94%-Effective Policies for a Two-Stage Serial Inventory System with Stochastic Demand

Fangruo Chen
Graduate School of Business, Columbia University, New York, New York 10027
fc26@columbia.edu

A two-stage inventory system is considered where Poisson demand occurs at Stage 1, and Stage 1 replenishes its inventory from Stage 2, which in turn orders from an outside supplier with unlimited stock. Each shipment, either to Stage 2 or to Stage 1, incurs a fixed setup cost. Under the assumption that the supply leadtime at Stage 2 is zero, we characterize a simple heuristic policy whose long-run average cost is guaranteed to be within 6% of optimality, i.e., a 94%-effective policy. The paper also provides heuristic policies for more general inventory systems and reports computational results.
(*Multi-Echelon Inventory*; *Stochastic Demand*; *Worst Case Analysis*; *Heuristic Policy*)

1. Introduction

It is now well known that for a wide range of multi-echelon, *deterministic* inventory systems, there exist policies of exceedingly simple structure whose performance is guaranteed to come within 2% of optimality (see, e.g., Maxwell and Muckstadt 1985 and Roundy 1985, 1986). These policies are usually referred to as power-of-two policies: Each item at each facility is replenished at constant (item- or facility-specific) intervals that are power-of-two multiples of a base period. Power-of-two policies are easy to implement, and an optimal power-of-two policy is easy to compute, even in production-distribution networks with a complex topology and/or general setup cost structures (see, e.g., Federgruen et al. 1992).

However, for multi-echelon, *stochastic* inventory systems, the state of the art is less satisfactory. The field started about three decades ago, when Clark and Scarf (1960, 1962) published their seminal papers. In the second paper, they attempted to characterize an optimal policy for a two-stage, serial system with economies of scale at both stages. The conclusion was that even for this simple system, the optimal policy, if one exists, must be extremely complex. This observation has driven subsequent research on multi-echelon, stochastic inventory systems with economies of scale to focus on heuristic policies. So far, no heuristic policy in any such model has been identified with a guaranteed worst-case optimality gap.

For years, researchers in the field of multi-echelon, stochastic inventory systems have hoped to establish heuristic policies with a guaranteed, worst-case performance that is still close to optimal. The above-mentioned recent and quite remarkable achievements for deterministic systems have only heightened this desire. Unfortunately, the task is extremely difficult and progress has been slow. This paper represents a small step towards the ultimate goal.

We consider the following two-stage inventory model. Poisson demand occurs at Stage 1, and Stage 1 replenishes its inventory from Stage 2, which in turn orders from an outside supplier with unlimited stock. There are economies of scale at both stages for placing orders. For this model, a simple heuristic policy is identified whose long-run average cost comes within

0025-1909/99/4512/1679$05.00
1526-5501 electronic ISSN

6% of optimality in the worst case, although the optimal policy itself is unknown. This policy is also called 94%-effective. The result is obtained by theoretically bounding the relative gap between the long-run average cost of the heuristic and a lower bound on the long-run average costs of all feasible policies. This represents, to our knowledge, the first worst-case analysis of a multi-echelon, stochastic inventory system.

The above result depends upon two key assumptions. First, we assume that the leadtime from the outside supplier to Stage 2 is negligible. This assumption allows us to strengthen an existing lower bound on the long-run average costs of all feasible policies, a critical step in the worst-case analysis. It is a reasonable assumption when the supplier is physically close to Stage 2 and guarantees a high service level. However, we allow the leadtime from Stage 2 to Stage 1 to be any nonnegative constant. The second assumption is that the discrete units of inventories can be approximated by continuous variables. This approximation simplifies the analysis significantly. It is also a common approach in the literature. For example, the EOQ formula is often used even when the order quantity must be discrete.

The 94%-effective policy has a simple structure. Stage 2 follows an echelon-stock (R, Q) policy: It orders a fixed quantity Q_2 from the outside supplier whenever its echelon stock reaches a reorder point r_2. Upon arrival of the order, a shipment is sent to Stage 1 to increase its inventory position to its target level $r_1 + Q_1$ $(\leq r_2 + Q_2)$. Moreover, $(r_2 + Q_2) - (r_1 + Q_1)$ is chosen to be an integer multiple m of Q_1. The inventory maintained at Stage 2, which has size mQ_1, is sent in batches of size Q_1 to Stage 1 each time its inventory position drops to a reorder point r_1 $(> r_2)$, i.e., Stage 1 follows the (r_1, Q_1) policy m times. (Thus there are $m + 1$ shipments in each order cycle, including the shipment to Stage 1 made when an order arrives at Stage 2.) Thereafter, Stage 2's inventory is depleted. When the system inventory level drops to r_2, the next order will be placed with the outside supplier, and the process repeats itself.

The paper proceeds to consider two more general inventory systems. One is the above two-stage system with a positive leadtime at Stage 2; the other allows multiple outlets at the lower echelon, i.e., one-warehouse multiretailer systems. For the former, we propose a heuristic policy based on the structure of the above 94%-effective policy. Under this heuristic, Stage 2 still follows an echelon-stock (R, Q) policy. When a batch arrives at Stage 2, a *dynamic shipping schedule* is determined based on the state of the system upon arrival. The schedule specifies how the incoming batch at Stage 2 is going to be shipped to Stage 1. (Different batches may have different shipping schedules.) It is designed to capture two essential features of the 94%-effective policy: (i) each shipment to Stage 1 increases its inventory position to the *same* level; and (ii) the first of these shipments is at least as large as the rest. We call this heuristic policy the Dynamic Shipping Policy or DSP. Numerical examples show that the DSP is cost-effective when compared with an existing heuristic policy. On the other hand, a similar heuristic is proposed for one-warehouse multiretailer systems. It has several attractive features which are absent in all of the heretofore proposed policies for such systems.

There is an extensive literature on various heuristic policies in different multi-echelon, stochastic inventory systems. Examples include Sherbrooke (1968), Eppen and Schrage (1981), Deuermeyer and Schwarz (1981), Federgruen and Zipkin (1984a, b), De Bodt and Graves (1985), Graves (1985), Moinzadeh and Lee (1986), Lee and Moinzadeh (1987a, b), Jackson (1988), Svoronos and Zipkin (1988, 1991), Axsater (1990, 1993a, b), and Chen and Zheng (1994a, 1997, 1998). Axsater (1993c) and Federgruen (1993) provide comprehensive reviews of this literature. Although most of the proposed heuristic policies make intuitive sense, it is unclear how suboptimal they can be. So far, all we have is some numerical evidence; see, e.g., Federgruen and Zipkin (1984b) and Chen and Zheng (1994a, b, 1997). But a numerical study, no matter how extensive it is, only makes a *posterior* statement on the cost-effectiveness of a heuristic policy; it is never clear whether the observations from the numerical study carry over to other instances. To date, Atkins and De (1992) appear to have made the only serious attempt at worst-case analysis of a multi-echelon, stochastic inventory system. They considered a two-stage model much like the model considered in this paper, but with a positive leadtime at the upper stage. They provided a lower bound, a heuristic, and some results on the gap between the two. Their lower bound is weaker than ours. Their heuristic is different. Most importantly, their analysis is incomplete due to several logical problems with their arguments.

A key building block of this paper is a recent result by Zheng (1992) for the single-stage (R, Q) model. It states that the long-run average cost of the (R, Q) model is insensitive to the choice of Q, assuming that the optimal value of R corresponding to each value of Q is used. This result is an extension of the well-known fact that the cost of the EOQ model is insensitive to variations of Q around its optimum (see Hadley and Whitin 1963). Zheng's result is as essential to this paper as the insensitivity result of the EOQ model is to, say, Roundy's work.

Six sections follow this introduction. Section 2 presents the two-stage model together with preliminary results. Section 3 establishes a lower bound on the long-run average costs of all feasible policies. Section 4 identifies a class of feasible policies and, for each policy in this class, provides an upper bound on its long-run average cost. Section 5 characterizes a 94%-effective policy. Section 6 allows a positive leadtime at Stage 2 of the two-stage model. Section 7 deals with one-warehouse multiretailer systems.

2. Preliminaries

2.1. Model and Notation

Consider the following two-stage inventory model. Customer demand arrives at Stage 1 according to a simple Poisson process with constant rate λ. Stage 1 is replenished by Stage 2, which in turn orders from an outside supplier with unlimited stock. Each shipment to Stage i incurs a fixed setup cost K_i, $i = 1, 2$. The leadtime from the outside supplier to Stage 2 is zero, and the leadtime from Stage 2 to Stage 1 is a constant $L \geq 0$. Let $h_i > 0$ be the echelon holding cost rate at Stage $i = 1, 2$. When Stage 1 runs out of stock, the excess demand is backlogged. The backorder cost rate is p. The planning horizon is infinite, and the objective is to minimize the long-run average systemwide cost.

For any time t, define

$I_1(t)$ = on-hand inventory at Stage 1,
$B(t)$ = backorder level at Stage 1,
$IL_1(t)$ = inventory level at Stage 1 $= I_1(t) - B(t)$,
$IP_1(t)$ = inventory position at Stage 1
= $IL_1(t)$ plus inventories in transit to Stage 1,
$I_2(t)$ = echelon inventory at Stage 2
= on-hand inventory at Stage 2 plus inventories in transit to, or on-hand at, Stage 1,
$IL_2(t)$ = echelon inventory level at Stage 2, or system inventory level
$= I_2(t) - B(t)$.

At time t, the systemwide holding and backorder costs accrue at rate (with t suppressed)

$$h_2 I_2 + h_1 I_1 + pB = h_2 IL_2 + h_1 IL_1 + (p + H_1)B$$

where $H_1 = h_1 + h_2$ is the installation holding cost rate at Stage 1. We call $h_2 IL_2$ the echelon holding cost, and $h_1 IL_1 + (p + H_1)B$ the holding and backorder costs at Stage 1.

2.2. Characterizing an Optimal Policy

Two decisions are made at any point in time: How much, if any, to order from the outside supplier and how much, if any, to ship to Stage 1. These decisions will be referred to as *ordering* and *shipping* decisions, respectively. We will restrict decision epochs to demand epochs, i.e., ordering and shipping decisions are made immediately after each demand arrival. Due to the memoryless property of the Poisson process, this restriction causes no loss of generality.

