* if the initial point is within the band. From the definition of $\mathscr{A}(\widetilde{X}_n)$, we have

$$X_n^1 \leqslant Y_n^1 \leqslant X_n^2, \tag{6}$$

$$Y_n^1 \leqslant X_n^1 + K_1, \tag{7}$$

$$X_n^2 \leqslant Y_n^2 \leqslant X_n^2 + K_2. \tag{8}$$

From Corollary 1(b), because the beginning inventory is within the band (i.e., $\widetilde{X}_n \in \mathscr{S}$) and $K_1 \leqslant K_2$,

$$Y_n^2 \leqslant Y_n^1 + K_1. \tag{9}$$

Because $K_1 \leqslant K_2$ and $Y_n^1 \leqslant X_n^2$ (from Equation (6)), we achieve $Y_n^1 + K_1 \leqslant X_n^2 + K_2$ and (8) can be replaced by (9). Because $X_n^2 - X_n^1 \leqslant K_1$, the upper bound in (6) is less than or equal to the upper bound in (7). Combining these facts, under the conditions in the theorem statement, we can restate the constraint conditions in period n as

$$X_n^1 \leqslant Y_n^1 \leqslant X_n^2 \leqslant Y_n^2 \leqslant Y_n^1 + K_1. \tag{10}$$

By starting within the band in period n, from Corollary 1, for all periods $m < n$, $X_m^2 - X_m^1 \leqslant K_1$ and thus the constraint set can be expressed by the simplified conditions (10) for all $m < n$. Most importantly, when imposing (10), we do not need to impose the capacity constraints anymore.

The proof of the theorem is by induction, and the claim holds trivially for $n = 0$. Assume it holds for $n - 1$ (induction assumption).

First, we need to demonstrate that $V_n(X_n^1, X_n^2) = V_n^1(X_n^1) + V_n^2(X_n^2)$.

$$V_n(\widetilde{X}_n) = \min_{\widetilde{Y}_n \in \mathscr{A}(\widetilde{X}_n)} \left\{ \begin{array}{l} \mathbb{E}_{D_n}\left[\left(\sum_{j=1}^{2} h_j\right)(Y_n^1 - D_n)^+ + p(D_n - Y_n^1)^+\right] \\ \quad + h_2(Y_n^2 - Y_n^1) + \beta \mathbb{E}_{D_n} V_{n-1}^2(Y_n^2 - D_n) \\ \quad + \beta \mathbb{E}_{D_n} V_{n-1}^1(Y_n^1 - D_n) \end{array} \right\}. \tag{11}$$

Define

$$f_n^2(Y_n^2) = h_2 Y_n^2 + \beta \mathbb{E}_{D_n} V_{n-1}^2(Y_n^2 - D_n).$$

Because f_n^2 is convex on $\Re$ (based on the induction assumption), from Lemma 2 we get

$$\min_{X_n^2 \leqslant Y_n^2 \leqslant Y_n^1 + K_1} f_n^2(Y_n^2) = f_n^{2L}(X_n^2) + f_n^{2U}(Y_n^1 + K_1),$$

where

$$f_n^{2L}(X_n^2) = \min_{X_n^2 \leqslant x} f_n^2(x)$$

and f_n^{2L} and f_n^{2U} are convex functions on $\Re$. Holding Y_n^1 constant and minimizing over Y_n^2,

$$\begin{aligned} V_n(\widetilde{X}_n) &= \min_{X_n^1 \leqslant Y_n^1 \leqslant X_n^2} \left\{ \begin{array}{l} \mathbb{E}_{D_n}\left[\left(\sum_{j=1}^{2} h_j\right)(Y_n^1 - D_n)^+ + p(D_n - Y_n^1)^+\right] \\ \quad - h_2 Y_n^1 + f_n^{2U}(Y_n^1 + K_1) \\ \quad + \beta \mathbb{E}_{D_n} V_{n-1}^1(Y_n^1 - D_n) + f_n^{2L}(X_n^2) \end{array} \right\} \\ &= \min_{X_n^1 \leqslant Y_n^1 \leqslant X_n^2} \left\{ \begin{array}{l} \mathbb{E}_{D_n}\left[\left(\sum_{j=1}^{2} h_j\right)(Y_n^1 - D_n)^+ + p(D_n - Y_n^1)^+\right] \\ \quad - h_2 Y_n^1 + f_n^{2U}(Y_n^1 + K_1) \\ \quad + \beta \mathbb{E}_{D_n} V_{n-1}^1(Y_n^1 - D_n) \end{array} \right\} \\ &\quad + f_n^{2L}(X_n^2) \\ &= \min_{X_n^1 \leqslant Y_n^1 \leqslant X_n^2} f_n^1(Y_n^1) + f_n^{2L}(X_n^2). \end{aligned} \tag{12}$$

Again, applying Lemma 2,

$$\begin{aligned} V_n(\widetilde{X}_n) &= f_n^{1L}(X_n^1) + f_n^{1U}(X_n^2) + f_n^{2L}(X_n^2) \\ &= V_n^1(X_n^1) + V_n^2(X_n^2), \end{aligned}$$

where

$$V_n^1(X_n^1) = \min_{X_n^1 \leqslant Y_n^1} \left\{ \begin{array}{l} \mathbb{E}_{D_n}\left[\left(\sum_{j=1}^{2} h_j\right)(Y_n^1 - D_n)^+ + p(D_n - Y_n^1)^+\right] \\ -h_2 Y_n^1 + f_n^{2U}(Y_n^1 + K_1) \\ +\beta \mathbb{E}_{D_n} V_{n-1}^1(Y_n^1 - D_n) \end{array} \right\},$$

$$V_n^2(X_n^2) = \min_{X_n^2 \leqslant Y_n^2} \left\{ f_n^{1U}(X_n^2) + h_2 Y_n^2 + \beta \mathbb{E}_{D_n} V_{n-1}^2(Y_n^2 - D_n) \right\},$$

and both V_n^1 and V_n^2 are convex from Lemma 2, which completes the induction. □

Let γ_n^{1*} and γ_n^{2*} be the minimizers of f_n^1 and f_n^2, respectively. We now formally define and discuss the induced penalty functions.

Definition 2 (Induced Penalty Functions). The induced penalty functions incurred in the two-echelon model are

$$f_n^{2U}(x) = \begin{cases} 0, & x \geqslant \gamma_n^{2*}, \\ h_2(x - \gamma_n^{2*}) \\ \quad + \beta \mathbb{E}_{D_n}[V_{n-1}^2(x - D_n) - V_{n-1}^2(\gamma_n^{2*} - D_n)], \\ & x < \gamma_n^{2*}, \end{cases}$$

and

$$f_n^{1U}(x) = \begin{cases} 0, & x \geqslant \gamma_n^{1*}, \\ l(x) - l(\gamma_n^{1*}) - h_2(x - \gamma_n^{1*}) \\ \quad + f_n^{2U}(x + K_1) - f_n^{2U}(\gamma_n^{1*} + K_1) \\ \quad + \beta \mathbb{E}_{D_n}[V_{n-1}^1(x - D_n) - V_{n-1}^1(\gamma_n^{1*} - D_n)], \\ & x < \gamma_n^{1*}, \end{cases}$$

where

$$l(x) = \mathbb{E}_{D_n}\left[\left(\sum_{j=1}^{2} h_j\right)(x - D_n)^+ + p(D_n - x)^+\right].$$

Just as Clark and Scarf (1960) determined that an induced penalty function acted upon the upper echelon as a punishment when it was unable to supply enough stock for the lower echelon to achieve its optimal inventory position (base-stock level), the capacitated case has similar induced penalties. First, echelon 1 incurs a penalty (f_n^{2U}) whenever its combination of order and capacity level ($Y_n^1 + K_1$) fails to reach a sufficient level. This can be interpreted as assuming (as a cost) the additional benefit that echelon 2 would have accrued if the lowest installation were not capacity limited. (Corollary 1(a) proves that the higher echelons' orders are limited as a result of the lowest installation's capacity being the "bottleneck" of the production system.) Second, the higher echelon incurs a penalty (f_n^{1U}) whenever its stock is insufficient to supply echelon 1 with an amount needed to achieve its base-stock level. This is analogous to the induced and penalty in Clark and Scarf (1960). Each installation potentially incurs a penalty for the limitation it imposes upon the other installation. Note that while f^{1U} and f^{2U} are "normalized," i.e., equal to 0 at γ_n^{1*} and γ_n^{2*}, f^{1L} and f^{2L} are not. Instead, $f^{1L} \min f^1$ and $f^{2L} \min f^2$ are the costs of deviating from the minimums.

Let us now construct a "capacitated" version of a base-stock policy. This policy is within the family of generalized kanban policies, where K_1 is the number of kanbans at stage 2.

Definition 3 (MEBS Policy). The modified echelon base-stock policy (MEBS) can be written as ($Y_n^{j*} \geqslant X_n^j$ for all j)

$$Y_n^{1*} = \min(z_n^1, X_n^1 + K_1, X_n^2),$$
$$Y_n^{j*} = \min(z_n^j, Y_n^{(j-1)*} + K_1, X_n^{j+1}) \quad \text{for } j = 2, \ldots, N-1,$$
$$Y_n^{N*} = \min(z_n^N, Y_n^{(N-1)*} + K_1).$$

Consider the function $J_n(\cdot)$, defined in (2). We define the following: $\zeta_n := \arg\min_{Y^1} J_n^1(Y^1)$ and $z_n^2 := \arg\min_{Y^2} J_n^2(Y^2)$, where $J_n^1(Y^1) = h_1(Y^1 - \mathbb{E}_{D_n} D_n) + (p + \sum_{i=1}^{N} h_i)\mathbb{E}_{D_n}(D_n - Y^1)^+ + \beta \mathbb{E}_{D_n} V_{n-1}^1(Y^1 - D_n)$ and $J_n^j(Y^j) = h_j(Y^j - \mathbb{E}_{D_n} D_n) + \beta \mathbb{E}_{D_n} V_{n-1}^j(Y^j - D_n)$ for $j > 1$. If $z_n^2 - \zeta_n \geqslant K_1$, then $z_n^1 := \arg\min_{Y^1} J_n(Y^1, Y^1 + K_1)$, else $z_n^1 := \zeta_n$. From these definitions we see that if the intersection of the two echelons' minimizing points is within the band, then each of them is the base-stock level. However, if the intersection does not occur within the band, then echelon 1's base-stock level will be found along the upper edge of the band, where $Y^2 = Y^1 + K_1$.

