

Lower Bounds for Multi-echelon Stochastic Inventory Systems

Fangruo Chen • Yu-Sheng Zheng

Graduate School of Business, Columbia University, New York, New York 10027

Operations and Information Management Department, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104

We establish lower bounds on the minimum costs of managing certain production-distribution networks with setup costs at all stages and stochastic demands. These networks include serial, assembly, and one-warehouse multi-retailer systems. We obtain the bounds through novel cost-allocation schemes. We evaluate the bounds' performance for one-warehouse multi-retailer systems by comparing them with simple, heuristic policies. The bounds are quite tight for systems with a small number of retailers. We also present simplified proofs of known optimality results for serial and assembly systems.

(Stationary Analysis; Cost Allocation; Economies of Scale; Multi-stage Production Systems; Distribution Networks; Replenishment Strategies)

1. Introduction

This paper establishes lower bounds on the minimum costs of managing certain production-distribution networks with setup costs at all stages and stochastic demands. These networks include serial, assembly, and one-warehouse multi-retailer systems. This study was motivated by our effort to identify cost-effective control policies for these systems. It is well known that optimal control of these systems is difficult due to the presence of setup costs, especially at downstream stages (Clark and Scarf 1962). As a result, heuristic control policies become attractive. Tight lower bounds on the minimum costs are necessary to evaluate the effectiveness of heuristic policies.

We obtain lower bounds through novel cost-allocation schemes. Cost allocation has been successfully applied to obtain lower bounds for several related multi-stage systems by Roundy (1985, 1986), Atkins and Iyoun (1987) and Zheng (1987), among others. (See Atkins 1990 for a review on cost allocation as well as other lower-bounding methodologies.) For more general systems, new lower bounds can be created by combining cost allocation with "physical decomposition." In particular, imagine that the product consists of a number

of fictitious components. Each component is supplied/produced through a subsystem—a part of the original system—and is allocated part of the costs. By assuming that the components can be replenished and sold separately (i.e., physical decomposition), the original system decomposes to a number of independent systems, one for each component. The sum of the minimum costs of these independent systems is a lower bound on the minimum cost of the original system.

The virtue of the cost-allocation, physical-decomposition framework is that it provides freedom in creating fictitious components and in allocating costs. By allocating one-period expected holding and backorder cost *functions* (hereafter referred to as loss functions), we re-derive and generalize the so-called "induced-penalty bounds" for serial systems (Clark and Scarf 1962 and Atkins and De 1992) and for one-warehouse multi-retailer systems (Rosling 1977). On the other hand, by allocating cost *rates*, we generate a new class of bounds, called "parameter-allocation bounds." These bounds are derived for both periodic-review and continuous-review, stationary systems operating over an infinite horizon.

Our lower-bounding framework depends upon sev-

eral key results in the multi-echelon inventory literature. These include the optimal policies for the periodic-review, serial system of Clark and Scarf (1960) and Federgruen and Zipkin (1984b) and the periodic-review, assembly system of Rosling (1989), and the lower bound for the one-warehouse multi-retailer system of Federgruen and Zipkin (1984a, b, c). In this paper, we provide a simpler proof for the above optimality results (for serial and assembly systems). We first establish (induced-penalty) lower bounds on the minimum costs of these systems, and then show that these lower bounds can actually be achieved by feasible policies. Therefore the feasible policies are optimal and the lower bounds are the minimum costs. This method also leads to parallel optimality results for the continuous-review counterparts of the above systems.

We have conducted a numerical test of the bounds' performance by comparing them with simple, heuristic policies in one-warehouse multi-retailer systems. The results indicate that the bounds are tight for systems with a small number of retailers.

The paper is organized as follows. Sections 2–4 focus on periodic-review systems. In §2, we define the systems to be considered and review related known results. In §3, we present the new optimality proofs for serial and assembly systems. In §4, we establish lower bounds for serial, assembly, and one-warehouse multi-retailer systems with setup costs at all stages. Section 5 describes briefly how the optimality results and lower bounds can be extended to continuous-review systems. In §6, we report the numerical study. Section 7 contains concluding remarks.