We next introduce a convention for charging holding and backorder costs. Let t_i be the arrival time of the ith demand. Let $IP_1(t_i) = y_1$ and $IL_2(t_i) = y_2$. (The inventory variables are assessed just after the decision epoch t_i.) Since the system inventory level stays at y_2 over the time interval $[t_i, t_{i+1})$ and the expected length of the interval is $1/\lambda$, the expected total echelon holding cost in the interval is $h_2 y_2/\lambda$.

Now consider the holding and backorder costs at Stage 1. Note that for any time t,

$$IL_1(t + L) = IP_1(t) - D$$

where D is the total demand in $(t, t + L]$. Therefore, as far as Stage 1's inventory level is concerned, the shipping decision at time t_i takes effect only after time $t_i + L$ and the next shipping decision (at time t_{i+1}) will not take effect until time $t_{i+1} + L$. In other words, the shipping decision at time t_i determines the expected holding and backorder costs at Stage 1 in the interval $[t_i + L, t_{i+1} + L)$, which are denoted by $g_1(y_1)$. From Federgruen and Schechner (1983),

$$g_1(y_1) = E[h_1(y_1 - D)^+ + (p + h_2)(y_1 - D)^-]/\lambda$$

where $(x)^+ = \max\{x, 0\}$ and $(x)^- = \max\{-x, 0\}$. The following costs are charged to t_i:

$$h_2 y_2/\lambda + g_1(y_1).$$

A policy is *nested* if whenever Stage 2 receives an order from the outside supplier, it sends a shipment to Stage 1. Since the leadtime at Stage 2 is zero, any optimal policy must be nested. The reason is that any order at Stage 2 can be delayed until just before the next shipment to Stage 1, saving holding costs at Stage 2. We will consider only nested policies.

Consider any nested policy. Suppose that Stage 2 places an order at time 0. This order is received immediately. Due to nestedness, a shipment is sent to Stage 1 at time 0. Call this the first shipment. Let $IL_2(0) = S_2$ and $IP_1(0) = S_1^1$. Let n be the number of shipments before the next order is placed. Let r_1^{i-1} (resp., S_1^i) be Stage 1's inventory position just before (resp., after) the ith shipment, $i = 2, \ldots, n$. Let r_1^n be Stage 1's inventory position, and r_2 the system inventory level, just before the next order. Since the on-hand inventory at Stage 2 is, by definition, always nonnegative, $r_1^n \le r_2$.

A cycle is the time interval between two consecutive orders at Stage 2. During the cycle beginning at time 0, the system inventory level decreases from S_2 to $r_2 + 1$ in unit step sizes. Similarly, from the ith to the $(i + 1)$th shipment, Stage 1's inventory position decreases from S_1^i to $r_1^i + 1$ in unit step sizes, $i = 1, \ldots, n$. (The $(n + 1)$st shipment is the first shipment in the next order cycle.) Therefore the expected total cost in the cycle is

$$K_2 + \sum_{y=r_2+1}^{S_2} h_2 y/\lambda + \sum_{i=1}^{n}\left(K_1 + \sum_{y=r_1^i+1}^{S_1^i} g_1(y)\right).$$

Define $Q_2 = S_2 - r_2$ and $Q_1^i = S_1^i - r_1^i$, $i = 1, \ldots, n$. Note that Q_2 is the total demand in the cycle and Q_1^i is the total demand between the ith shipment and the next. Therefore, the average expected cost *per unit of demand* in the cycle is

$$\frac{K_2 + \sum_{y=r_2+1}^{r_2+Q_2} h_2 y/\lambda + \sum_{i=1}^{n}\left(K_1 + \sum_{y=r_1^i+1}^{r_1^i+Q_1^i} g_1(y)\right)}{Q_2}.$$

The minimum value of the above expression over $(r_2, Q_2, n, r_1^i, Q_1^i, i = 1, \ldots, n)$ is thus a lower bound on the average expected cost per unit of demand of any nested policy. Multiplying this lower bound by the demand arrival rate λ, one has a lower bound on the long-run average cost (*per unit of time*) of any nested policy. More specifically, this lower bound is the minimum value of the objective function in

$$\mathbf{P_0}: \quad \text{MIN} \quad \frac{\lambda K_2 + \sum_{y=r_2+1}^{r_2+Q_2} h_2 y + \sum_{i=1}^{n}\left(\lambda K_1 + \sum_{y=r_1^i+1}^{r_1^i+Q_1^i} G_1(y)\right)}{Q_2}$$

$$\text{ST:} \quad \sum_{i=1}^{n} Q_1^i = Q_2$$

$$r_1^i < r_1^{i+1} + Q_1^{i+1}, \qquad i = 1, \ldots, n-1,$$

$$r_1^n \le r_2,$$

$$Q_2 \ge 1, \quad Q_1^i \ge 1, \quad i = 1, \ldots, n,$$

$$n \ge 1, \quad \text{all variables integral},$$

where $G_1(y) = \lambda g_1(y)$. The constraints follow by definition. For example, the first constraint says that the total demand in a cycle is the sum of the demands in the n disjoint intervals between consecutive shipments, and the second constraint follows because Stage 1's inventory position increases after each shipment. Note that any optimal solution to $\mathbf{P_0}$ must have $r_1^n = r_2$ since if $r_1^n < r_2$ then the objective function value can be reduced by decreasing r_2 while fixing all the other decision variables. In other words, Stage 2 should have zero on-hand inventory when it orders. This is often called the zero-inventory-ordering property in the literature.

Figure 1 A Cyclic Policy

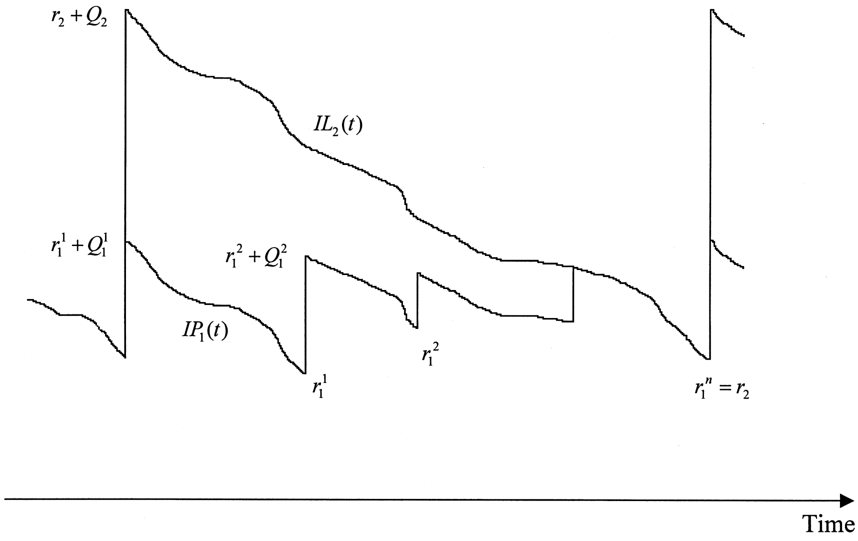

Let $\mathbf{r_1} = (r_1^1, \ldots, r_1^n)$ and $\mathbf{Q_1} = (Q_1^1, \ldots, Q_1^n)$. Any feasible solution $(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1})$ to $\mathbf{P_0}$ specifies a feasible policy which operates as follows. Suppose the initial state of the system is that Stage 2 has zero on-hand inventory and Stage 1's inventory position is r_2. (This state can always be reached: If the initial system inventory level is less than r_2, place an order to increase it to r_2 and send the entire order to Stage 1; otherwise, ship all the units at Stage 2, if any, to Stage 1 and wait until Stage 1's inventory position reaches r_2.) Whenever the system inventory level drops to r_2, Stage 2 orders Q_2 units from the outside supplier. This order is shipped to Stage 1 in n shipments. The ith shipment raises Stage 1's inventory position from r_1^{i-1} to $r_1^i + Q_1^i$, $i = 1, \ldots, n$, with $r_1^0 = r_2$. This policy will be referred to as a *cyclic policy*; see Figure 1 for an illustration. Note that a cyclic policy is nested and has the zero-inventory-ordering property. Furthermore, its long-run average cost is exactly equal to the objective function of $\mathbf{P_0}$. Therefore, an optimal solution to $\mathbf{P_0}$ leads to a cyclic policy that is optimal since the minimum value of the objective function is also a lower bound on the long-run average costs of all feasible policies. The goal of this paper, however, is not to find such an optimal solution, but to seek a simple feasible solution that is close to optimal.