Theorem 2. *The modified echelon base-stock policy* (*MEBS*) *with parameters* (z_n^1, z_n^2) *defined above is optimal for an* $N = 2$ *system where* $K_1 \leqslant K_2$ *and* $X^2 - X^1 \leqslant K_1$.

Proof. See the appendix.

Recall that, based on Lemma 1, the optimal policy outside the above-defined band, $\mathscr{S}$, is to order nothing until the inventory position is drawn into the band. While in the band, the MEBS policy differs from the "uncapacitated" echelon base-stock policy (demonstrated in Clark and Scarf 1960) only in that it constrains the system to operate within the inventory band. Let us revisit the "counterexample" presented in Table 1 (§3). The parameters are found to be $z^1 = 15$ and $z^2 = 27$. Figure 3 illustrates this case. Note that set $\mathscr{A}(X_1, X_2) = \{(Y_1, Y_2) \mid X_1 \leqslant Y_1 \leqslant X_2 \leqslant Y_2, Y_1 \leqslant X_1 + K_1, Y_2 \leqslant K_2\}$ is truncated to the band $\{Y_1 \leqslant Y_2 \leqslant Y_1 + K_1\}$; i.e., desired inventory at echelon 2 is limited by $\min(X^1 + K_1, z^2)$. Within this feasible set, MEBS policy first chooses Y_1 closest to the target z_1 and then Y_2 closest to z_2. Thus, $z^1 = 15$ implies that for $X_1 \leqslant 15$, optimal $Y_1 \leqslant 15$ and, consequently, $Y_2 \leqslant 15 + 10 = 25 < z^2$. Higher starting levels of inventory ($X_1 > 15$) allow, however, Y_2 to climb towards and eventually reach the desired $z^2 = 27$.

Note that this result indicates that the form of the base-stock policy, which Glasserman and Tayur (1995) explore numerically, could be suboptimal because it does not incorporate the censoring effect described in MEBS. Speck and Van der Wal (1991b) consider a heuristic approach to determining approximate values of the optimal average cost

Figure 3. Illustration of MEBS policy with target $(z_1, z_2) = (15, 27)$. Starting inventory is at the lower-left corner of the feasible region and the optimal inventory (Y_1, Y_2) is denoted by a black circle.

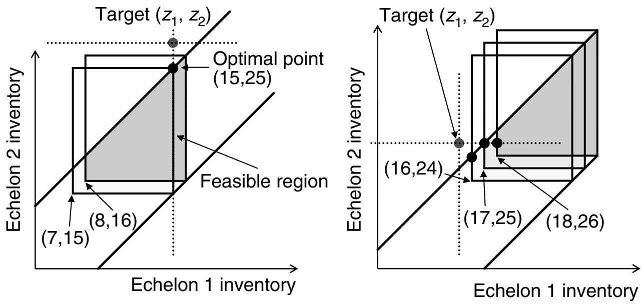

Note. Left: Starting from either (7, 15) or (8, 16) it is optimal to raise inventory levels to (15, 25).

Right: The optimal decisions for (16, 24), (17, 25), and (18, 26) are (16, 26), (17, 27), and (18, 27), respectively.

value functions. They do not, however, consider the possibility that the gap between the base-stock levels (i.e., $z^2 - z^1$) exceeds K_1 while $K_1 \leqslant K_2$. While there is similarity in the direction of adjustments, we can also see that the Speck and Van der Wal (1991a) heuristic is different in that they suggest a change in echelon 1's base-stock level based on installation 2's inventory position. That is, at some level as installation 2's inventory position increases, the base-stock level at echelon 1 is also forced to increase. This is a contrast to the optimal policy where there is an increase in echelon 2's base-stock level when there is an increase in installation 1's inventory position.

Theorem 3. *The optimal base-stock levels, z_n^1 and z_n^2, are*
(a) *nondecreasing in period n*; *and*
(b) *nonincreasing in K_1 so long as $K_1 \leqslant K_2$ is maintained.*

Proof. See the online appendix at http://or.pubs.informs.org/pages/collect.html.

The optimal base-stock levels increase in the horizon length, which is consistent with other inventory models. Later, we show that these levels stabilize to steady-state values. If a starting inventory is above the base-stock levels, but still within the band, $\mathcal{S}$, no material will be ordered and demand will progressively draw down the echelon inventory levels equally until the inventory levels fall below the base-stock levels, whereupon it will order up to the base-stock levels, if possible. All inventory paths will then continue to be below the base-stock levels, and the inventory territory above the base-stock levels within the band will not be revisited.[3]

4.2. Discussion of the Model Assumptions

The restrictions that $N = 2$ and installation 2 has no lead times other than those derived naturally from the periodicity of the model are somewhat limiting, but necessary for our proof to hold. The proof we provide is based on the equivalence of the two capacity constraints to a constraint on installation 2 inventory. The same equivalence does not hold for more than two echelons or when arbitrary lead times are introduced (but we do not claim that MEBS or an optimal policy similar to MEBS is not optimal in generalizations of the model considered in this paper).

The conditions under which Theorem 1 is proven include $X^2 - X^1 \leqslant K_1$, which cannot be guaranteed if lead times are permitted at installation 2 unless the base stocks differ by K_1 or less. Especially with long lead times, the base-stock levels are unlikely to differ by less than K_1. Let us demonstrate the difficulty of echelon 2 lead times with a simple scenario. Suppose a situation has arisen where the pipeline inventory leading to installation 2 sums to a level greater than K_1. Now a succession of very low consumer demand realizations occurs so that installation 1's pipeline inventory is sufficient to cater for expected demands, and no additional orders are placed with installation 2. Inexorably, all the pipeline inventory arrives at installation 2 and the installation inventory stock exceeds K_1, thus contravening the condition $X^2 - X^1 \leqslant K_1$ and not allowing the capacity constraints to be omitted.

The constraint that $N = 2$ is closely related to the lead-time restrictions, as demonstrated by Glasserman and Tayur (1994). Another restriction imposed is that $K_1 \leqslant K_2$. While this is representative of many real systems, especially where we can also imbed this two-stage system in longer systems of a specific type (see §4.3.3), it is worth considering the opposite for the two-stage system, i.e., $K_2 < K_1$. A numerical example in §4.3.4 illustrates a behavior with multiple "plateaus" for each echelon, which does not resemble MEBS policy.

It is interesting to note that our MEBS policy may be interpreted as a special case of generalized kanban policies. Axsäter and Rosling (1993) described kanban policies as a special case of base-stock policies. Such a characterization holds when considering the chance of backorders as being extremely low. However, Veatch and Wein (1994) show that there do exist cases when kanban and base-stock policies perform significantly differently. Later, Axsäter and Rosling (1999) extend their previous classification of inventory policies and allow kanban-like constraints to be imposed on top of base-stock policies. The MEBS policy proposed here is a special case of their structure. It is also the same as that described in Buzacott and Shanthikumar (1992). The relationship with the results of Veatch and Wein is the most interesting. As noted in §1, Veatch and Wein (1994) consider a capacity-limited, two-installation serial system under the assumptions of exogenous Poisson demand and controllable production rates. Using simulation, they find that kanban policies tend to be superior (although not necessarily optimal) to base-stock policies when the downstream installation is the bottleneck. When the bottleneck is at the upstream installation, base-stock policies are superior. While the setting modeled in their paper differs from our model, the spirit of the heuristic is the same—limit the inventory at the upstream installation. While kanban policy seems to be superior in continuous settings with Poisson arrivals, we find that a generalized kanban policy is optimal in a discrete-time general-demand setting. It is also important to note that kanban systems—sometimes justified as a suboptimal inventory policy but a good incentive mechanism—in situations we consider, actually are exact applications of the optimal inventory policy.

Despite the specific assumptions we make in this paper, some generalizations are possible. In §§4.3.1 and 4.3.2 we show, among other things, how some of the lead times can be incorporated and that the MEBS policy is optimal in a Markov-modulated demand model. Clearly, the optimal base-stock levels will fluctuate according to the Markov-chain state and inventory levels above base-stock levels may occur given such fluctuations. In addition, §4.3.3 demonstrates how our core results may be applied in other supply chain configurations.

4.3. Extensions

4.3.1. Lead Times. The core model has no lead times, other than those that occur through the natural dynamics of the supply chain. Here we expand the model to include analysis of delivery lead times.

We assume that the holding cost is constant throughout the pipeline at the level of the higher installation. We could choose a separate pipeline holding cost, higher than h_2 but lower than (h_1+h_2), to reflect the reality that additional costs have been incurred by this delivery (e.g., transportation cost, insurance coverage, etc.), but the value-added costs at the delivery destination have yet to be incurred. (Analytically, any linear holding cost for pipeline inventory, even outside the interval (h_2, h_1+h_2), can be incorporated into the model. This would not, however, add particular insight into the inclusion of lead times into the core model.)