2. The Models

In this section, we describe the models to be studied, review relevant known results, and define general notation. Consider serial, assembly, and one-warehouse multi-retailer systems with stochastic demand. (Since one-warehouse multi-retailer systems are the only distribution systems considered in this paper, they will be referred to as distribution systems in the sequel.) Here we concentrate on periodic-review systems. The following assumptions are common to all the systems:

- Each system produces or distributes a single final product through multiple stages. A stage can be a distinctive location in distribution systems or a buffer for

a particular intermediate item in production systems. The stage(s) that receives stock from an outside supplier is called the "highest" stage. The stage(s) where customer demand arises is called the "lowest" stage.

- The transportation (or production) leadtime at stage i is a constant of L_i periods, i.e., a batch released to stage i at the beginning of period t will be received at the beginning of period $t + L_i$.

- Customer demands at different periods are independent and identically distributed. F denotes the cumulative distribution function (cdf) of one-period demand. If customer demand arises at more than one stage, F denotes the joint cdf. Demand is discrete.

- Replenishment decisions are centralized and based on system-wide inventory information, which is available and free. The objective is to minimize long-run average system-wide costs.

- A fixed setup cost K_i is incurred for each batch released to stage i .

- The echelon holding cost at stage i is h_i per unit · period. It is more expensive to hold inventory at a lower stage than at an upper stage, i.e., $h_i \geq 0$.

- Unsatisfied customer demand is fully backlogged at the lowest stage with penalty cost p per unit · period. For systems where demand occurs at several stages, the penalty cost is p_i at stage i .

Serial systems. The system consists of N stages arranged in series with stage 1 receiving stock from stage 2, 2 from 3, etc., and stage N from an outside supplier with infinite stock. Customer demand occurs only at stage 1. Denote the system by "Series(N).". When only stage N has a setup cost K , i.e., $K_N = K$ and $K_i = 0$ for $1 \leq i \leq N - 1$, denote the system by "Series(N, K).".

Assembly systems. The system produces a final product through a tree structure, denoted by \mathcal{G} , with N stages. More specifically, each stage has exactly one successor stage, except for the lowest stage (the root of the tree) where customer demand arises. The stages without any predecessor stage (the leaves of the tree) are replenished by outside suppliers with infinite stock. Without loss of generality, we assume that exactly one unit of each intermediate item is used to produce its immediate successor item. We will refer to the above assembly system as "Assembly(\mathcal{G}).". Let $S(i)$ be the set consisting of stage i and all its successor stages. Define M_i as the total leadtime at stage i , i.e., $M_i = \sum_{j \in S(i)} L_j$.

The stages are numbered such that M_i is nondecreasing in i . (Thus stage 1 is the lowest stage.) When only stage N , the stage with the longest total leadtime, has a setup cost K , the above system will be referred to as "Assembly(\mathcal{G} , K)."

Distribution systems. The system has one warehouse and N retailers. The warehouse (stage 0) receives stock from an outside supplier with infinite stock and replenishes the retailers (stages 1, . . . , N). Customer demands occur only at the retailers. The above system will be denoted by "Distribution(N)," or by "Distribution(N , K)" when only the warehouse has a setup cost K .

We proceed to briefly review the existing results on Series(N , K), Assembly(\mathcal{G} , K), and Distribution(N , K), which are the building blocks of our lower bounds.

Note that Series(1, K) or Series(1) is the extensively studied single-location system. For this system, it is well known that (s, S) policies are optimal. (See Scarf 1960 and Iglehardt 1963 for the original proof, and Zheng 1991 for a simple proof.) Under an (s, S) policy, an order is placed to increase the inventory position (= inventory on hand + outstanding orders - backorders) to S whenever it drops to or below s . The minimum cost of the system is therefore the cost associated with the optimal (s, S) policy. For a simple and efficient algorithm to compute optimal (s, S) policies, see Zheng and Federgruen (1991).