For tractability, we assume that the variables $(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1})$ are *continuous*. We will focus on the following approximation of $\mathbf{P_0}$:

$$\mathbf{P}: \quad \text{MIN} \quad C(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1})$$

$$\text{ST:} \quad \sum_{i=1}^{n} Q_1^i = Q_2$$

$$r_1^i < r_1^{i+1} + Q_1^{i+1}, \qquad i = 1, \ldots, n-1,$$

$$r_1^n = r_2,$$

$$Q_2 > 0, \quad Q_1^i > 0, \quad i = 1, \ldots, n,$$

$$n \geq 1, \quad \text{integer},$$

where

$$C(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1}) = \frac{\lambda K_2 + \int_{r_2}^{r_2+Q_2} h_2 y\, dy + \sum_{i=1}^{n} \left(\lambda K_1 + \int_{r_1^i}^{r_1^i+Q_1^i} G_1(y)\, dy\right)}{Q_2}.$$

Remark. Note that replacing $\sum_{y=r_2+1}^{r_2+Q_2} h_2 y$ by $\int_{r_2}^{r_2+Q_2} h_2 y dy$ leads to an underestimate of the objective function by a constant $\frac{1}{2}$. This step does not really represent an approximation. In fact, if this

Figure 2 **Illustration of Z2, Z3, and Z4**

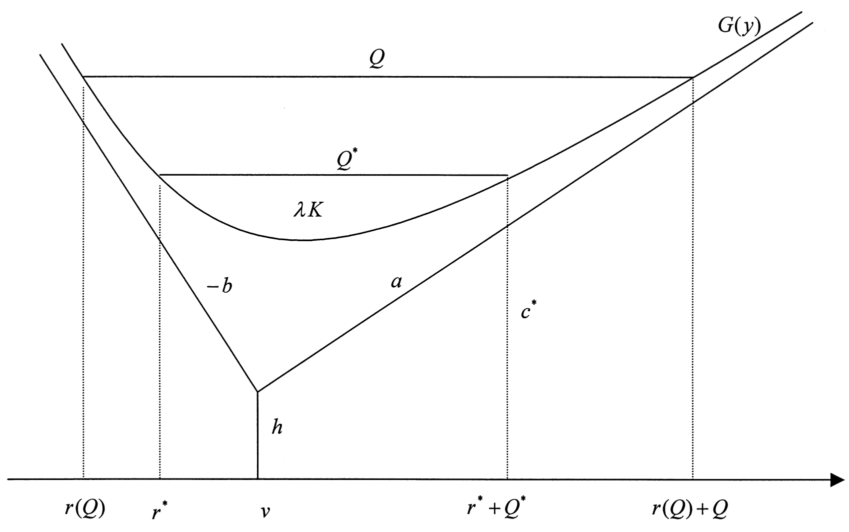

were the only change from $\mathbf{P}_0$ to $\mathbf{P}$, then the worst-case bound (to be developed later) on the relative gap between the long-run average cost of a heuristic and a lower bound on the long-run average costs of all policies would be *conservative*, since increasing both the cost of the heuristic and the lower bound by the same amount only *decreases* their relative gap. On the other hand, the replacement of $\sum_{y=r_1^i+1}^{r_1^i+Q_1^i} G_1(y)$ by $\int_{r_1^i}^{r_1^i+Q_1^i} G_1(y)\,dy$ does represent an approximation. But the effect should be small for large values of the Q's. This type of approximation is quite common in the inventory literature; see Zheng (1992) for a fuller discussion on this point.

2.3. Existing Results

For any $Q > 0$ and r, define

$$C(r, Q) = \frac{\lambda K + \int_r^{r+Q} G(y)\,dy}{Q}$$

where $K > 0$ and $G(\cdot)$ satisfies the following three conditions:

(i) $G(\cdot)$ is convex;

(ii) $\exists a > 0$, $b > 0$, $h \geq 0$ and v so that $G(y) \geq a(y - v)^+ + b(y - v)^- + h \stackrel{\text{def}}{=} G_d(y)$, $\forall y$;

(iii) $G_d(\cdot)$ is an asymptote of $G(\cdot)$, i.e., $\lim_{y \to \pm\infty}[G(y) - G_d(y)] = 0$.

Define $C(Q) = C(r(Q), Q) = \min_r C(r, Q)$ and $C^* = C(r^*, Q^*) = C(r(Q^*), Q^*) = \min_{r,Q} C(r, Q)$. If there are multiple solutions to the problem $\min_r C(r, Q)$, then let $r(Q)$ be the largest such solution. This convention will be used throughout the paper.

The following results are useful for later developments:

(Z1) $C(r, Q)$ is jointly convex in r and Q;

(Z2) $G(r(Q)) = G(r(Q) + Q)$;

(Z3) $G(r^*) = G(r^* + Q^*) = C^*$;

(Z4) $\lambda K = \int_{r^*}^{r^*+Q^*} [C^* - G(y)]\,dy$;

(Z5) $C(Q)/C^* \leq \epsilon(Q/Q^*)$ where $\epsilon(q) = (q + q^{-1})/2$.

These results can be obtained from Zheng (1992). Note that the above conditions (i), (ii), and (iii) are slightly more general than those given in Zheng. Moreover, the function G encountered in this paper is not differentiable at every point, unlike in Zheng. Therefore, his original proofs need to be modified. This can be done easily. We omit the details here; they can be obtained from the author upon request. Also note that (Z1) first appeared in Zipkin (1986) and later in Zhang (1996), who provided a simpler proof. See Figure 2 for an illustration of (Z2), (Z3), and (Z4).

3. Lower Bound

We first introduce three single-stage (r, Q) models, which are then used to specify a lower bound on the

Figure 3 **Illustration of the Induced-Penalty Cost**

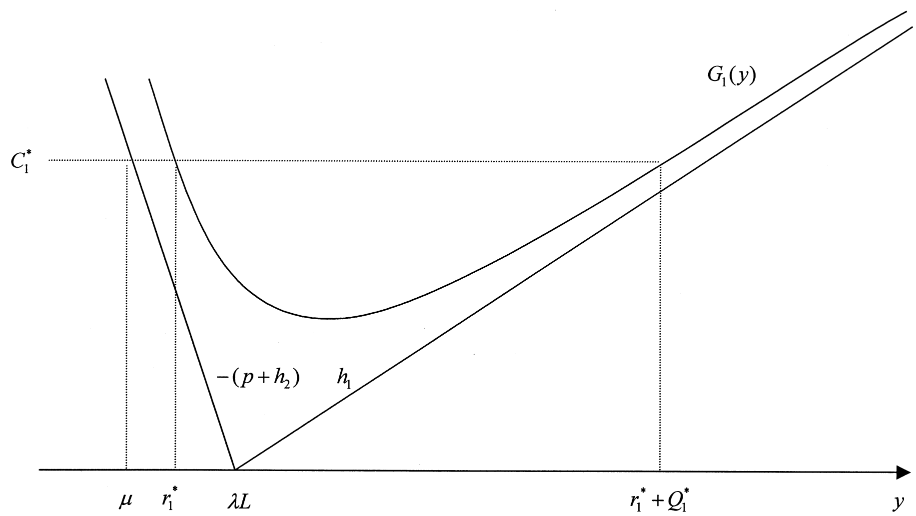

long-run average costs of all feasible policies in the two-stage model.

3.1. Stage-1 Model

Suppose Stage 2 has infinite on-hand inventory. In this case, Stage 1 reduces to the standard single-stage model with setup cost K_1 and holding-backorder cost rate function $G_1(\cdot)$. The long-run average cost of an (r, Q) policy in this model is

$$C_1(r, Q) = \frac{\lambda K_1 + \int_r^{r+Q} G_1(y)\, dy}{Q}.$$

Note that

$$G_1(y) = E[h_1(y - D)^+ + (p + h_2)(y - D)^-]$$
$$\geq h_1(y - \lambda L)^+ + (p + h_2)(y - \lambda L)^-,$$

where the inequality is Jensen's. Thus $G_1(\cdot)$ satisfies (i) and (ii). The third condition can also be verified easily. Consequently, (Z1)~(Z5) apply. Define $C_1(Q) = C_1(r_1(Q), Q) = \min_r C_1(r, Q)$ and $C_1^* = C_1(r_1^*, Q_1^*) = C_1(r_1(Q_1^*), Q_1^*) = \min_{r,Q} C_1(r, Q)$.

3.2. Stage-2 Model

Let y be the system inventory level. Define

$$G_{12}(y) = \begin{cases} G_1(y) - C_1^* & \text{if } y < r_1^*, \\ 0 & \text{otherwise,} \end{cases}$$

which is often called the induced-penalty cost function. It can be interpreted as the cost increase at Stage 1 due to the lack of inventory at Stage 2: When $y < r_1^*$, Stage 1 can only increase its inventory position to y, and its holding and backorder costs accumulate at an expected rate $G_1(y)$, which is higher than the minimum cost rate C_1^* achievable if Stage 2 had ample stock. It is clear that $G_{12}(\cdot)$ is convex and nonincreasing (see Figure 3 for an illustration). By charging to Stage 2 the echelon holding cost and the induced-penalty cost, we have the holding-backorder cost rate function for Stage 2:

$$G_2(y) = h_2 y + G_{12}(y) + (1 - \alpha)C_1^*,$$

where $\alpha = p/(p + h_2)$. (The constant term is added so that (ii) is satisfied, as we will see shortly.) Now Stage 2 becomes the standard single-stage model with setup cost K_2 and holding-backorder cost rate function $G_2(\cdot)$. The long-run average cost of an (r, Q) policy in this model is

$$C_2(r, Q) = \frac{\lambda K_2 + \int_r^{r+Q} G_2(y)\, dy}{Q}.$$

Note that $G_2(\cdot)$ is convex, satisfying condition (i). Now consider conditions (ii) and (iii). Define $\mu = \lambda L - C_1^*/(p + h_2)$. From Figure 3, it is easy to see that

$G_{12}(y) \geq (p + h_2)(y - \mu)^-$ and that $(p + h_2)(y - \mu)^-$ is an asymptote of $G_{12}(y)$. Therefore,

$$G_2(y) \geq h_2 y + (p + h_2)(y - \mu)^- + (1 - \alpha)C_1^*$$
$$= h_2(y - \mu)^+ + p(y - \mu)^- + h_2\mu + (1 - \alpha)C_1^*$$
$$\overset{\text{def}}{=} G_{2d}(y).$$

Note that $h_2\mu + (1 - \alpha)C_1^* = h_2\lambda L \geq 0$. Thus condition (ii) is satisfied. Finally, $G_{2d}(\cdot)$ is an asymptote of $G_2(\cdot)$ because $(p + h_2)(y - \mu)^-$ is an asymptote of $G_{12}(y)$. Thus $G_2(\cdot)$ also satisfies condition (iii). As a result, (Z1)~(Z5) apply. Define $C_2(Q) = C_2(r_2(Q), Q) = \min_r C_2(r, Q)$ and $C_2^* = C_2(r_2^*, Q_2^*) = C_2(r_2(Q_2^*), Q_2^*) = \min_{r,Q} C_2(r, Q)$.