Consider, as before, a two-stage multiechelon inventory system with capacity limits $K_1 \leqslant K_2$. There is a delivery lead time of (integer) λ periods from installation 2 to installation 1. Installation 2 has no lead times. Let

$$\tilde{a}_n = (a^1_{n+1}, a^1_{n+2}, \ldots, a^1_{n+\lambda}).$$

The transition equations are

$$x^1_{n-1} = x^1_n + a^1_{n+\lambda} - D_n, \quad x^2_{n-1} = x^2_n + a^2_n - a^1_n.$$

Let the pipeline inventory costs be h_2 per unit. Installation 1 assumes the cost of the stock, h_1+h_2, once it arrives at the site. The systemwide inventory at the end of period n is

$$\begin{aligned} &x^2_n + a^2_n + a^1_{n+1} + a^1_{n+2} + a^1_{n+3} \\ &\quad + \cdots + a^1_{n+\lambda-1} + [x^1_n + a^1_{n+\lambda} - D_n]^+ \\ &= X^2_n + a^2_n - x^1_n - a^1_{n+\lambda} + [x^1_n + a^1_{n+\lambda} - D_n]^+. \end{aligned}$$

We account for all costs when they occur except for the costs at installation 1. (This is similar to the development of Federgruen 1993.) At installation 1 the costs are incurred in the period in which the order is delivered, i.e., λ periods after the order is triggered and, therefore, are discounted by β^λ to bring the costs back to the ordering period. The cost equations become

$$\begin{aligned} &h_2(X^2_n + a^2_n) - \beta^\lambda h_2 \mathbb{E}\left[x^1_n + a^1_n + \sum_{i=1}^{\lambda}(a^1_{n+i} - D_{n-i})\right] \\ &\quad + \beta^\lambda \mathbb{E}\left\{(h_1+h_2)\left[x^1_n + \sum_{i=0}^{\lambda}(a^1_{n+i} - D_{n-i})\right]^+ \right. \\ &\qquad \left. + p\left[-x^1_n - \sum_{i=0}^{\lambda}(a^1_{n+i} - D_{n-i})\right]^+ \right\} \\ &= h_2(X^2_n + a^2_n) - \beta^\lambda h_2 \mathbb{E}(X^1_n + a^1_n - D^{(\lambda)}_n) \\ &\quad + \beta^\lambda \mathbb{E}\{(h_1+h_2)[X^1_n + a^1_n - D^{(\lambda+1)}_n]^+ \\ &\qquad + p[D^{(\lambda+1)}_n - X^1_n - a^1_n]^+\} \\ &=: L^2(X^2_n + a^2_n) + L^1(X^1_n + a^1_n), \end{aligned}$$

where $X^1_n := x^1_n + \sum_{i=1}^{\lambda} a^1_{n+i}$ and $X^2_n := x^1_n + x^2_n + \sum_{i=1}^{\lambda} a^1_{n+i}$. $D^{(j)}_n$ is the convolution of the demand random variable over j periods; the expectation operators apply to these convolutions. Evidently, the terms $X^1_n + a^1_n$ and $X^2_n + a^2_n$ are isolated and can be labeled $Y^1_n := X^1_n + a^1_n$ and $Y^2_n := X^2_n + a^2_n$, respectively. The action set for this model is

$$\mathcal{A}(x^1_n, x^2_n, \tilde{a}_n) = \{a^1_n, a^2_n \mid 0 \leqslant a^1_n \leqslant K_1, a^1_n \leqslant x^2_n, 0 \leqslant a^2_n \leqslant K_2\}$$

or equivalently

$$\begin{aligned} \mathcal{A}(\widetilde{X}_n) = \{\widetilde{Y}_n \in \Re^2 \mid X^1_n \leqslant Y^1_n \leqslant X^2_n \leqslant Y^2_n \leqslant X^2_n + K_2, \\ Y^1_n \leqslant X^1_n + K_1\}. \end{aligned}$$

With respect to the dynamic programming recursion, the last λ periods will be constants because no decisions made in those periods will have any cost effects during the time remaining. Note that we continue to have $X^2_{n-1} - X^1_{n-1} = x^2_{n-1} = Y^2_n - Y^1_n$, and the basis upon which Corollary 1 sits remains valid: It is unprofitable to order up to an amount above a level of K_1 at installation 2 because installation 1 cannot order any more than this amount in a single period. Formal translation of the lead-time model to the original model is as follows:

Definition 4 (Lead-Time Model).

$$V_n(\widetilde{X}_n) = \min_{\widetilde{Y}_n \in \mathcal{A}(\widetilde{X}_n)} J_n(\widetilde{Y}_n), \tag{13}$$

$$J_n(\widetilde{Y}_n) = \{L^1(Y^1_n) + L^2(Y^2_n) + \beta \mathbb{E}_{D_n} V_{n-1}(\widetilde{Y}_n - D_n)\}, \tag{14}$$

where

$$\begin{aligned} \mathcal{A}(\widetilde{X}_n) = \{\widetilde{Y}_n \in \Re^2 \mid X^1_n \leqslant Y^1_n \leqslant X^2_n \leqslant Y^2_n \leqslant X^2_n + K_2, \\ Y^1_n \leqslant X^1_n + K_1\}, \end{aligned}$$

$$L^2(Y_n) = h_2(Y_n),$$

$$\begin{aligned} L^1(Y_n) = &-\beta^\lambda h_2 \mathbb{E}(Y_n - D^{(\lambda)}_n) \\ &+ \beta^\lambda \mathbb{E}\{(h_1+h_2)[Y_n - D^{(\lambda+1)}_n]^+ + p[D^{(\lambda+1)}_n - Y_n]^+\}. \end{aligned}$$

The proof of the following theorem is identical to that of Theorems 1 and 2.

Theorem 4. *The lead-time model follows the MEBS policy for the conditions stated in Theorem* 1.

4.3.2. Markov-Modulated Demand. Consider a Markov-modulated demand. That is, there are M states with Markovian transitions in matrix P and a demand distribution associated with each of the states. Chen and Song (2001) show that demand-distribution-dependent echelon base-stock policies are optimal for an uncapacitated serial multiechelon system facing nonstationary demands. Their proof was necessarily complicated by the fact that they were illustrating an algorithm that had desirable economic interpretations. We now demonstrate a similar result, but use the decomposition proof used in previous sections.

In the Markov-modulated demand model, the demand expectation is augmented by $p_{mm'}$, the probability that the next state will be m' *given* the current-period state is m.

Definition 5 (Markov-Modulated Model). The value function of the nonstationary demand model is

$$V_n(\widetilde{X}_n, m) = \min_{\widetilde{Y}_n \in \mathcal{A}} \left\{ \begin{aligned} &\sum_{i=2}^{N}\sum_{j=i}^{N} h_j (Y_n^i - Y_n^{i-1}) \\ &\quad + \sum_{m'} p_{mm'} \mathbb{E}_{D_{m'}} \left[\sum_{j=1}^{N} h_j (Y_n^1 - D_{m'})^+ \right. \\ &\qquad\qquad \left. + p(D_{m'} - Y_n^1)^+ \right] \\ &\quad + \beta \sum_{m'} p_{mm'} \mathbb{E}_{D_{m'}} V_{n-1}(\widetilde{Y}_n - D_{m'}, m'), \end{aligned} \right\} \tag{15}$$

where

$$\mathcal{A}(\widetilde{X}_n) = \{\widetilde{Y}_n \mid X_n^i \leqslant Y_n^i \leqslant X_n^{i+1}, Y_n^i - X_n^i \leqslant K_i, \quad i = 1, \ldots, N-1; X_n^N \leqslant Y_n^N \leqslant X_n^N + K_N\}.$$

$m \in \{1, 2, \ldots, M\}$ is the state of the Markov chain. It is easy to verify that Corollary 1(a) remains valid for this model (because it does not require that the realizations be drawn from the same distribution). It is possible to demonstrate that the MEBS policy is optimal for this model. Because the proof is very similar to the proofs of Theorems 1 and 2, it is omitted here.

Theorem 5. *Assume that the conditions stated in Theorem* 1 *hold. For the nonstationary demand model for $N = 2$, the MEBS policy is optimal. The parameters of the optimal MEBS policy depend, however, on both the period number and state of the system.*

4.3.3. Other Capacity Conditions. While the assumption that the lowest installation must have the lowest capacity level may appear restrictive, there exist other situations for which a base-stock ordering policy is optimal. It can be easily demonstrated that a base-stock policy remains optimal for (a) a serial multiechelon system where only the two highest installations have capacity constraints $K_{N-1} \leqslant K_N \leqslant \infty$ with general lead times permitted at all stages except the highest installation, and (b) a serial multiechelon system, with general lead times at all installations, where the only capacity restriction is at the uppermost installation, $K_N < \infty$. These results are summarized in the following theorem.

Theorem 6. (a) *Consider an N-stage system without capacity limits at stages $j < N - 1$ and finite capacities $K_{N-1} \leqslant K_N < \infty$. Under the inventory condition $X_n^N - X_n^{N-1} \leqslant K_{N-1}$, $V_n(\widetilde{X}_n) = \sum_{j=1}^N V_n^j(X_n^j)$. The optimal policy is as follows: MEBS holds for the capacitated installations and the remaining installations follow an echelon base-stock policy.*

(b) *Consider an N-stage system without capacity limits at stages $j < N$ and finite capacity K_N. For this system, $V_n(\widetilde{X}_n) = \sum_{j=1}^N V_n^j(X_n^j)$. An echelon base-stock policy is optimal for this system.*

Proof. See the appendix.

Now consider an N-stage system with capacity limits at each stage and identical holding costs. Note that when holding costs for all installations are equal, the problem becomes very easy—for all stages $1 < i \leqslant N$, one cannot be worse off by forwarding the inventory to echelon 1. Thus, the system is equivalent to a one-stage capacitated system with lead time N, for which optimal policy is a modified base-stock (MBS) policy.