In their seminal paper, Clark and Scarf (1960) characterized the optimal policies for the finite-horizon version of Series(N , K). Their result was extended by Federgruen and Zipkin (1984b) to infinite-horizon systems. An optimal policy for Series(N , K) uses an echelon stock (s, S) policy at stage N and an echelon stock order-up-to policy at each downstream stage. We denote the minimum cost of Series(N , K) by a function of the system parameters, $C_N^s(K, \mathbf{h}, p, \mathbf{L}, F)$, where $\mathbf{h} = \{h_i\}_1^N$ and $\mathbf{L} = \{L_i\}_1^N$. A new derivation of the optimal policy and its minimum cost will be provided in §3.

For Assembly(\mathcal{G} , 0), Rosling (1989) characterized the optimal policies by showing that the system is equivalent to a serial system by introducing the notion of total leadtimes. Section 3 provides a simple proof that this equivalence holds for Assembly(\mathcal{G} , K). Denote the minimum cost of Assembly(\mathcal{G} , K) by $C_N^a(K, \mathbf{h}, p, \mathbf{L}, F)$, where $\mathbf{h} = \{h_i\}_1^N$ and $\mathbf{L} = \{L_i\}_1^N$.

Optimal policies for Distribution(N , K) are unknown. The difficulty is due to possible stock imbalance among different retailers (Eppen and Schrage 1981, Zipkin 1984). However, Federgruen and Zipkin (1984a, b, c) provided a lower bound on the minimum cost of the system by allowing a *free* inventory (position) rebalance among the retailers. Under such a relaxation, the original system reduces to a single-location system whose minimum cost can be easily computed. This minimum cost is a lower bound on the minimum cost of the original system. Denote this lower bound by $C_N^d(K, \mathbf{h}, \mathbf{p}, \mathbf{L}, F)$, where $\mathbf{h} = \{h_i\}_0^N$, $\mathbf{p} = \{p_i\}_1^N$ and $\mathbf{L} = \{L_i\}_0^N$. (Notice that \mathbf{h} and \mathbf{L} are being used to represent different vectors in different contexts.)

We proceed to define key state variables. *Echelon inventory level* at stage i is the inventory on hand at stage i plus inventories at or in transit to all its successor stages minus total customer backorder at stage i and its successor stages. *Echelon inventory position* at stage i is the sum of echelon inventory level at stage i and the inventories *in transit* to stage i . Assume that the activities in a period happen in the following sequence: (a) at the beginning of the period, replenishment batches are released and received; (b) during the period, customer demand occurs; and (c) at the end of the period, holding and backorder costs are assessed based on the ending inventory and backorder levels. Fix a period t . Let $IL_i^-(t)$ and $IP_i(t)$ be the (beginning) echelon inventory level and the echelon inventory position at stage i after replenishment batches are released and received but before demand occurs. At the end of period t (after demand), let

$I_i(t)$ = echelon inventory at stage i
 = inventory on hand at stage i plus inventories at or in transit to all its successor stages

$B_i(t)$ = customer backorder level at stage i , defined only if customer demand arises at stage i ; use $B(t)$ if customer demand arises only at one stage

$IL_i(t)$ = (ending) echelon inventory level at stage i

For $t_1 \leq t_2$, we use $[t_1, t_2]$ to denote the time interval of periods t_1, \dots, t_2 . Similarly, we use $[t_1, t_2)$ to denote the time interval of periods $t_1, \dots, t_2 - 1$. For Series(N) and Assembly(\mathcal{G}), let $D[t_1, t_2]$ and $D[t_1, t_2)$ be the total customer demand in $[t_1, t_2]$ and $[t_1, t_2)$ respectively for any $t_1 \leq t_2$ with $D[t, t) \equiv 0$. For Distribution(N), let