3.3. Combined Model

Suppose Stage 2 does not hold any inventory. That is, when an order is received at Stage 2 it is directly sent to Stage 1. In this case, the effective setup cost for each order is $K_1 + K_2$ and $IL_2(t) = IP_1(t)$ for all t. Thus, given $IL_2(t) = y$, the systemwide holding and backorder costs accrue at an expected rate equal to

$$G_0(y) \overset{\text{def}}{=} h_2 y + G_1(y).$$

The two-stage model reduces to the standard single-stage model with setup cost $K_1 + K_2$ and holding-backorder cost rate function $G_0(\cdot)$. The long-run average cost of an (r, Q) policy in this model is

$$C_0(r, Q) = \frac{\lambda(K_1 + K_2) + \int_r^{r+Q} G_0(y)\, dy}{Q}.$$

Note that $G_0(\cdot)$ is convex. It also satisfies (ii) since

$$G_0(y) = E[H_1(y - D) + (p + H_1)(y - D)^-] + h_2\lambda L$$
$$= E[H_1(y - D)^+ + p(y - D)^-] + h_2\lambda L$$
$$\geq H_1(y - \lambda L)^+ + p(y - \lambda L)^- + h_2\lambda L.$$

The third condition also holds here. Consequently, (Z1)~(Z5) apply. Define $C_0^* = C_0(r_0^*, Q_0^*) = \min_{r,Q} C_0(r, Q)$.

3.4. Lower Bound

Now we are ready to derive a lower bound on the long-run average costs of all feasible policies. This is achieved by deriving a lower bound on the minimum value of the objective function in **P**. The following lemma establishes a relationship between the objective function and the above single-stage models. (All omitted proofs can be obtained from the author upon request.)

Lemma 1. *For any feasible solution* $(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1})$ *to* **P**, *we have* (a) $C(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1}) \geq C_2(r_2, Q_2) + \alpha C_1^*$. *If* $r_2 + Q_2 \leq r_1^* + Q_1^*$ *then we also have* (b) $C(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1}) \geq C_0(r_2, Q_2)$.

Corollary 1. $C_0^* \geq C_2^* + \alpha C_1^*$.

Proof. First, note that C_0^* is the long-run average cost of a cyclic policy $(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1})$ with $n = 1$, $r_1^1 = r_2 = r_0^*$ and $Q_1^1 = Q_2 = Q_0^*$. Under this policy, Stage 2 does not hold any inventory and the two-stage system functions essentially as a single stage. The corollary follows from Lemma 1(a). □

Theorem 1. *B^* is a lower bound on the long-run average cost of any feasible policy, where*

$$B^* \overset{\text{def}}{=} \min\{C_0^*, \min_{r_2+Q_2 \geq r_1^*+Q_1^*} C_2(r_2, Q_2) + \alpha C_1^*\}.$$

Moreover,

$$B^* = \begin{cases} C_2^* + \alpha C_1^* & \text{if } r_2^* + Q_2^* > r_1^* + Q_1^*, \\ C_0^* & \text{otherwise.} \end{cases}$$

Proof. Take any feasible solution $(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1})$ to **P**. If $r_2 + Q_2 \geq r_1^* + Q_1^*$, we have from Lemma 1(a)

$$C(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1}) \geq \min_{r_2+Q_2 \geq r_1^*+Q_1^*} C_2(r_2, Q_2) + \alpha C_1^* \geq B^*.$$

Otherwise, if $r_2 + Q_2 \leq r_1^* + Q_1^*$, we have from Lemma 1(b)

$$C(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1}) \geq C_0^* \geq B^*.$$

The theorem follows by combining these cases.

To show the alternative form of B^*, we distinguish between two cases.

Case 1. $r_2^* + Q_2^* > r_1^* + Q_1^*$. Since

$$\min_{r_2+Q_2 \geq r_1^*+Q_1^*} C_2(r_2, Q_2) + \alpha C_1^* = C_2^* + \alpha C_1^*$$

we have $B^* = C_2^* + \alpha C_1^*$ from Corollary 1.

Case 2. $r_2^* + Q_2^* \le r_1^* + Q_1^*$. Let

$$\min_{r_2+Q_2\ge r_1^*+Q_1^*} C_2(r_2, Q_2) = C_2(\bar{r}_2, \bar{Q}_2).$$

Since $C_2(r_2, Q_2)$ is jointly convex in r_2 and Q_2 (see (Z1)), we have

$$\bar{r}_2 + \bar{Q}_2 = r_1^* + Q_1^*.$$

Now consider $C_2(r_2, r_1^* + Q_1^* - r_2)$ as a function of r_2. Since $G_2(y)$ is increasing for $y \ge r_1^*$ by definition, $C_2(r_2, r_1^* + Q_1^* - r_2)$ is increasing in r_2 for $r_2 \ge r_1^*$. Consequently, $\bar{r}_2 \le r_1^*$. Thus, from the definition of $G_{12}(\cdot)$,

$$\int_{\bar{r}_2}^{\bar{r}_2+\bar{Q}_2} (G_{12}(y) + C_1^*)\, dy = \int_{\bar{r}_2}^{r_1^*} G_1(y)\, dy + \int_{r_1^*}^{r_1^*+Q_1^*} C_1^*\, dy.$$

By (Z4), the right side of the above equation is equal to

$$\int_{\bar{r}_2}^{r_1^*} G_1(y)\, dy + \lambda K_1 + \int_{r_1^*}^{r_1^*+Q_1^*} G_1(y)\, dy = \lambda K_1 + \int_{\bar{r}_2}^{\bar{r}_2+\bar{Q}_2} G_1(y)\, dy.$$

Therefore,

$$C_2(\bar{r}_2, \bar{Q}_2) + \alpha C_1^* = \frac{\lambda K_2 + \int_{\bar{r}_2}^{\bar{r}_2+\bar{Q}_2} h_2 y\, dy + \int_{\bar{r}_2}^{\bar{r}_2+\bar{Q}_2} (G_{12}(y) + C_1^*)\, dy}{\bar{Q}_2}$$
$$= \frac{\lambda K_2 + \int_{\bar{r}_2}^{\bar{r}_2+\bar{Q}_2} h_2 y\, dy + \lambda K_1 + \int_{\bar{r}_2}^{\bar{r}_2+\bar{Q}_2} G_1(y)\, dy}{\bar{Q}_2}$$
$$= C_0(\bar{r}_2, \bar{Q}_2) \ge C_0^*.$$

Thus $B^* = C_0^*$ by definition. □

Remarks.

(1) Note that $C_2^* + \alpha C_1^*$ is the induced-penalty bound established in Chen and Zheng (1994b). From Corollary 1 and Theorem 1, B^* is stronger than the induced-penalty bound.

(2) If $r_2^* + Q_2^* \le r_1^* + Q_1^*$ then by Theorem 1, C_0^* is a lower bound on the long-run average cost of any feasible policy. But from the proof of Corollary 1, a feasible policy achieves C_0^*. This policy is optimal for this case.

(3) B^* is easy to compute. First, minimize $C_i(r, Q)$ to determine r_i^*, Q_i^* and C_i^*, $i = 1, 2$. Here one can use the efficient algorithm by Federgruen and Zheng (1992). If $r_2^* + Q_2^* > r_1^* + Q_1^*$ then $B^* = C_2^* + \alpha C_1^*$. Otherwise, minimize $C_0(r, Q)$ to obtain C_0^*, and $B^* = C_0^*$. This procedure is parallel to the one developed to compute a lower bound for the deterministic counterpart of the two-stage serial model. Suppose the demand at Stage 1 is deterministic and arrives continuously at a constant rate. To compute a lower bound for this model, we first compute two EOQs (or equivalently, reorder intervals), one for the upper stage and one for the lower stage. If the upper-stage EOQ is smaller than the lower-stage EOQ, then it is optimal to combine the two stages, and the lower bound is equal to the minimum cost of a combined EOQ model. It is interesting that for the deterministic system, the rank order of order quantities (or reorder intervals) determines whether or not it is optimal to combine; while for the stochastic system, it is the rank order of *order-up-to levels* that matters.

4. A Class of Cyclic Policies

This section defines a class of feasible policies, and then for each policy in this class, establishes an upper bound on its long-run average cost.

Consider the following feasible policy. It has four control parameters (r_2, Q_2, r_1, Q_1) so that

$$Q_1 > 0, \quad Q_2 > 0, \quad r_2 < r_1, \quad r_2 + Q_2 = r_1 + nQ_1$$

for some positive integer n. The system starts at time 0 with $IL_2(0) = IP_1(0) = r_2$. Whenever the system inventory level reaches r_2, Stage 2 orders Q_2 units from the outside supplier. Every incoming order to Stage 2 (of size Q_2) is sent to Stage 1 in n shipments: The first shipment raises Stage 1's inventory position from r_2 to $r_1 + Q_1$, and the $n - 1$ subsequent

Figure 4 **A Policy in $\mathcal{S}$**

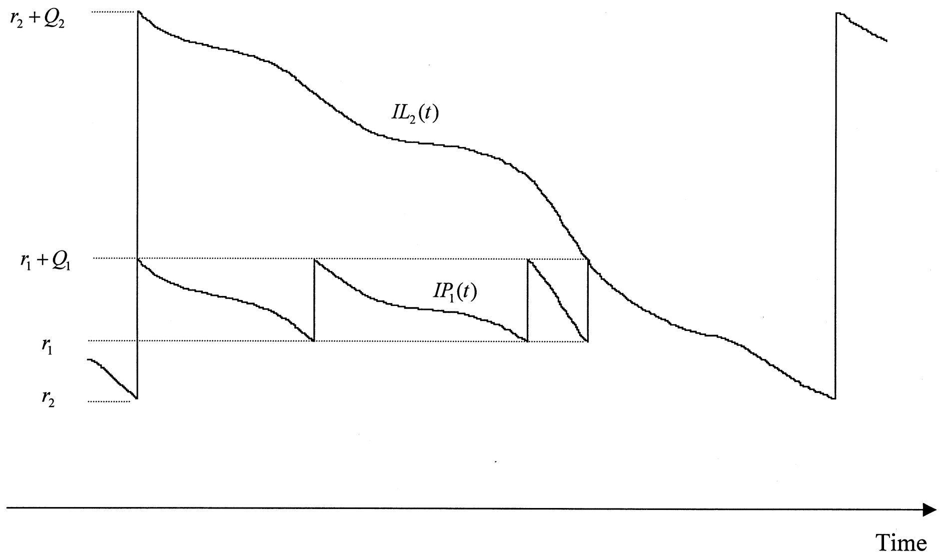

shipments each raise Stage 1's inventory position from r_1 to $r_1 + Q_1$. Thus the size of the first shipment is $r_1 + Q_1 - r_2$, and the size of each of the subsequent shipments is Q_1. We also require the policy parameters to satisfy

$$G_1(r_1) = G_1(r_1 + Q_1) \quad \text{and}$$
$$G_1(r_2) = h_2Q_2 + G_1(r_1). \tag{1}$$

Let $\mathcal{S}$ be the set of policies that satisfy the above conditions.