Corollary 2. *Consider an N-echelon model with installation holding costs $H_i = h_1$ for all i and $K_1 \leqslant K_j$ for all $j > 1$. The optimal policy is MBS.*

4.3.4. Lower Capacity at the Higher Echelon: $K_2 < K_1$. Considering that we have dealt with the $K_2 \geqslant K_1$ configuration in the $N = 2$ case, the obvious question is whether base-stock policies are optimal in the $K_2 < K_1$ configuration. Table 2 contains a numerical counterexample detailing the optimal orders for a number of starting inventory positions. This 10-period example is achieved with the following parameter values:

$$K_1 = 11, \quad K_2 = 10, \quad c = 0, \quad r = 0, \quad p = 10,$$
$$h_1 + h_2 = 1, \quad h_2 = 0.05, \quad \beta = 0.9, \quad \Pr(D=2) = 0.1,$$
$$\Pr(D=3) = 0.2, \quad \Pr(D=9) = 0.25, \quad \Pr(D=10) = 0.1,$$
$$\Pr(D=13) = 0.2, \quad \Pr(D=18) = 0.1, \quad \Pr(D=22) = 0.05.$$

The mean demand for this distribution is 9.55, which is less than the lowest capacity level, $K_2 = 10$.

It appears that echelon 1 is striving to order up to levels of 22, 23, and 24, while echelon 2 appears to be trying to reach levels of 43, 47, 48, and 49 (not shown). If this does demonstrate a base-stock policy, it is one more ornate and intricate than we can envision.

Table 2. Counterexample for $K_2 < K_1$ in the $N = 2$ case.

Initial installation inventory		Initial echelon inventory		Installation orders		Ending echelon inventory	
x^1	x^2	X^1	X^2	a^1	a^2	Y^1	Y^2
10	15	10	25	11	10	21	35
11	15	11	26	11	10	22	36
12	15	12	27	10	10	22	37
13	15	13	28	10	10	23	38
14	15	14	29	9	10	23	39
15	15	15	30	9	10	24	40
16	15	16	31	8	10	24	41
17	15	17	32	7	10	24	42
18	15	18	33	6	10	24	43
.	.	.	.	.	.	.	.
24	15	24	39	0	4	24	43
25	15	25	40	0	3	25	43
26	15	26	41	0	3	26	44
27	15	27	42	0	3	27	45
28	15	28	43	0	3	28	46
29	15	29	44	0	3	29	47
30	15	30	45	0	2	30	47
31	15	31	46	0	2	31	48
32	15	32	47	0	1	32	48

4.3.5. Effect of Variance upon Costs and Recommended Capacity Levels. In this subsection, we illustrate the effect demand variance may have upon costs and optimal capacity levels with a numerical example. We numerically solve the system with a discrete demand distribution for various coefficients of variation. As expected, the costs at the best operating level[4] increase with the coefficient of variation and, as K_1 increases (maintaining the condition, $K_1 \leqslant K_2$), this best operating level cost also decreases in a convexlike manner—see Figure 4. The parameters used for this example are $p = 5$, $h_1 + h_2 = 1$, $h_2 = 0.5$, $r = 0$, $c = 0$, and $\beta = 0.9$, with $K_1 \in \{10, 11, 12, 15, 20\}$. The demand distributions were designed so that the mean was kept at 9.8 and the coefficient of variation was 0.1, 0.2, and 0.4. Using these data, the algorithm was run until the difference of successive value functions was less than 0.005. This resulted in horizon lengths that differed for different K_1s and demand distributions. A summary of these horizon lengths appears in Table 3.

Consider the three curves denoting the best operating-cost curves for three levels of coefficient of variation (cv). We observe that doubling the cv from 0.1 to 0.2 has a smaller effect upon costs than doubling the cv from 0.2 to 0.4, and safety stocks are not proportional to standard deviations, as they are for one-stage systems. Lastly, the percentage labels in Figure 4 indicate the percentage cost above the uncapacitated system. Confirming intuition, the number of units of additional capacity required to get within 1% of uncapacitated costs is smaller for systems with less variance.

Now, given an appropriate capacity acquisition-cost function, we can draw conclusions about optimal one-time capacity investment. Two elements of such a cost function could be a fixed part (independent of the level of K_1) and a variable part (dependent on K_1). If the variable part is convex increasing (including a linear function), then clearly the sum of the best operating-cost function (V_n) and the capacity acquisition costs is also convex in K_1 with a finite minimizing point, K_1^*, and recommendations about optimal capacity investments to minimize total costs (operating and investment costs) may be made.

The method of numerical computation is based on the value-iteration algorithm. The analytical results offer

Figure 4. The effect of increasing capacity and the coefficient of variation on the best operating level.

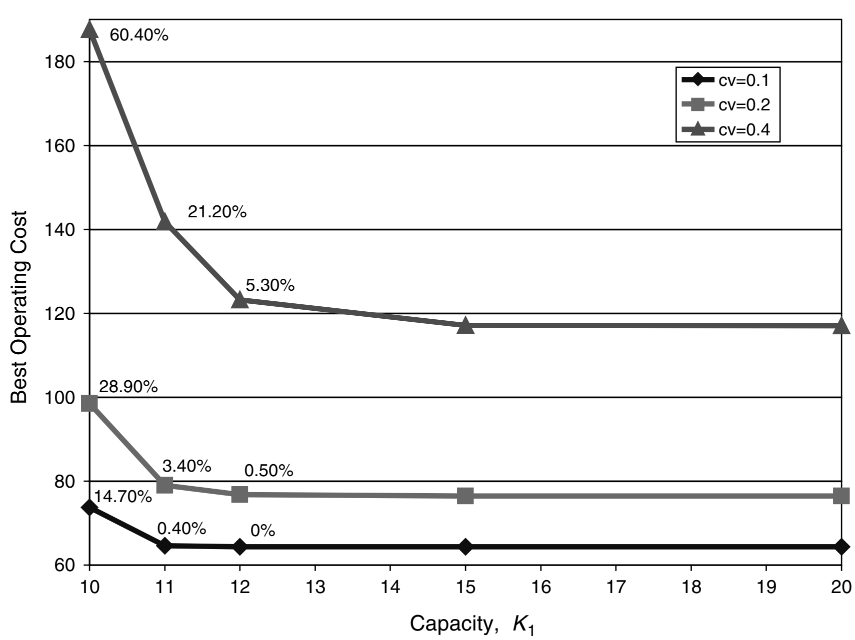

Table 3. Example of horizon length required for the convergence of up-to levels.

	Horizon length for		
Capacity	cv = 0.1	cv = 0.2	cv = 0.4
10	33	33	38
11	31	29	33
12	31	29	31
15	31	29	31
20	31	29	31

additional opportunities for computational efficiencies. Namely, the convexity result permits us to search for the minimums efficiently, while Lemma 1 enables us to restrict the search to a subset of the state space, the band. Using a zero salvage value function and a discretized state space, the convexity of the value function is exploited to determine the optimal decisions.

5. Infinite-Horizon Results

In this section, we demonstrate that the key results found for the finite-horizon model in §4 also hold for the infinite-time horizon. Federgruen and Zipkin (1984) extend the decomposition of Clark and Scarf (1960) into separate echelon-based dynamic programs and show that the optimal policy holds for the infinite horizon by demonstrating it for each of the decomposed programs. We show that the original combined finite-horizon model described in §4 converges in cost and policy in the infinite horizon. For notational simplicity we demonstrate these for the basic model, but the same results can be derived for the extensions described in §4.2. (The results achieved in Federgruen and Zipkin 1984 could be achieved more easily using the techniques illustrated here. The closure is easily demonstrated by bounding the optimal base-stock levels. Despite not having natural bounds as in case of the capacitated problem, Theorem 8 below establishes a bound on undominated target levels, and the problem can be translated into an equivalent one with the feasible actions limited to the states within this bounded area.) In this section, we additionally assume that $\mathbb{E}D < K_1$.

5.1. Discounted Cost

Let $0 \leqslant \beta < 1$. Consider the following infinite-horizon cost for policy δ, which defines the order quantities, a_k^i, $i \in \{1, \ldots, N\}$:

$$V_\delta(\widetilde{X}_0) = \sum_{k=1}^{\infty} \beta^{k-1} \mathbb{E}_{D_k} \left\{ \begin{array}{l} H_1(x_k^1 + a_k^1 - D_k)^+ + p(D_k - x_k^1 - a_k^1)^+ \\ \quad + \sum_{i=2}^{N} H_i(x_k^i + a_k^i - a_k^{i-1}) \end{array} \right\},$$

where $x_{k+1}^j = x_k^j - a_k^{j-1} + a_k^j$ for $j > 1$, $x_{k+1}^1 = x_k^1 + a_k^1 - D_k$, $H_i = \sum_{j \geqslant i} h_j$ (as defined before), and all other dynamic relationships are as before. (Note that we count time forward in this section, using k rather than n.) Let us define the minimal infinite-horizon cost as

$$V^*(\widetilde{X}) = \inf_{\delta} V_\delta(\widetilde{X}).$$

Theorem 7. *The finite-horizon function, V_n, converges to a finite-valued infinite-horizon counterpart; that is,*

$$V^*(\widetilde{X}) = \lim_{n \to \infty} V_n(\widetilde{X}) < \infty \quad \textit{for all } \widetilde{X} \in S.$$