$D_i[t_1, t_2]$ and $D_i[t_1, t_2]$ be the total customer demand at retailer i , $1 \leq i \leq N$, in $[t_1, t_2]$ and $[t_1, t_2]$ respectively for any $t_1 \leq t_2$ with $D_i[t, t] \equiv 0$. Let $D_0[t_1, t_2] = \sum_{i=1}^N D_i[t_1, t_2]$ and $D_0[t_1, t_2] = \sum_{i=1}^N D_i[t_1, t_2]$.

3. Optimality Proofs

In this section, we rederive the optimality results for Series(N, K) and Assembly(\mathcal{G}, K) by first establishing lower bounds on the minimum costs of the systems and then showing that the bounds can be reached by feasible policies.

3.1. Series(N, K)

Consider any feasible policy for Series(N, K). The total holding and backorder cost incurred at period t is

$$\sum_{i=1}^N h_i I_i(t) + pB(t).$$

Since $IL_i(t) = I_i(t) - B(t)$, the above expression can be rewritten as

$$\sum_{i=1}^N h_i IL_i(t) + (p + H_1)B(t)$$

where $H_1 = \sum_{i=1}^N h_i$, the installation holding cost at stage 1. For convenience, we charge the following cost to period $t - M_N$

$$\sum_{i=1}^N h_i IL_i(t - M_{i-1}) + (p + H_1)B(t) \quad (1)$$

where $M_0 = 0$ and $M_i = \sum_{j=1}^i L_j$ for $i = 1, \dots, N$. Notice that M_i is the total leadtime at stage i . This cost accounting scheme only shifts costs across time periods. Therefore, it does not affect the long-run average holding and backorder costs. The accounting scheme has the following intuitive interpretation: Since $IL_i(t - M_{i-1}) = IP_i(t - M_i) - D[t - M_i, t - M_{i-1}]$, we see that $IL_i(t - M_{i-1})$ is statistically determined by $IP_i(t - M_i)$. Moreover, by definition, $IP_i(t - M_i)$ is constrained by the beginning echelon inventory level at stage $i + 1$, i.e. $IP_i(t - M_i) \leq IL_{i+1}^-(t - M_i)$, with the difference being the on-hand inventory at stage $i + 1$. But $IL_{i+1}^-(t - M_i)$ is, in turn, statistically determined by $IP_{i+1}(t - M_{i+1})$: $IL_{i+1}^-(t - M_i) = IP_{i+1}(t - M_{i+1}) - D[t - M_{i+1}, t - M_i]$. A simple induction shows that $IP_N(t - M_N)$ determines, directly or indirectly, $IL_i(t - M_{i-1})$

for $i = 1, \dots, N$. Therefore it is meaningful to charge the holding and backorder cost in (1) to period $t - M_N$ when the decision on $IP_N(t - M_N)$ is made. For simplicity, write IP_i for $IP_i(t - M_i)$, IL_i for $IL_i(t - M_{i-1})$, IL_i^- for $IL_i^-(t - M_{i-1})$, and B for $B(t)$. Consequently, the total holding and backorder cost charged to period $t - M_N$ is

$$\sum_{i=1}^N h_i IL_i + (p + H_1)B. \quad (2)$$

Next, we identify a lower bound on the expected value of (2). Define

$$G_1(y) = E[h_1(y - D[t - M_1, t]) + (p + H_1)(y - D[t - M_1, t])^-].$$

Notice that $G_1(\cdot)$ is convex. Let Y_1 be a minimum point of $G_1(\cdot)$ and $C_1 = G_1(Y_1)$. Define

$$G_1^1(y) = \begin{cases} C_1 & \text{if } y \leq Y_1 \\ G_1(y) & \text{otherwise} \end{cases}$$