Note that for any given $Q_1 > 0$ and any positive integer n, there exist r_1, r_2 and Q_2 so that (r_2, Q_2, r_1, Q_1) is a member of $\mathcal{S}$. First, determine r_1 from the equation $G_1(r_1) = G_1(r_1 + Q_1)$. A solution always exists since $G_1(y)$ is convex and goes to infinity as $|y| \to +\infty$. (If there are multiple solutions, choose the largest one. That is, $r_1 = r_1(Q_1)$. Let us always follow this convention.) Second, determine r_2 so that $r_2 < r_1$ and

$$G_1(r_2) = h_2(r_1 + nQ_1 - r_2) + G_1(r_1),$$

which is equivalent to

$$G_0(r_2) = G_0(r_1) + h_2nQ_1.$$

Here a unique solution exists because $G_0(y)$ is convex and goes to infinity as $y \to -\infty$. Finally, set $Q_2 = r_1 + nQ_1 - r_2$, which is clearly positive. The resulting policy (r_2, Q_2, r_1, Q_1) is in $\mathcal{S}$. Figure 4 illustrates such a policy.

Now take any policy $(r_2, Q_2, r_1, Q_1) \in \mathcal{S}$. Let $C(r_2, Q_2, r_1, Q_1)$ be its long-run average cost. Since this policy is cyclic $(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1})$ with $r_1^i = r_1$ and $Q_1^i = Q_1$ for $i = 1, \ldots, n - 1$, $r_1^n = r_2$, and $Q_1^n = r_1 + Q_1 - r_2$, its long-run average cost is equal to the corresponding objective function of **P**. After some algebra, we have

$$C(r_2, Q_2, r_1, Q_1) = \frac{\lambda K_2 + \int_{r_2}^{r_2+Q_2} h_2y\,dy + \int_{r_2}^{r_1} [G_1(y) - C_1(Q_1)]\,dy}{Q_2} + C_1(Q_1). \tag{2}$$

Remarks.

(1) Using Lagrange multipliers to relax the two equality constraints in **P**, we have the following Lagrangian

$$\mathcal{L} = C(r_2, Q_2, \mathbf{r_1}, \mathbf{Q_1}) - \eta_1\left(\sum_{i=1}^{n} Q_1^i - Q_2\right) - \eta_2(r_1^n - r_2).$$

This leads to the following necessary conditions:

$$\frac{\partial \mathcal{L}}{\partial Q_1^i} = \frac{G_1(r_1^i + Q_1^i)}{Q_2} - \eta_1 = 0, \quad i = 1, \ldots, n,$$

$$\frac{\partial \mathcal{L}}{\partial r_1^i} = \frac{G_1(r_1^i + Q_1^i) - G_1(r_1^i)}{Q_2} = 0, \quad i = 1, \ldots, n-1,$$

$$\frac{\partial \mathcal{L}}{\partial r_1^n} = \frac{G_1(r_1^n + Q_1^n) - G_1(r_1^n)}{Q_2} - \eta_2 = 0,$$

$$\frac{\partial \mathcal{L}}{\partial r_2} = h_2 + \eta_2 = 0.$$

These conditions translate into the following equations:

$$G_1(r_1^i + Q_1^i) = G_1(r_1^j + Q_1^j), \qquad i, j = 1, \ldots, n,$$

and

$$G_1(r_1^i) = G_1(r_1^i + Q_1^i), \qquad i = 1, \ldots, n-1,$$

and

$$G_1(r_1^n) = h_2 Q_2 + G_1(r_1^n + Q_1^n).$$

These equations are all satisfied by the policies in $\mathcal{S}$.

(2) A policy in $\mathcal{S}$ induces a *nonstationary* flow of inventory from Stage 2 to Stage 1, since the first shipment to Stage 1 in a cycle is larger than the subsequent shipments. This has an intuitive explanation. Suppose Stage 1 has a backorder at some time t. If Stage 2 has inventory on hand at the same time, the backorder could have been avoided (by having an item at the right place at the right time), saving $p + h_2$. We can think of $p + h_2$ as the *effective* penalty cost when Stage 2 has on-hand inventory. On the other hand, the effective penalty cost is just p if Stage 2 has zero inventory. Since Stage 2 is more likely to run out of stock at the end of an order cycle, backlogging becomes *cheaper* toward the end of a cycle. As a result, the end-of-cycle backlog is larger than the previous ones. Satisfying this larger-than-normal backlog is why the first shipment in a cycle is larger than the subsequent shipments. Similar patterns have been observed in some deterministic inventory systems; see Mitchell (1987) and Atkins and Sun (1995). Moreover, maintaining a stationary flow from Stage 2 to Stage 1 can sometimes be very costly (Chen 1998).

Lemma 2. *For any* $(r_2, Q_2, r_1, Q_1) \in \mathcal{S}$,
(a) $r_2 + Q_2 \geq r_1 + Q_1 > r_1^*$;
(b) $r_2(Q_2) + Q_2 \geq r_1^* \geq r_2(Q_2)$;
(c) $(r_1 - r_2)/Q_2 \geq 1 - \alpha$.

Lemma 3. *Let $f(\cdot)$ be a convex and nonincreasing function. Take any $u_1 \leq u_2$. For any $v_1 < v_2$ with $f(u_1) - f(v_1) = f(u_2) - f(v_2)$,*

$$\int_{u_1}^{v_1} [f(u_1) - f(y)]\, dy \leq \int_{u_2}^{v_2} [f(u_2) - f(y)]\, dy.$$

Theorem 2. *For any* $(r_2, Q_2, r_1, Q_1) \in \mathcal{S}$, $C(r_2, Q_2, r_1, Q_1) \leq C_2(Q_2) + \alpha C_1(Q_1)$.

Proof. Note that

$$\int_{r_2}^{r_1} [G_1(y) - C_1(Q_1)]\, dy$$

$$= \int_{r_2}^{r_1} [G_1(y) - C_1^*]\, dy + (r_1 - r_2)[C_1^* - C_1(Q_1)]$$

$$\leq \int_{r_2}^{r_1} [G_1(y) - C_1^*]\, dy + (1 - \alpha) Q_2 [C_1^* - C_1(Q_1)]$$

where the inequality follows from Lemma 2(c) and the fact that $C_1^* - C_1(Q_1) \leq 0$. Using the above inequality in (2),

$$C(r_2, Q_2, r_1, Q_1) \leq C_2(Q_2) + \alpha C_1(Q_1) + (A + B)/Q_2,$$

where

$$A = \int_{r_2}^{r_1} [G_1(y) - C_1^*]\, dy - \int_{r_2}^{r_2+Q_2} G_{12}(y)\, dy,$$

and

$$B = \int_{r_2}^{r_2+Q_2} G_2(y)\, dy - \int_{r_2(Q_2)}^{r_2(Q_2)+Q_2} G_2(y)\, dy.$$

It suffices to show $A + B \leq 0$.

Consider A. Since $r_2 + Q_2 > r_1^*$ from Lemma 2(a) and $G_{12}(y) = 0$ for $y \geq r_1^*$ by definition,

$$\int_{r_2}^{r_2+Q_2} G_{12}(y)\, dy = \int_{r_2}^{r_1^*} G_{12}(y)\, dy.$$

Therefore,

$$\begin{aligned} A &= \int_{r_2}^{r_1^*} [G_1(y) - C_1^*]\, dy \\ &\quad + \int_{r_1^*}^{r_1} [G_1(y) - C_1^*]\, dy - \int_{r_2}^{r_1^*} G_{12}(y)\, dy \\ &= \int_{r_2}^{r_1^*} [G_1(y) - C_1^* - G_{12}(y)]\, dy \\ &\quad + \int_{r_1^*}^{r_1} [G_1(y) - C_1^*]\, dy. \end{aligned} \tag{3}$$

Now consider B. Note that

$$\begin{aligned} B &= \int_{r_2(Q_2)+Q_2}^{r_2+Q_2} G_2(y)\, dy - \int_{r_2(Q_2)}^{r_2} G_2(y)\, dy \\ &= \int_{r_2(Q_2)}^{r_2} G_2(y + Q_2)\, dy - \int_{r_2(Q_2)}^{r_2} G_2(y)\, dy \\ &= \int_{r_2(Q_2)}^{r_2} [G_2(y + Q_2) - G_2(y)]\, dy \\ &= \int_{r_2(Q_2)}^{r_2} [h_2 Q_2 + G_{12}(y + Q_2) - G_{12}(y)]\, dy, \end{aligned} \tag{4}$$

where the last equality follows from the definition of $G_2(\cdot)$. From Lemma 2(a, b) and the definition of $G_{12}(\cdot)$, $G_{12}(r_2 + Q_2) = 0$ and $G_{12}(r_2(Q_2) + Q_2) = 0$. Thus

$$\int_{r_2(Q_2)}^{r_2} G_{12}(y + Q_2)\, dy = 0, \tag{5}$$

and

$$\begin{aligned} h_2 Q_2 &= G_{12}(r_2(Q_2)) - G_{12}(r_2(Q_2) + Q_2) \\ &= G_{12}(r_2(Q_2)) \\ &= G_1(r_2(Q_2)) - C_1^*, \end{aligned} \tag{6}$$

where the first equality follows from $G_2(r_2(Q_2)) = G_2(r_2(Q_2) + Q_2)$ (Z2) and the definition of $G_2(\cdot)$ and the last equality follows from Lemma 2(b) and the definition of $G_{12}(\cdot)$. Using (5) and (6) in (4),

$$B = \int_{r_2(Q_2)}^{r_2} [G_1(r_2(Q_2)) - C_1^* - G_{12}(y)]\, dy. \tag{7}$$

Combining (3) and (7),

$$\begin{aligned} A + B &= \int_{r_2(Q_2)}^{r_2} G_1(r_2(Q_2))\, dy + \int_{r_2}^{r_1^*} G_1(y)\, dy \\ &\quad - \int_{r_2(Q_2)}^{r_1^*} [C_1^* + G_{12}(y)]\, dy \\ &\quad + \int_{r_1^*}^{r_1} [G_1(y) - C_1^*]\, dy \\ &= \int_{r_2(Q_2)}^{r_2} [G_1(r_2(Q_2)) - G_1(y)]\, dy \\ &\quad - \int_{r_1^*}^{r_1} [C_1^* - G_1(y)]\, dy \end{aligned} \tag{8}$$

since by definition, $C_1^* + G_{12}(y) = G_1(y)$ for $y \leq r_1^*$.