Proof. The cost of any feasible policy in the infinite horizon is bounded from above,

$$\begin{aligned} V_\delta(\widetilde{X}_0) &= \sum_{k=1}^{\infty} \beta^{k-1} \mathbb{E}_{D_k} \Bigg[\sum_{i=2}^{N} H_i x_k^i + \sum_{i=2}^{N} h_i a_k^i - H_2 a_k^1 \\ &\qquad + H_1(x_k^1 + a_k^1 - D_k)^+ \\ &\qquad + p(D_k - x_k^1 - a_k^1)^+ \Bigg] \\ &\leqslant \sum_{k=1}^{\infty} \beta^{k-1} \mathbb{E}_{D_k} \Bigg[\sum_{i=2}^{N} H_i x_k^i + \sum_{i=2}^{N} h_i a_k^i \\ &\qquad + H_1(|x_k^1 + a_k^1 - D_k|) \\ &\qquad + p(|D_k - x_k^1 - a_k^1|) \Bigg] \\ &\leqslant \sum_{k=1}^{\infty} \beta^{k-1} \Bigg[\sum_{i=2}^{N} H_i(x_0^i + kK_i) + \sum_{i=2}^{N} h_i K_i \\ &\qquad + H_1(|x_0^1| + kK_1) + p(|x_0^1| + k\mathbb{E}D) \Bigg] \\ &< \infty. \end{aligned}$$

The bounds are justified based on $\sum_{k=1}^{\infty} k\beta^k = \beta/(1-\beta)^2 < \infty$, $\sum_{k=0}^{\infty} \beta^k = 1/(1-\beta) < \infty$ for $\beta < 1$, and the assumption that $\mathbb{E}D < \infty$, and are independent of policy δ. Clearly, the value of the optimal policy, if it exists, is also bounded. Now consider the set

$$\mathscr{B}_n(\widetilde{X}, \lambda) = \{\widetilde{Y} \in \mathscr{A}(\widetilde{X}) \mid \mathbb{E}_D[L(\widetilde{Y}) + \beta T^n V_0(\widetilde{Y} - D)] \leqslant \lambda\},$$

where T is the one-period mapping denoted by (1) and (2). We have $V_n = TV_{n-1} = T^n V_0$. Because $\mathscr{A}(\widetilde{X}_n)$ is bounded, $\mathscr{B}_n$ must be bounded. The function in the expectation is continuous in $\widetilde{Y}$ (from the continuity of finite sums; see §4.1) and $\mathscr{A}(\widetilde{X}_n)$ is closed. Therefore, $\mathscr{B}_n$ must be closed. Based on boundedness of values and the fact that $\mathscr{B}_n$ is closed, invoking Proposition 1.7 of §3.1 of Volume II of Bertsekas (1995) (equivalently, Bertsekas and Shreve 1996, Proposition 9.17),

$$\lim_{n \to \infty} V_n(\widetilde{X}) = V^*(\widetilde{X}) \quad \text{for all } \widetilde{X} \in S. \quad \square$$

We are interested as to whether the optimal policy from the finite-horizon problem converges to the optimal policy for the infinite horizon.

Theorem 8. (a) *There exists a finite upper bound on inventory targets for* $V(\widetilde{X})$, *which is independent of* $\beta \in (0,1]$; *and* (b) *there exists a finite upper bound on inventory targets for* $\mathbb{E}V(\widetilde{Y}-D)$, *which is independent of* $\beta \in (0,1]$.

An inventory target is any $Y(X) > X$ (*i.e., an up-to level where we actually increase inventory*).

Proof. See the appendix.

Theorem 9. *The optimal policy for the finite-horizon function,* V_n, *converges to its infinite-horizon counterpart. Consequently,* $\widetilde{Y}^* = \lim_{n\to\infty} \widetilde{Y}_n^*$ *exists, and* $\widetilde{Y}^*$ *minimizes* $V^*(\cdot)$.

Proof. In Theorem 7, we demonstrated that there exists $V^*(\widetilde{X}) = \lim_{n\to\infty} V_n(\widetilde{X})\ \forall \widetilde{X} \in S$, which in turn implies $J(\widetilde{Y}) := \lim_{n\to\infty} J_n(\widetilde{Y})$ for all $\widetilde{Y} \in \mathscr{A}(\widetilde{X})$. The limit exists because, from Property 1(c) in §4.1, $J_{n+1}(\widetilde{Y}) \geqslant J_n(\widetilde{Y})$ for each n and $\widetilde{Y} \in \mathscr{A}(\widetilde{X})$, and, from Theorem 7, $J_{n+1}(\widetilde{Y})$ are bounded. $\mathscr{A}(\widetilde{X})$ is convex and compact $\forall \widetilde{X} \in S$, and $J_n(\cdot)$ is convex (see Property 1(a) in §4.1). The proof of Theorem 8 verifies that the optimal base-stock levels are bounded, and hence so are the optimal decisions. Consequently, all the conditions of Theorem 8–15 in Heyman and Sobel (1984) are satisfied and we get

$$\widetilde{Y}^* = \lim_{n\to\infty} \widetilde{Y}_n^* \text{ exists} \quad \text{and} \quad \widetilde{Y}^* = \arg\min V^*(\cdot). \quad \square$$

Corollary 3. *The modified echelon base-stock policy is optimal in the discounted-cost infinite-horizon setting in the* $N=2$ *system when* $K_1 \leqslant K_2$ *and the beginning inventory satisfies* $X^2 - X^1 \leqslant K_1$.

Proof. Theorem 9 demonstrates that the optimal decisions converge in the infinite horizon. However, we must also establish that the base-stock levels also converge. Because the base-stock levels are monotonically nondecreasing in n (Theorem 3) and they are bounded (Theorem 8), $\tilde{z}_n$ converge (due to pointwise convergence). From Theorem 9, this results in the MEBS policy structure being optimal in the discounted-cost infinite horizon. $\square$

Note that other than $\mathbb{E}D < K_1$, there have been no additional restrictions on the demand distributions to demonstrate that the MEBS policy extends to the infinite horizon. To numerically evaluate these models, due to the analytical results of Theorem 9, we can use value-iteration algorithms until the differences of the value functions converge to a predefined quantity.

The case of Markov-modulated demand may be easily extended to the infinite horizon. As for the finite-horizon case (§4.3.2), the cost function depends additionally on the state $m \in \{1,2,\dots,M\}$ of the underlying Markov chain. All theorems within this section continue to hold. For Theorem 8 to apply, the underlying Markov chain needs to be ergodic, each demand distribution must satisfy $0 < \mathbb{E}D_m < \infty$ for all m, and $\sum_m p_m \mathbb{E}D_m < K_1$, where p_m is the long-term probability for state m. Theorems 7 and 9 are easily modifiable.

5.2. Average Cost

In this section, we demonstrate that the base-stock policy structure, optimal for the discounted expected cost model, is also optimal under an expected average-cost criterion in the infinite horizon. To satisfy this, the state space is now the set of all integers, and therefore the action space for each initial inventory is finite. We also assume that the demand is i.i.d. nonnegative and integer and $\mathbb{E}[D^2] < \infty$. The previous results remain valid for this more restrictive model. Note that for problems that have *finite* state and action spaces, demonstration of the convergence of the discounted-cost optimal policy to the average-cost case is quite standard (see Sennott 1989, 1999; Bertsekas 1995). Due to the inclusion of backlogging of unsatisfied demand, we cannot assume a finite state space, or bound the state space. Whereas the discounted-cost model has infinite nondenumerable state sets, we assume infinite denumerable state sets in the average-cost model. This is different than Federgruen and Zipkin (1984), who do not require discrete demand, except for the computation of the optimal policies. Discreteness of demand does simplify the analysis, and our proof does apply directly to a capacitated single-stage case analyzed by Federgruen and Zipkin (1986a), who also assume discrete demand. (For the uncapacitated case, however, to bound the inventory at a higher installation, instead of K_1, a bound based on Theorem 8 needs to be used.)

There exist a few versions of sufficient conditions that guarantee the convergence for the average-cost criterion—see Sennott (1999) for an excellent review. We focus on the conditions in Schäl (1993), which are based on the optimal discounted value function for an infinite-state Markov decision process with *unbounded* costs.

Schäl (1993) suggests two sets of sufficient conditions for the convergence of the optimal discounted value function and policy to the average-cost equivalents. We use the first set, which consists of two conditions. The first condition is straightforward. To satisfy the second condition, some notation is useful. For the sake of clarity, the infinite-horizon value function under the discounted-cost criterion is now labeled V_β.

$\Phi(\delta, \widetilde{X}) :=$ average cost of policy δ given state $\widetilde{X}$.
$\Delta :=$ set of randomized policies.
$g := \inf_{\widetilde{X}\in S} \inf_{\delta\in\Delta} \Phi(\delta, \widetilde{X}) < \infty$.
$m_\beta := \inf_{\widetilde{X}\in S} V_\beta(\widetilde{X})$.
$\underline{V}_\beta(\widetilde{X}) := \inf_{\widetilde{Y}\in\mathscr{A}(\widetilde{X})} \mathbb{E}V_\beta(\widetilde{Y}-D)$.
$\underline{m}_\beta := \inf_{\widetilde{X}\in S} \underline{V}_\beta(\widetilde{X})$.

Schäl's second condition is the following:

$$\text{(B)} \qquad \sup_{\beta<1} w_\beta(\widetilde{X}) < \infty \quad \text{for } \widetilde{X} \in S,$$

where $w_\beta(\widetilde{X}) := V_\beta(\widetilde{X}) - m_\beta$.

Let us restate a lemma from Schäl (1993), rephrased with our notation and simplified (general $\eta \geqslant 0$ is replaced with $\eta = 0$). Use of this lemma allows us to bound $w_\beta(\widetilde{X})$ in condition (B).