We next show that the right side of (8) is less than or equal to zero. Let y_0 be the largest minimum point of $G_1(\cdot)$. Since $G_1(y)$ is convex, it is nonincreasing for $y \leq y_0$. Since $G_1(r_1) = G_1(r_1 + Q_1)$ by definition and $G_1(r_1^*) = G_1(r_1^* + Q_1^*)$ from (Z3), we have $r_1 \leq y_0$ and $r_1^* \leq y_0$. Since $r_2 < r_1$ by definition and $r_2(Q_2) \leq r_1^*$ from Lemma 2(b), $r_2 \leq y_0$ and $r_2(Q_2) \leq y_0$. Moreover, since $C_1^* = G_1(r_1^*)$ from (Z3) and $G_1(r_2) - G_1(r_1) = h_2 Q_2$ from (1), we have from (6),

$$G_1(r_1) - G_1(r_1^*) = G_1(r_2) - G_1(r_2(Q_2)).$$

Figure 5 **The Case with $r^*_1 < r_1$**

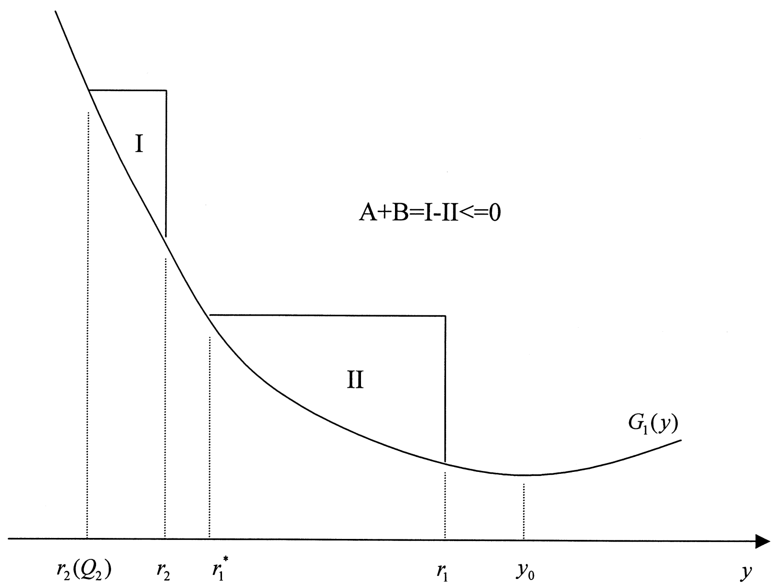

Now substitute $G_1(\cdot)$ for $f(\cdot)$ and $(r_2(Q_2), r_2, r_1^*, r_1)$ for (u_1, v_1, u_2, v_2) in Lemma 3, which implies $A + B \leq 0$ since $r_2(Q_2) \leq r_1^*$ (Lemma 2(b)). Figure 5 provides a graphic illustration for the case with $r_1^* < r_1$. □

5. 94%-Effective Policies

Suppose $r_1^* + Q_1^* < r_2^* + Q_2^*$. Following Roundy (1985), we define the *effectiveness* of a policy to be 100% times the ratio of the infimum of the long-run average cost over all policies to the long-run average cost of the policy in question. This section identifies a heuristic policy that is at least 94%-effective. (The case with $r_1^* + Q_1^* \geq r_2^* + Q_2^*$ has already been solved; see §3 for an optimal policy.)

We restrict attention to a subset of $\mathcal{S}$. Take any $Q_1 > 0$ and any integer $k \geq 0$. Let $n = 2^k$, a positive integer. Let $(r_2^k(Q_1), Q_2^k(Q_1), r_1(Q_1), Q_1)$ be the policy in $\mathcal{S}$ that corresponds to Q_1 and n. As mentioned in the previous section, such a policy always exists and its parameters are determined by solving

$$G_0(r_2^k(Q_1)) = G_0(r_1(Q_1)) + h_2 n Q_1$$

for $r_2^k(Q_1) < r_1(Q_1)$ and setting

$$Q_2^k(Q_1) = r_1(Q_1) + nQ_1 - r_2^k(Q_1).$$

Under the above policy, there are exactly n (a power-of-two integer) shipments in each order cycle.

Lemma 4.

(a) $Q_2^0(Q_1)$ *is nondecreasing in* $Q_1 > 0$;
(b) *For any* $Q_1 > 0$ *and any* $\beta \geq 1$, $Q_2^0(\beta Q_1) \leq \beta Q_2^0(Q_1)$;
(c) *For any* $Q_1 \in [\sqrt{0.5}Q_1^*, \sqrt{2}Q_1^*]$, $Q_2^0(Q_1) \leq \sqrt{2}Q_2^*$.

Lemma 5. *For any* $Q_1 > 0$, $Q_2^k(Q_1)$ *is increasing in* k *with* $Q_2^{k+1}(Q_1)/Q_2^k(Q_1) \leq 2$, $k = 0, 1, \ldots$.

Now take any $Q_1 \in [\sqrt{0.5}Q_1^*, \sqrt{2}Q_1^*]$. Let $m \geq 0$ be an integer so that

$$\sqrt{0.5}Q_2^* \leq Q_2^m(Q_1) \leq \sqrt{2}Q_2^*.$$

Lemma 4(c) and Lemma 5 ensure the existence of such an m.

Theorem 3. *The policy* $(r_2^m(Q_1), Q_2^m(Q_1), r_1(Q_1), Q_1)$ *is* 94%-*effective.*

Proof. *Let* $\alpha_1 = Q_1/Q_1^*$ and $\alpha_2 = Q_2^m(Q_1)/Q_2^*$. Thus $\alpha_i \in [\sqrt{0.5}, \sqrt{2}]$ and $\epsilon(\alpha_i) = (\alpha_i + \alpha_i^{-1})/2$

≤ 1.06, $i = 1, 2$. Since $(r_2^m(Q_1), Q_2^m(Q_1), r_1(Q_1), Q_1) \in \mathcal{S}$, we have

$$\begin{aligned} C(r_2^m(Q_1), Q_2^m(Q_1), r_1(Q_1), Q_1) &\leq C_2(Q_2^m(Q_1)) + \alpha C_1(Q_1) \quad \text{(Theorem 2)} \\ &\leq \epsilon(\alpha_2)C_2^* + \alpha\epsilon(\alpha_1)C_1^* \quad \text{(Z5)} \\ &\leq 1.06(C_2^* + \alpha C_1^*) \\ &= 1.06B^*. \quad \square \end{aligned}$$

Remark. One can search for the optimal Q_1 in the root-two interval of Q_1^* that minimizes $C_2(Q_2^m(Q_1)) + \alpha C_1(Q_1)$. However, it seems that $Q_2^m(Q_1)$ depends on Q_1 in a rather complex way. This prevents us from obtaining the 98% bound that has been achieved for many deterministic systems.

6. Positive Leadtime at Stage 2

Assume that there is a positive, constant leadtime at Stage 2. This section proposes a heuristic policy and uses numerical examples to demonstrate its performance.

The heuristic policy has two pairs of control parameters, (R_1, Q_1) for Stage 1 and (R_2, Q_2) for Stage 2, where $Q_2 \geq Q_1$ and $R_1 = r_1(Q_1)$. Like the 94%-effective policy, Stage 2 follows the echelon-stock (R_2, Q_2) policy: Whenever the system inventory position falls to R_2, order Q_2 units from the outside supplier. (The system inventory position is equal to the system inventory level plus the outstanding orders from the outside supplier.) For each batch arriving at Stage 2, an optimization problem is solved to determine how this batch is going to be shipped to Stage 1. This policy is called the DSP (Dynamic Shipping Policy).

To see how a shipping schedule is determined, consider an arbitrary batch arriving at Stage 2. Just before its arrival, let u be the on-hand inventory at Stage 2 and v the echelon inventory position at Stage 1 (i.e., IP_1).

(i) If $u > 0$ or $v \geq R_1$, then the batch is shipped to Stage 1 in m shipments where m is the largest integer less than or equal to Q_2/Q_1. Let r^i and Q^i, $i = 1, \ldots, m$, be a solution to the following problem:

$$\begin{aligned} \min \quad & \sum_{i=1}^{m} \sum_{y=r^i+1}^{r^i+Q^i} G_1(y) \\ \text{s.t.} \quad & \sum_{i=1}^{m} Q^i = Q_2, \\ & Q^i \geq 1, \ \text{integer}, \ i = 1, \ldots, m. \end{aligned}$$

This problem can be solved easily by dividing the Q_2 units into m sub-batches as equally as possible. The size of the ith shipment is Q^i, and it is sent to Stage 1 when IP_1 reaches r^i. Note that each Q^i is at least as large as Q_1. Thus $r^i \leq R_1$ and $r^i + Q^i \geq R_1 + Q_1$ for all i.

(ii) If $u = 0$ and $v < R_1$, then solve the following problem to determine r^i and Q^i, $i = 1, \ldots, n$, where n is also a decision variable:

$$\begin{aligned} \min \quad & \sum_{i=1}^{n} \left(\lambda K_1 + \sum_{y=r^i+1}^{r^i+Q^i} G_1(y) \right) \\ \text{s.t.} \quad & \sum_{i=1}^{n} Q^i = Q_2, \\ & Q^i \geq 1, \quad \text{integer}, \quad i = 1, \ldots, n, \\ & r^1 \leq v, \\ & n \geq 1 \quad \text{integer}. \end{aligned}$$

The batch arriving at Stage 2 is split into n sub-batches. The ith sub-batch has size Q^i and is sent to Stage 1 when IP_1 reaches r^i, $i = 1, \ldots, n$. If $n = 1$ and $r^1 = v$ then the entire batch is shipped to Stage 1 upon its arrival.

To test the performance of the above DSP, we compared it with an existing policy in numerical examples. Consider the echelon-stock (R, Q) policy, which also has two pairs of control parameters, (R_1, Q_1) for Stage 1 and (R_2, Q_2) for Stage 2, where Q_2 must be a positive integer multiple of Q_1. Stage 2 orders from the outside supplier according to the (R_2, Q_2) policy based on the system inventory position. Whenever Stage 1's inventory position falls to or below R_1, Stage 2 ships a minimum integer multiple of Q_1 to Stage 1 to raise its inventory position to above R_1. If Stage 2 does not have sufficient on-hand stock, then ship as much as possible. Note that a key difference between the DSP and the echelon-stock (R, Q) policy is that the former *dynamically* customizes a shipping schedule for each batch arriving at Stage 2.

Let L_i be the leadtime at Stage i, $i = 1, 2$. The model has eight parameters: λ, L_1, L_2, K_1, K_2, h_1, h_2, and p. Fix $L_1 = L_2 = 0.01$, $h_1 = 10$, and $p = 100$. Taking all combinations of $\lambda = 5, 50, 500, 5000$, $K_1 = 25, 30$, and $(K_2, h_2) = (0.01, 0.001), (0.1, 0.01), (1, 0.1), (10, 1)$, we have a total of 32 examples.

For each example, we used an algorithm in Chen and Zheng (1998) to determine an optimal power-of-two echelon-stock (R, Q) policy where both Q_1 and Q_2 are restricted to integer powers of two. For the DSP, we used the following simple procedure to determine its parameters. First, minimize

$$\frac{\lambda K_1 + \sum_{y=R+1}^{R+Q} G_1(y)}{Q}$$

over R and Q. Let $r_1(Q)$ be the optimal R for the given Q. Let Q_1^* be the optimal value of Q. For each value of Q_1, let $R_1 = r_1(Q_1)$ and

$$R_2 = R_1 + \lambda L_2 \quad \text{and} \quad Q_2 = 2\sqrt{\lambda L_2} + mQ_1,$$

where $m = \sqrt{(K_2/h_2)/(K_1/h_1)}$. Since both R_2 and Q_2 are integers, try rounding up and down to obtain the best combination. The parameters for the DSP are obtained via a search over Q_1 in the neighborhood of Q_1^*. The above choice of R_2 and Q_2 is quite intuitive. Note that λL_2 is both the mean and variance of the leadtime demand at Stage 2, which is a Poisson random variable. The above value of R_2 ensures, on average, that a delivery to Stage 2 coincides with a shipment to Stage 1 (i.e., nestedness). On the other hand, the value of Q_2 is based on the following hypothetical scenario. The leadtime demand at Stage 2 is two standard deviations *above* its mean. As a result, when a batch arrives at Stage 2, it sees $IP_1 = R_1 - 2\sqrt{\lambda L_2}$. In response, Stage 2 sends $(2\sqrt{\lambda L_2} + Q_1)$ units to Stage 1 immediately to raise IP_1 to $R_1 + Q_1$. The remaining $(m - 1)$ sub-batches, each of size Q_1, are sent to Stage 1 to raise IP_1 from R_1 to $R_1 + Q_1$. Note that under this hypothetical scenario, the DSP has exactly the same structure as the 94%-effective policy obtained earlier for a simpler model. Finally, the value of m is based on a deterministic, two-stage model; see, e.g., Muckstadt and Roundy (1993).

Table 1 summarizes the computational results. The long-run average cost of each DSP was obtained by simulation with its 95% confidence half-interval in the parentheses. The relative deviation between the two policies is defined as

$$\text{Rel Dev} = \frac{(R, Q)\ \text{cost} - \text{DSP cost}}{\text{DSP cost}}.$$

Notice that the DSP outperforms the (R, Q) policy for most of the examples with savings as high as 6.54%. For three examples, the DSP is slightly worse. This is likely due to the above heuristic procedure for determining the DSP parameters. It remains to be seen if the DSP is 94%-effective.

7. One-Warehouse Multiretailer Systems

This section considers a common distribution system with one central warehouse and multiple retailers where random demands arise. This system is more general than the (basic) serial system in that the upper stage (the warehouse) serves to replenish multiple retailers instead of a single outlet and that the order leadtime at the upper stage may be positive.

To design an effective replenishment policy for the above system, the insights from this paper alone are not enough. They should be combined with lessons from the deterministic literature. In serial systems with deterministic demand, it is optimal to apply a nested policy, i.e., orders are placed only when a shipment to the downstream facility is planned (see, e.g., Schwarz 1973). In systems with multiple outlets replenished by a single source, this is no longer the case. In fact, Roundy (1985) shows that a nested policy can be arbitrarily bad. Following Roundy, we partition the retailers in our system into two sets, R^+ and R^-, where R^+ (resp., R^-) contains retailers with average shipping frequencies larger (resp., smaller) than the warehouse's average ordering frequency. (A plausible

Table 1 **DSP and Echelon-Stock** (R, Q) **Policy**

					DSP					Echelon-Stock (R, Q) Policy					
No.	Lambda	$K1$	$K2$	$h2$	$R1$	$Q1$	$R2$	$Q2$	SimCost	$R1$	$Q1$	$R2$	$Q2$	Cost	Rel Dev (%)
1	5	25	0.01	0.001	−1	5	−1	10	46.0147 (0.0595)	−1	4	1	8	47.1713	2.51
2	5	25	0.1	0.01	−1	5	−1	10	46.0724 (0.0493)	−1	4	1	8	47.2767	2.61
3	5	25	1	0.1	−1	5	−1	10	46.9466 (0.0370)	−1	4	0	8	48.2426	2.76
4	5	25	10	1	−1	5	−1	10	55.4179 (0.0367)	−1	4	−1	8	57.3680	3.52
5	5	30	0.01	0.001	−1	6	−1	12	50.6853 (0.0685)	−1	8	1	8	53.9665	6.47
6	5	30	0.1	0.01	−1	6	−1	12	50.7866 (0.0679)	−1	8	0	8	54.0675	6.46
7	5	30	1	0.1	−1	6	−1	12	51.6521 (0.0591)	−1	8	0	8	55.0311	6.54
8	5	30	10	1	−1	6	−1	12	60.3701 (0.0596)	−1	8	−1	8	63.9069	5.86
9	50	25	0.01	0.001	−1	16	−1	33	152.9192 (0.1031)	−1	16	2	32	152.4565	−0.30
10	50	25	0.1	0.01	−1	16	−1	33	153.1360 (0.1320)	−1	16	1	32	152.7584	−0.25
11	50	25	1	0.1	−1	16	−1	33	155.9734 (0.1811)	−1	16	0	32	155.6690	−0.20
12	50	25	10	1	−1	16	−1	33	183.5978 (0.1818)	−1	16	−1	32	183.6211	0.01
13	50	30	0.01	0.001	−2	18	−2	37	167.8313 (0.1494)	−1	16	2	32	168.0815	0.15
14	50	30	0.1	0.01	−2	19	−2	39	167.9428 (0.2285)	−1	16	1	32	168.3834	0.26
15	50	30	1	0.1	−2	18	−2	37	170.8751 (0.2779)	−1	16	0	32	171.2940	0.25
16	50	30	10	1	−2	18	−2	37	198.0188 (0.1865)	−1	16	−1	32	199.2461	0.62
17	500	25	0.01	0.001	0	53	5	110	482.9642 (0.4103)	−1	64	13	128	490.5026	1.56
18	500	25	0.1	0.01	0	53	5	110	484.0272 (0.3618)	−1	64	10	128	491.4842	1.54
19	500	25	1	0.1	0	52	5	108	493.2051 (0.5004)	−1	64	8	128	501.1132	1.60
20	500	25	10	1	0	52	5	108	583.8744 (0.5442)	−1	64	4	128	594.2374	1.77
21	500	30	0.01	0.001	−1	57	4	118	528.2474 (0.5531)	−1	64	12	128	529.5645	0.25
22	500	30	0.1	0.01	−1	57	4	118	529.3426 (0.5981)	−1	64	10	128	530.5467	0.23
23	500	30	1	0.1	−1	59	4	122	538.3565 (0.6096)	−1	64	7	128	540.1331	0.33
24	500	30	10	1	−1	57	4	118	629.4388 (0.6529)	−1	64	4	128	633.2999	0.61
25	5000	25	0.01	0.001	34	167	84	181	1534.5556 (1.6409)	38	128	107	256	1579.9145	2.96
26	5000	25	0.1	0.01	34	167	84	181	1538.7238 (2.3392)	38	128	102	256	1583.3790	2.90
27	5000	25	1	0.1	34	167	84	348	1570.1617 (2.9671)	38	128	95	256	1617.4042	3.01
28	5000	25	10	1	34	168	84	350	1889.2586 (2.4272)	38	128	85	256	1949.4063	3.18
29	5000	30	0.01	0.001	33	182	83	196	1675.3595 (2.7624)	26	256	106	256	1760.6544	5.09
30	5000	30	0.1	0.01	33	184	83	198	1678.5555 (2.4732)	26	256	100	256	1764.0248	5.09
31	5000	30	1	0.1	33	184	83	382	1713.3783 (3.1637)	26	256	93	256	1797.1060	4.89
32	5000	30	10	1	33	182	83	378	2032.9613 (3.2877)	26	256	84	256	2120.0713	4.28

partition can be found by solving a deterministic version of the problem.)