Lemma 3 (Schäl 1993). *For $\beta < 1$, $\widetilde{X} \in S$:*

$$w_\beta(\widetilde{X}) \leqslant \inf_{\delta\in\Delta} \mathbb{E}_{\widetilde{X}_n,\delta}\left[\sum_{n=0}^{\sigma-1} L(\widetilde{X}_n) + L(\varphi_\beta(\widetilde{X}_\sigma))\right],$$

where φ_β is the policy minimizing $\mathbb{E}V_\beta(\widetilde{Y}-D)$, $\widetilde{X}_0=\widetilde{X}$, and $\widetilde{X}_n$ is the inventory position n periods later, when policy δ is used and random variable σ is any upper bound on $\sigma^\beta := \inf\{n \geqslant 0, \beta\underline{V}_\beta(\widetilde{X}_n) \leqslant m_\beta\}$.

The following theorem is the main result of this section.

Theorem 10. *There exists a stationary policy δ^* that is average optimal in the sense that $\Phi(\delta^*, x) := \limsup_{n\to\infty} (1/n)\mathbb{E}_{\widetilde{X},\delta^*}[\sum_{m=0}^{n-1} L(\widetilde{X}_m)] = \inf_{\widetilde{X}\in S}\inf_\delta \Phi(\delta, \widetilde{X}) =: g$ for all $\widetilde{X}$. δ^* is limit discount optimal in that for any $\widetilde{X}\in S$ for all $\beta \nearrow 1$ such that $\delta^*(\widetilde{X}) = \lim_{\beta\nearrow 1}\delta^*_\beta(\widetilde{X})$. Additionally, $g = \lim_{\beta\nearrow 1}(1-\beta)m_\beta = \lim_{\beta\nearrow 1}(1-\beta)V_\beta$.*

The underlying idea of the proof is to show that for any starting point and any discount factor β, the extra cost for not starting at the "best" point is bounded and the bound is independent of the discount factor. The proof of this theorem consists of several steps. Utilizing Lemma 3, we establish a finite upper bound on the relative cost difference on its right-hand side. The bound is determined by constructing an alternative policy that deliberately visits every point within a finite subset of the state space in which the optimal base-stock levels are guaranteed to reside. The cost of the policy, while reaching each of the points, is an upper bound on the cost until the optimal point is reached. We show that the cost of this policy is finite and that the bound is independent of β. The complete proof is in the online appendix at http://or.pubs.informs.org/pages/collect.html.

As for the discounted cost infinite-horizon section, numerical evaluation of the average cost may be achieved using the value-iteration algorithms (see, for example, Puterman 1994), justified by the results in this section. Markov-modulated demand requires a simple modification, adding the state of the underlying Markov chain, and a connectiveness requirement similar to that in Kapuscinski and Tayur (1998).

6. Conclusions

We have analyzed a two-echelon supply chain with capacity constraints. Under quite general conditions—lower capacity at the lower installation—we have shown that the cost function can be decomposed into echelon-dependent components, which leads to full characterization of the optimal inventory policy. The optimal policy, the modified echelon base-stock (MEBS) policy, is straightforward to describe; the lower installation attempts to reach a base-stock level, if possible, as it is constrained by capacity, K_1, and availability of stock at its immediate supplier, installation 2. Echelon 2 also attempts to reach a base-stock level, but is limited by installation 1's capacity. That is, the optimal echelon ordering decisions are restricted to the "band," $\mathscr{S}$. The result of this is that there are "induced penalty functions" applied to the system once the value function is decomposed. These induced penalty functions are analogous to those of Clark and Scarf (1960), except that each installation receives a separate function. Echelon 1 accrues an induced cost for potentially limiting the system by possessing the bottleneck operation. Echelon 2 accrues an induced cost for potentially not providing sufficient materials to keep its immediate customer sufficiently stocked.

We extend this structural result to models incorporating lead times, Markov-modulated demand (which may have a variety of nonstationary demand processes imbedded into it), and to other system configurations. We also extend the main result to the infinite-time horizon for discounted-cost and average-cost criteria. This is done without significant restrictions upon the demand process.

Appendix. Additional Proofs

Lemma 1. *Let $K_1 \leqslant K_j$. For any $\widetilde{X}_n$ all optimal $\widetilde{Y}_n$ satisfy $y_n^j \leqslant \max(K_1, x_n^j - a_n^{j-1})$ for $j > 1$.*

Proof. Assume that there exists an optimal policy π, such that for certain n and j, $y_n^j > \max(K_1, x_n^j - a_n^{j-1})$. Without loss of generality, we choose minimal n and j, i.e., assume that n is the shortest horizon for which there exists such a j, and for that n, j is the smallest among the candidate installations. Let π' be an alternative policy such that $a_n^{j'} = a_n^j - 1$ and $a_{n-1}^{j'} = a_{n-1}^j + 1$, but otherwise follows policy π. Clearly, $x_{n-1}^{(j+1)'} = x_{n-1}^{j+1} + 1$ and $x_{n-1}^{j'} = x_{n-1}^j - 1$.

Because π is optimal, it is feasible. Before comparing the costs of the two policies, we need to check the feasibility of π'. It is easy to justify that the sufficient conditions are: (i) $a_n^{j'} \leqslant K_j$, (ii) $a_n^{j'} \leqslant x_n^{j+1}$, (iii) $a_{n-1}^{j'} \leqslant x_{n-1}^{(j+1)'}$, (iv) $a_{n-1}^{j-1} \leqslant x_{n-1}^{j'}$, and (v) $a_{n-1}^{j'} \leqslant K_j$.

Because $a_n^{j'} < a_n^j \leqslant K_j$, we get (i). Because $a_n^{j'} < a_n^j \leqslant x_n^{j+1}$, we get (ii).

(iii) $a_{n-1}^{j'} = a_{n-1}^j + 1 \leqslant x_{n-1}^{j+1} + 1 = x_{n-1}^{(j+1)'}$.

(iv) $x_{n-1}^j = y_n^j > K_1$ implies $x_{n-1}^{j'} = x_{n-1}^j - 1 \geqslant K_1 \geqslant a_{n-1}^{j-1}$.

(v) $x_{n-2}^j = x_{n-1}^j - a_{n-1}^{j-1} + a_{n-1}^j = x_{n-2}^{j'} = x_{n-1}^{j'} - a_{n-1}^{j-1} + a_{n-1}^{j'}$ and (from (iv)), $a_{n-1}^{(j-1)'} \leqslant x_{n-1}^{j'}$, so $x_{n-1}^{j'} - a_{n-1}^{(j-1)'} \geqslant 0$. Thus, $a_{n-1}^{j'} \leqslant K_1 \leqslant K_j$.

The difference in cost between the policies is

$$\begin{aligned}\mathrm{Cost}(\pi) - \mathrm{Cost}(\pi') &= \left(\sum_{i=j}^N h_i\right)(x_{n-1}^j - x_{n-1}^{j'}) \\ &\quad + \left(\sum_{i=j+1}^N h_i\right)(x_{n-1}^{j+1} - x_{n-1}^{(j+1)'}) \\ &= \left(\sum_{i=j}^N h_i\right)(a_n^j) + \left(\sum_{i=j+1}^N h_i\right)(-a_n^{j'}) \\ &= h_j(1) > 0,\end{aligned}$$

which is a contradiction with optimality of π. □

Theorem 2. *The modified echelon base-stock policy (MEBS), with parameters (z_n^1, z_n^2) defined above, is optimal for an $N = 2$ system where $K_1 \leqslant K_2$ and $X^2 - X^1 \leqslant K_1$.*

Proof. Theorem 1 states that the value function of the system is separable into individual value functions, each dependent upon the echelon starting inventory. It also states that each of these separate value functions is convex with (unrestricted) minima at ζ_n and z_n^2. In addition, Corollary 1(a) states that for these conditions, $Y^2 - Y^1 \leqslant K_1$ for optimal Y^1 and Y^2. There are two possible cases: (a) $z_n^2 - \zeta_n \leqslant K_1$ and (b) $z_n^2 - \zeta_n > K_1$.

(a) In this case, the global minimizing point is within the band $\mathscr{S}$ and $(z_n^1, z_n^2) = (\zeta_n, z_n^2)$. Depending upon the starting inventory, this point may not always be achievable. If $X^1 \leqslant z_n^1 \leqslant X^2$, then we would order to z_n^1; if z_n^1 lies beyond these limits, we would order the closer limit due to the aforementioned convexity. Likewise, if $X^2 \leqslant z_n^2 \leqslant Y^1 + K_1$, we would order up to z_n^2. However, if z_n^2 lay outside these limits, we would choose the closer limit due to the convexity of the function. Therefore, MEBS is optimal for case (a).

(b) $z_n^2 - \zeta_n > K_1$. Due to the result of Corollary 1(a), $Y^2 \leqslant Y^1 + K_1$. Now, because $z_n^2 - \zeta > K_1$, and due to the joint convexity and separability, there will exist z_n^1 such that $\zeta_n \leqslant z_n^1 \leqslant z_n^2 - K_1$, which minimizes the value function along $(Y^1, Y^1 + K_1)$. The minimizing z_n^1 occurs within this interval using the following logic. J_n is convex decreasing in $Y^1 \leqslant \zeta_n$ and convex decreasing in $Y^2 \leqslant \zeta_n + K_1$, implying that $J_n(Y^1, Y^1 + K_1)$ is convex decreasing in $Y^1 \leqslant \zeta_n$. Similarly, J_n is convex increasing in $Y^1 \geqslant z_n^2 - K_1$ and convex increasing in $Y^2 \geqslant z_n^2$, implying that $J_n(Y^1, Y^1 + K_1)$ is convex increasing in $Y^1 \geqslant z_n^2 - K_1$. When $X^1 \leqslant z_n^1 \leqslant X^2$, the minimizing point $(z_n^1, z_n^1 + K_1)$ is reachable. When $X^2 < z_n^1$, the convexity implies that the point $(X^2, X^2 + K_1)$ gives the lowest cost under the constraint set. (z_n^1, z_n^2) may not always be achievable; this point will not be achievable if $z_n^1 < z_n^2 - K_1$. If $X^1 > z_n^1$, the lower installation will order nothing, $Y^1 = X^1$; echelon 2 will order up to either $X^1 + K_1$ or z_n^2, whichever is the smaller, due to the convexity. Thus, MEBS is optimal for case (b). □