We propose the following policy for the one-warehouse multiretailer system. Let N be the number of retailers. The policy requires $N + 1$ pairs of control parameters: (r_i, Q_i), $i = 0, 1, \ldots, N$, where index 0 represents the warehouse and index $i \geq 1$ represents retailer i. Orders at the warehouse are triggered by the inventory position of the *subsystem* consisting of the warehouse and the retailers in R^+, i.e., whenever the subsystem's inventory position drops to r_0, the warehouse places an order. To determine the size of the order, let x_i be the inventory position of retailer i at the time of order placement. Then the order size is

$$Q_0 + \sum_{\substack{i \in R^- \\ x_i \leq r_i}} (r_i + Q_i - x_i).$$

Therefore, Q_0 is the *base quantity*. The additional quantity is the total amount required to restore the inventory positions of those retailers in R^- whose inventory positions are at or below their reorder points r_i to their order-up-to levels $r_i + Q_i$. Note that different retailers in R^- may contribute to warehouse orders with different frequencies. Let λ_i be the average demand rate at retailer i. For retailer $i \in R^-$, if Q_i/λ_i is chosen to be a multiple m of $Q_0/\sum_{j\in R^+} \lambda_j$ then retailer i will on average contribute once every m warehouse orders. Consequently, different warehouse orders may have different sizes, as in the deterministic version of the system.

When an order arrives at the warehouse, shipments are sent to the retailers. The exact sizes of the shipments depend on the inventory positions of the retailers. In the case of a serial system with zero order leadtime (and only one retailer), the retailer's inventory position is exactly r_0 when an order arrives at the warehouse, and the shipment size can be *predetermined*. However, when there are multiple retailers, the retailers' inventory positions are random variables. We nevertheless determine the sizes of the shipments at the arrival epoch of a warehouse order, so as to restore retailer *i*'s (in R^+ or R^-) inventory position to its order-up-to level $r_i + Q_i$ if retailer *i*'s inventory position is at or below r_i. If the warehouse's on-hand inventory is insufficient to achieve this goal, then some rationing is required. A plausible rationing scheme can be obtained by solving myopic allocation problems such as those formulated in Eppen and Schrage (1981) and Federgruen and Zipkin (1984a, b). (Those retailers whose inventory positions are above their reorder points at the arrival epoch of a warehouse order will not receive any shipment immediately.)

In between order arrivals at the warehouse, shipments are made to the retailers in R^+ in accordance with retailer-specific (r, Q) policies. That is, a shipment of size Q_i is sent to retailer i, $i \in R^+$ each time the retailer's inventory position drops to r_i, with the clear modification that if the warehouse's on-hand inventory is less than Q_i, only the partial quantity is shipped. (As with the serial system, we should have $\sum_{i\in R^+} (r_i + Q_i) \le r_0 + Q_0$. The retailers in R^- will not get any shipment between orders.)

The above class of policies is new. Compared with the (R, Q) policies suggested in the literature for the distribution network (see, e.g., Deuermeyer and Schwarz 1981), it has three salient features: (i) a systematic design to differentiate the initial shipment to a retailer within a warehouse order cycle from subsequent shipments to the same retailer within the same cycle, if any; (ii) nonstationary order quantities at the warehouse; and (iii) the dynamic allocation of warehouse orders at time of arrival. The second and third features are especially important when the retailer characteristics are significantly different. Indeed, the classical (R, Q) policies for this system were originally designed for systems with identical retailers. We shall leave the empirical test of this new class of policies to a future study.[1]

[1] The author would like to thank Awi Federgruen and the anonymous reviewers for their helpful comments and suggestions that have significantly improved the exposition of this paper. Partial support was provided by the Faculty Research Fund, Graduate School of Business, Columbia University.

References

Atkins, D., S. De. 1992. 94%-effective lot-sizing for a two-stage serial production/inventory system with stochastic demand. Working Paper, Faculty of Commerce, University of British Columbia, Vancouver, B.C., Canada.

——, D. Sun. 1995. 98%-effective lot sizing for series inventory systems with backlogging. *Oper. Res.* **43** 335–345.

Axsater, S. 1990. Simple solution procedures for a class of two-echelon inventory problems. *Oper. Res.* **38** 64–69.

——. 1993a. Exact and approximate evaluation of batch ordering policies for two-level inventory systems. *Oper. Res.* **41** 777–785.

——. 1993b. Simple evaluation of echelon stock (*R*, *Q*) policies for two-level inventory systems. Working Paper, Lulea University of Technology, Sweden.

——. 1993c. Continuous review policies for multi-level inventory systems with stochastic demand. S. Graves, A. Rinnooy Kan and P. Zipkin, eds. *Handbook in Operations Research and Management Science, Vol. 4, Logistics of Production and Inventory*, North Holland.

Chen, F. 1998. Stationary policies in multi-echelon inventory systems with deterministic demand and backlogging. *Oper. Res.* **46** (S3) S26–S34.

——, Y.-S. Zheng. 1994a. Evaluating echelon stock (*R*, *nQ*) policies in serial production/inventory systems with stochastic demand. *Management Sci.* **40** 1262–1275.

——, ——. 1994b. Lower bounds for multi-echelon stochastic inventory systems. *Management Sci.* **40** 1426–1443.

——, ——. 1997. One-warehouse multi-retailer systems with centralized stock information. *Oper. Res.* **45** 275–287.

——, ——. 1998. Near-optimal echelon-stock (R, Q) policies in multi-stage serial systems. *Oper. Res.* **46** (4) 592–602.

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

——, ——. 1962. Approximate solutions to a simple multi-echelon inventory problem. K. J. Arrow, S. Karlin and H. Scarf, eds. *Studies in Applied Probability and Management Science*. Stanford University Press, Stanford, CA. 88–110.

De Bodt, M., S. Graves. 1985. Continuous review policies for a multi-echelon inventory problem with stochastic demand. *Management Sci.* **31** 1286–1295.

Deuermeyer, B., L. Schwarz. 1981. A model for the analysis of system service level in warehouse/retailer distribution systems: The identical retailer case. L. Schwarz, ed. *Studies in the Management Sciences*: *The Multi-Level Production/Inventory Control Systems*, Vol. 16. North-Holland, Amsterdam. 163–193.

Eppen, G., L. Schrage. 1981. Centralized ordering policies in a multiwarehouse system with leadtimes and random demand. L. Schwarz, ed. *Multi-Level Production/Inventory Control Systems*: *Theory and Practice*. North Holland, 51–69.

Federgruen, A. 1993. Centralized planning models for multi-echelon inventory systems under uncertainty. S. Graves, A. Rinnooy Kan and P. Zipkin, eds. *Handbook in Operations Research and Management Science, Vol. 4, Logistics of Production and Inventory*. North-Holland, Amsterdam.

——, M. Queyranne, Y.-S. Zheng. 1992. Simple power of two policies are close to optimal in general class of production/distribution networks with general joint setup costs. *Math. Oper. Res.* **17** 951–963.

——, Z. Schechner. 1983. Cost formulas for continuous review inventory models with fixed delivery lags. *Oper. Res.* **31** 957–965.

——, Y.-S. Zheng. 1992. An efficient algorithm for computing an optimal (r, Q) policy in continuous review stochastic inventory systems. *Oper. Res.* **40** 808–813.

——, P. Zipkin. 1984a. Approximation of dynamic, multi-location production and inventory problems. *Management Sci.* **30** 69–84.

——, P. Zipkin. 1984b. Allocation policies and cost approximation for multi-location inventory systems. *Naval Res. Logist. Quart.* **31** 97–131.

Graves, S. 1985. A multi-echelon inventory model for a reparable item with one-for-one replenishment. *Management Sci.* **31** 1247–1256.

Hadley, G., T. Whitin. 1963. *Analysis of Inventory Systems*. Prentice-Hall, Englewood Cliffs, NJ.

Jackson, P. 1988. Stock allocation in a two-echelon distribution system or 'What to do until your ship comes in'. *Management Sci.* **34** 880–895.

Lee, H., K. Moinzadeh. 1987a. Two-parameter approximations for multi-echelon reparable inventory models with batch ordering policy. *IIE Trans.* **19** 140–149.

——, ——. 1987b. Operating characteristics of a two-echelon inventory system for reparable and consumable items under batch ordering and shipment policy. *Naval Res. Logist. Quart.* **34** 365–380.

Maxwell, W., J. Muckstadt. 1985. Establishing consistent and realistic reorder intervals in production-distribution systems. *Oper. Res.* **33** 1316–1341.

Mitchell, J. 1987. 98%-effective lot-sizing for one-warehouse multi-retailer inventory systems with backlogging. *Oper. Res.* **35** 399–404.

Moinzadeh, K., H. Lee. 1986. Batch size and stocking levels in multi-echelon reparable systems. *Management Sci.* **32** 1567–1581.

Muckstadt, J., R. Roundy. 1993. Analysis of multistage production systems. S. Graves, A. Rinnooy Kan and P. Zipkin, eds. *Handbook in Operations Research and Management Science, Vol. 4, Logistics of Production and Inventory*. North Holland, Amsterdam.

Roundy, R. 1985. 98%-effective integer-ratio lot-sizing for one-warehouse multi-retailer systems. *Management Sci.* **31** 1416–1430.

——. 1986. 98%-effective lot-sizing rule for multi-product, multi-facility production-inventory systems. *Math. Oper. Res.* **11** 699–727.

Schwarz, L. 1973. A simple continuous review deterministic one-warehouse multi-retailer problem. *Management Sci.* **19** 555–566.

Sherbrooke, C. 1968. METRIC: A multi-echelon technique for recoverable item control. *Oper. Res.* **16** 122–141.

Svoronos, A., P. Zipkin. 1988. Estimating the performance of multi-level inventory systems. *Oper. Res.* **36** 57–72.

——1991. Evaluation of one-for-one replenishment policies for multi-echelon inventory systems. *Management Sci.* **37** 68–83.

Zhang, H. 1996. A note on the convexity of service-level measures of the (r, Q) system. Working Paper, Hong Kong University of Science and Technology, Hong Kong.

Zheng, Y.-S. 1992. On properties of stochastic inventory systems. *Management Sci.* **38** 87–103.

Zipkin, P. 1986. Inventory service-level measures: Convexity and approximation. *Management Sci.* **32** 975–981.

Accepted by Awi Federgruen; received May 28, 1996. *This paper has been with the author* 15 *days for* 1 *revision*.