Theorem 6. (a) *Consider an N-stage system without capacity limits at stages $j < N - 1$ and finite capacities $K_{N-1} \leqslant K_N$. Under the inventory condition $X_n^N - X_n^{N-1} \leqslant K_{N-1}$, $V_n(\widetilde{X}_n) = \sum_{j=1}^N V_n^j(X_n^j)$. The optimal policy is as follows: MEBS holds for the capacitated installations, and the remaining installations follow an echelon base-stock policy.*

(b) *Consider an N-stage system without capacity limits at stages $j < N$ and finite capacity K_N. For this system, $V_n(\widetilde{X}_n) = \sum_{j=1}^N V_n^j(X_n^j)$. An echelon base-stock policy is optimal for this system.*

Proof. (a) The form of this proof follows similarly to that of Theorem 1, but the bottleneck is found at echelon $N-1$. Invoking Lemma 2, the sequence of decomposition of the value function begins from echelon 1. The proof mirrors that of Clark and Scarf (1960) up until echelon $N-2$. At this point, the value functions for these installations are similar to those in Theorem 1, and the proof is identical from this point onward. We utilize Corollary 1(a) for these final two installations, because the decomposition of installation $N-2$ from installations $N-1$ and N merely adds a convex function to the value function for the final two installations. (Essentially, it is shown that the decomposed dynamic programs for all echelons contain induced penalty functions for potentially not supplying enough material, and echelon $N-1$ has the additional induced penalty cost for limiting the ability of echelon N to reach a desirable inventory level as a result of the capacity limitation, K_{N-1}.) The optimality of the inventory policy follows from the fact that the convexity of each of the individual value functions of the lowest $N-2$ installations implies that an echelon base-stock level is desirable, identical to that of Clark and Scarf (1960). The optimal policy of the final two installations follows identically from Theorems 1 and 2.

(b) The proof of this element of the theorem is even simpler because the imposition of a capacity limitation upon echelon N can be incorporated into the original Clark and Scarf (1960) proof without any problem. This can be viewed from the perspective of the constraints upon the echelon order-up-to decision variables, Y^j for $j \in \{1, 2, \ldots, N\}$. The upper bound of each decision variable except the highest echelon is the stock availability at the next higher echelon; that is, $X^j \leqslant Y^j \leqslant X^{j+1}$ for $j < N$. The limitation upon echelon N is the capacity limitation, $X^N \leqslant Y^N \leqslant X^N + K_N$. The proof continues as in Clark and Scarf (1960), decomposing the dynamic program from the lowest echelon to the highest. Once the highest echelon is reached, the resulting dynamic program is a convex operand being minimized with the decision variable Y^N bounded between X^N and $X^N + K_N$; that is, it is a function of X^N only. Consequently, the structure of the optimal policy is identical to that of Clark and Scarf (1960), namely an echelon base-stock policy, although the actual base-stock levels will differ from those of Clark and Scarf. □

All the proofs for Technical Lemmas A1, A2, A3, *and* A4 *appear on the Operations Research website at http://or.pubs.informs.org/pages/collect.html.*

Lemma A1. *Consider two functions, J^A and J^B, each jointly convex and separable in their variables, Y^1 and Y^2, which satisfy*

$$\frac{\partial}{\partial Y^1} J^B(\widetilde{Y}) \leqslant \frac{\partial}{\partial Y^1} J^A(\widetilde{Y}) \quad \text{and} \quad \frac{\partial}{\partial Y^2} J^B(\widetilde{Y}) \leqslant \frac{\partial}{\partial Y^2} J^A(\widetilde{Y}).$$

If $K^B < K^A$, then $z^{1A} \leqslant z^{1B}$.

Lemma A2. *Let an $N = 2$ supply chain with $K_2 \geqslant K_1$ operate under a base-stock policy with levels $\tilde{z}$ such that $z^2 - z^1 = K_1$, and let $\widetilde{X} \leqslant z$. If the lower echelon base-stock level is reached, then the upper echelon can also achieve its base-stock level in the same period.*

LEMMA A3. *Consider an N-stage serial supply chain operating under a MEBS policy, where $z^j - z^{j-1} = K_1$ and $K_j \geqslant K_1$ for all $j > 1$. Assume for a given period $X^j - X^{j-1} = K_1$ for all $j > 1$. Then,* (a) *$X^j - X^{j-1} = K_1$ for all $j > 1$ in all future periods, and* (b) *once $z^1 - K_1 \leqslant X^1 \leqslant z^1$, then $\tilde{z}$ is reachable in a single period for all echelons.*

LEMMA A4. *Consider an N-stage serial multiechelon system with capacity limits of $K_i \geqslant K_1$ at stage i and $0 < \mathbb{E}[D] < K_1$. Consider the echelon base-stock policy, $\pi(K)$, ordering up to $(i-1)K_1$ at echelon i. Define T as the random variable representing the number of periods between subsequent visits to these base-stock levels. Then,* (a) *$\mathbb{E}[T] < \infty$, and* (b) *there exists $A \in \Re_+$, such that $J_n(0, K_1, 2K_1, \ldots, (N-1)K_1) \leqslant An$.*

THEOREM 8. (a) *There exists a finite upper bound on inventory targets for $V(\widetilde{X})$, which is independent of $\beta \in (0,1]$, and* (b) *there exists a finite upper bound on inventory targets for $\mathbb{E}V(\widetilde{Y} - D)$, which is independent of $\beta \in (0,1]$. An inventory target is any $Y(X) > X$ (i.e., an up-to level where we actually increase inventory).*

PROOF. We will show that for any finite-horizon problem and for any $\beta \in (0,1]$, the finite-horizon up-to levels are uniformly bounded. The proof for part (b) is similar to that of part (a) below, but the modified policy defined below would operate $\pi(K)$ for a different number of periods related to different target inventory levels; however, the principle of operating for $l(k)$ periods will be the same but the actual number will differ. This proof follows the standard type of proof seen in Kapuscinski and Tayur (1998). Based on Lemma A4, there exists a policy, $\pi(K)$, with base-stock levels $(0, K_1, 2K_1, \ldots, (N-1)K_1)$, such that there is a finite expected number of periods until this base-stock vector is revisited (regeneration).

We demonstrate that there is a finite upper bound on the optimal base-stock values by contradiction. Assume that $\limsup \tilde{y}_n^* = \infty$. Therefore, an increasing sequence $n_k \in N$ such that $y_{n_k}^* \to \infty$ and $y_{n_k}^* > y_n^*$ for all $n < n_k$. The cost of following this optimal policy will be compared to an alternative policy. The alternative, or modified policy, π_m, is described by operating $\pi(K)$ for $l(k)$ periods, followed by π, the optimal policy. Let

$$l(k) := \left\lfloor \frac{y_{n_k}^{N*} - (N-1)K_1}{\mathbb{E}D} \right\rfloor,$$

where $\lfloor\ \rfloor$ denotes rounding down to the nearest integer. Let m denote the random number of periods defined as

$$m = \min\{i \in \mathbf{N} \mid \tilde{y}_{n_k - l(k) - i}(\pi) = \tilde{y}_{n_k - l(k) - i}(\pi_m)\}$$

or $n_k - l(k)$ if such an i does not exist. Using Wald's theorem, we can bound $\mathbb{E}m$,

$$\mathbb{E}m \leqslant \frac{\sum_{i=1}^N y_{n_k}^{i*}}{K_1 - \mathbb{E}D}.$$

Now consider the cost of policy π for the $l(k)$ periods. Given that the cost function is convex, we can employ Jensen's inequality, so

$$\begin{aligned}
&\sum_{i=1}^{l(k)} \mathbb{E}L_{n_k - i}(x_{n_k - i}) \\
&\quad \geqslant \sum_{i=1}^{l(k)} \left[\left(\sum_{j=1}^N h_j\right)\left(y_{n_k}^{1*} - \sum_{j=n_k-i+1}^{n_k} \mathbb{E}D\right) \right. \\
&\qquad \left. + \sum_{j=2}^N \left(\sum_{p=j}^N h_p\right)(y_{n_k}^{j*} - y_{n_k}^{(j-1)*}) \right] \\
&\quad = \sum_{i=1}^{l(k)} \left[\sum_{j=1}^N h_j \left(y_{n_k}^{1*} - \sum_{j=n_k-i+1}^{n_k} \mathbb{E}D\right) + \sum_{j=2}^N h_j y_{n_k}^{j*} \right] \\
&\quad \geqslant l(k) \sum_{j=2}^N h_j y_{n_k}^{j*} \qquad \text{(A1)} \\
&\quad \geqslant \left(\frac{y_{n_k}^{N*} - (N-1)K_1}{\mathbb{E}D} - 1 \right) \sum_{j=2}^N h_j y_{n_k}^{j*} \\
&\quad = \left(\frac{y_{n_k}^{N*}}{\mathbb{E}D} - \frac{(N-1)K_1 + \mathbb{E}D}{\mathbb{E}D} \right) \sum_{j=2}^N h_j y_{n_k}^{j*} \\
&\quad = \frac{h_N}{\mathbb{E}D} (y_{n_k}^{N*})^2 + \sum_{j=2}^{N-1} \frac{h_j}{\mathbb{E}D} y_{n_k}^{j*} y_{n_k}^{N*} \\
&\qquad - \frac{(N-1)K_1 + \mathbb{E}D}{\mathbb{E}D} \sum_{j=2}^N h_j y_{n_k}^{j*}. \qquad \text{(A2)}
\end{aligned}$$

Inequality (A1) is justified from (a) the definition of $l(k)$, and (b) Corollary 1(a). Specifically, it is shown in Corollary 1(a) that the difference between ending inventory levels at neighboring echelons higher than echelon 1 will not be greater than K_1. This is then reflected in the definition of $l(k)$, where $N-1$ units of K_1 are subtracted from echelon N's target level; this difference is ensured to be lower than $y_{n_k}^{1*}$, thus justifying (A1).

Let us now focus on (A2). Because $y_{n_k}^{1*} \leqslant y_{n_k}^{2*} \leqslant \cdots \leqslant y_{n_k}^{(N-1)*} \leqslant y_{n_k}^{N*}$ (due to state space), the term growing most quickly as k increases is the $(y_{n_k}^{N*})^2$ term. The coefficient of this term is positive because $h_N > 0$ and $0 < \mathbb{E}D < K_1$. We now have an expression that acts as a lower bound on operating π for $l(k)$ periods. The cost of operating $\pi(K)$ for n periods is bounded above by An (Lemma A4).

Given that policy π_m acts as policy π from period $l(k)+1$ to period $l(k)+m$, the process operating under policy π_m is the same process as the process operating under policy π starting from different initial states in period $l(k)+1$. Glasserman and Tayur (1994) demonstrate that this process is Harris ergodic if the stability condition $\mathbb{E}D < K_1 \leqslant \min_i K_i$ (our notation) is satisfied, which our model does. Consequently, our model admits *coupling*, a direct result of the regenerative structure of a Harris ergodic Markov chain. This means that the processes starting from different initial states will coincide after a finite (random) time; that is, $m < \infty$. The cost of operating under either π_m or π during these m periods is, therefore, finite.

The cost of operating policy π_m during periods $l(k)+1$ to $l(k)+m$ consists of penalty and holding costs. Because π will be conducted during this interval, the penalty costs will not be larger than those of policy $\pi(K)$. Similarly, the inventory levels at echelon 1 achieved during this interval will be lower than the echelon 1 inventories realized from the original policy, the cost of which forms a bound on the echelon 1 holding cost. The bound for the holding costs at higher echelons can be regarded as a constant in every period: Corollary 1(a) demonstrates that the optimal policy will require no more than K_1 units at these higher installations and, thus, $\sum_{i=2}^{N}\sum_{j=1}^{N} h_j K_1 = (N-1)\sum_{j=1}^{N} h_j K_1$ serves as a bound in each period. Therefore, the cost of operating π_m during the first $l(k)+m$ periods is bounded by

$$A(l(k)+\mathbb{E}m)+(N-1)\sum_{j=1}^{N} h_j K_1 \mathbb{E}m$$

$$\leqslant A'(l(k)+\mathbb{E}m) \leqslant A'\left(\sum_{i=1}^{N} y_{n_k}^{i*}\left(\frac{1}{\mathbb{E}D}+\frac{1}{K_1-\mathbb{E}D}\right)\right). \tag{A3}$$

Consequently, the difference in costs between policies π and π_m over period $l(k)+m$ is at least (A2) minus (A3), which is quadratic in $y_{n_k}^{N*}$. So, as k increases, this term dominates and the cost difference goes to infinity. This suggests that π_m is less costly than π, contradicting our initial assumption. □

Endnotes

1. Clark and Scarf (1960) assign ordering decisions to the end of the current period instead of the beginning of the following one. This is equivalent to our sequence, which is more commonly used.
2. Porteus (2002, p. 137) shows a more general result. If a function f changes sign $j < n$ times, when the Pólya random variable of order n translates f resulting in $g(x) := \mathbb{E}f(x-X)$, then g changes sign at most j times.
3. An exception to this is when the optimal base-stock levels fall more quickly than a particular demand realization (e.g., when the end of the horizon approaches or when the demand is nonstationary).
4. The cost at the best operating level is defined as $V^* := \min_{\widetilde{X}\in\mathcal{S}} V(\widetilde{X})$.

Acknowledgments

The authors thank the area editor and two referees for their substantive comments, which helped them to improve this paper. They also appreciate comments from Yao Zhao and Rich Metters. This research was partially supported by the Tauber Manufacturing Institute of the University of Michigan.

References

Aviv, Y., A. Federgruen. 1997. Stochastic inventory models with limited production capacity and periodically varying parameters. *Probab. Engrg. Inform. Sci.* **11** 107–135.

Axsäter, S., K. Rosling. 1993. Installation vs. echelon stock policies for multilevel inventory control. *Management Sci.* **39** 1274–1280.

Axsäter, S., K. Rosling. 1999. Ranking of generalized multi-stage KANBAN policies. *Eur. J. Oper. Res.* **113** 560–567.

Bertsekas, D. 1995. *Dynamic Programming and Optimal Control*, Vols. I and II. Athena Scientific, Belmont, MA.

Bertsekas, D., S. Shreve. 1996. *Stochastic Optimal Control: The Discrete Time Case.* Athena Scientific, Belmont, MA.

Buzacott, J. A., J. G. Shanthikumar. 1992. A general approach for coordinating production in multiple-cell manufacturing systems. *Production Oper. Management* **1**(1) 34–52.

Chen, F., J.-S. Song. 2001. Optimal policies for multi-echelon inventory problems with Markov-modulated demand. *Oper. Res.* **4**(2) 226–234.

Chen, F., Y.-S. Zheng. 1994. Lower bounds for multi-echelon stochastic inventory systems. *Management Sci.* **40**(11) 1426–1443.

Clark, A. J., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* **6** 475–490.

Dallery, Y., G. Liberopoulos. 2000. Extended kanban control system: Combining kanban and base-stock. *IIE Trans.* **32** 369–386.

Federgruen, A. 1993. Centralized planning models for multi-echelon inventory systems under uncertainty. S. C. Graves, A. H. G. Rinnooy Kan, P. H. Zipkin, eds. *Logistics of Production and Inventory.* Elsevier Science Publishers B.V., Amsterdam, The Netherlands.

Federgruen, A., P. Zipkin. 1984. Computational issues in an infinite-horizon multiechelon inventory model. *Oper. Res.* **32**(4) 818–836.

Federgruen, A., P. Zipkin. 1986a. An inventory model with limited production capacity and uncertain demands I: The average-cost criterion. *Math. Oper. Res.* **11**(2) 193–207.

Federgruen, A., P. Zipkin. 1986b. An inventory model with limited production capacity and uncertain demands II: The discounted-cost criterion. *Math. Oper. Res.* **11**(2) 208–215.

Federgruen, A., P. Schweitzer, H. Tijms. 1983. Denumerable undiscounted decision processes with unbounded rewards. *Math. Oper. Res.* **8**(2) 298–314.

Glasserman, P., S. Tayur. 1994. The stability of a capacitated, multi-echelon production-inventory system under a base-stock policy. *Oper. Res.* **42**(5) 913–925.

Glasserman, P., S. Tayur. 1995. Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Sci.* **41**(2) 263–281.

Heyman, D., M. Sobel. 1984. *Stochastic Models in Operations Research*, Vol. II. McGraw Hill, New York.

Kapuscinski, R., S. Tayur. 1995. A capacitated production-inventory model with periodic demand. Working paper, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA.

Kapuscinski, R., S. Tayur. 1998. A capacitated production-inventory model with periodic demand. *Oper. Res.* **46**(6) 899–911.

Karush, W. 1959. A theorem in convex programming. *Naval Res. Logist. Quart.* **6**(3) 245–260.

Liberopoulos, G., Y. Dallery. 2000. A unified framework for pull control mechanisms in multi-stage manufacturing systems. *Ann. Oper. Res.* **93** 325–355.

Metters, R. 1997. Production planning with stochastic seasonal demand and capacitated production. *IIE Trans.* **29** 1017–1029.

Muharremoglu, A., J. N. Tsitsiklis. 2003. A single-unit decomposition approach to multi-echelon inventory systems. Working paper, MIT, Cambridge, MA.

Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory.* Stanford University Press, Stanford, CA.

Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming.* John Wiley and Sons, New York.

Rosling, K. 1989. Optimal inventory policies for assembly systems under random demands. *Oper. Res.* **37** 565–579.

Roundy, R. O., J. A. Muckstadt. 2000. Heuristic computation of periodic-review base stock inventory policies. *Management Sci.* **46**(1) 104–109.

Schäl, M. 1993. Average optimality in dynamic programming with general state space. *Math. Oper. Res.* **18**(1) 163–172.

Sennott, L. I. 1989. Average cost optimal stationary policies in infinite state Markov decision processes with unbounded costs. *Oper. Res.* **37**(4) 626–633.

Sennott, L. I. 1999. *Stochastic Dynamic Programming and the Control of Queueing Systems.* John Wiley and Sons, New York.

Simchi-Levi, D., Y. Zhao. 2001. The value of information sharing in a two-stage supply chain with production capacity constraints: The infinite horizon case. *Probab. Engrg. Inform. Sci.*, forthcoming.

Speck, C. J., J. van der Wal. 1991a. The capacitated multi-echelon inventory system with serial structure: 1. The "push ahead"-effect. Memorandum COSOR 91-39, Eindhoven University of Technology, Eindhoven, The Netherlands, 1–10.

Speck, C. J., J. van der Wal. 1991b. The capacitated multi-echelon inventory system with serial structure: 2. An average cost approximation method. Memorandum COSOR 91-40, Eindhoven University of Technology, Eindhoven, The Netherlands, 1–14.

Veatch, M. H., L. M. Wein. 1994. Optimal control and a two-station tandem production/inventory system. *Oper. Res.* **42**(2) 337–350.

Veinott, A. F. 1966. The status of mathematical inventory theory. *Management Sci.* **12**(11) 745–777.

Zipkin, P. 1989. Critical number policies for inventory models with periodic data. *Management Sci.* **35**(1) 71–80.