$$G_1^2(y) = G_1(y) - G_1^1(y)$$

and

$$G_2(y) = E[h_2(y - D[t - M_2, t - M_1]) + G_1^2(y - D[t - M_2, t - M_1])].$$

Note that both $G_1^2(\cdot)$ and $G_2(\cdot)$ are convex. The above definitions can be extended inductively. In particular, suppose that $G_i(\cdot)$, $2 \leq i < N$, are defined and convex. Let Y_i be the minimum point of $G_i(\cdot)$ and $C_i = G_i(Y_i)$. Define

$$G_i^i(y) = \begin{cases} C_i & \text{if } y \leq Y_i \\ G_i(y) & \text{otherwise} \end{cases}$$

$$G_i^{i+1}(y) = G_i(y) - G_i^i(y)$$

and

$$G_{i+1}(y) = E[h_{i+1}(y - D[t - M_{i+1}, t - M_i]) + G_i^{i+1}(y - D[t - M_{i+1}, t - M_i])].$$

By induction, it is easy to see that the $G_i(\cdot)$'s are all convex. (The above functions were introduced by Rosling 1989 under a slightly different form.) Now consider a single-stage system with setup cost K , loss function

$G_N(\cdot)$, and the demand process of the original system. We know that (s, S) policies are optimal. Let (s_N, S_N) be the optimal (s, S) policy and C_N the minimum cost of the single-stage system. (The above definitions will also be used in discussing assembly systems.)

LEMMA 1. For Series (N, K) , we have $E[\sum_{i=1}^N h_i IL_i + (p + H_1)B] \geq E[\sum_{i=1}^{N-1} G_i^i(IP_i) + G_N(IP_N)]$.

PROOF. Since $IL_1 = IP_1 - D[t - M_1, t]$ and $B = (IL_1)^-$, from the definition of $G_1(\cdot)$ we have

$$E[h_1 IL_1 + (p + H_1)B] = EG_1(IP_1).$$

Since $IP_1 \leq IL_2^-$ and $G_1^2(\cdot)$ is nonincreasing, we have $G_1(IP_1) = G_1^1(IP_1) + G_1^2(IP_1) \geq G_1^1(IP_1) + G_1^2(IL_2^-)$. Therefore

$$E\left[\sum_{i=1}^2 h_i IL_i + (p + H_1)B\right] \geq E[G_1^1(IP_1) + h_2 IL_2 + G_1^2(IL_2^-)].$$

Notice that $G_1^2(IL_2^-)$ corresponds to the induced-penalty cost in Clark and Scarf (1960). It represents the expected holding and backorder cost increment at stage 1 when stage 2's on-hand stock is insufficient to bring stage 1's inventory position up to Y_1 , i.e. $IL_2^- < Y_1$. Charging this induced-penalty cost to stage 2, we have

$$\begin{aligned} E[h_2 IL_2 + G_1^2(IL_2^-)] &= E[h_2(IP_2 - D[t - M_2, t - M_1]) \\ &\quad + G_1^2(IP_2 - D[t - M_2, t - M_1])] \\ &= EG_2(IP_2). \end{aligned}$$

Therefore,

$$E\left[\sum_{i=1}^2 h_i IL_i + (p + H_1)B\right] \geq E[G_1^1(IP_1) + G_2(IP_2)].$$

The lemma is now true for $N = 2$. A simple induction verifies the lemma for any value of N . \square

LEMMA 2. $\sum_{i=1}^N C_i$ is a lower bound on the minimum long-run average cost of Series (N, K) .

PROOF. Recall that C_i is the minimum value of $G_i(\cdot)$ for $i < N$. Thus, it follows from Lemma 1 that

$$E\left[\sum_{i=1}^N h_i IL_i + (p + H_1)B\right] \geq \sum_{i=1}^{N-1} C_i + EG_N(IP_N).$$

In other words, given $IP_N = y$, the expected systemwide holding and backorder costs charged to period $t - M_N$ under any policy are bounded below by $\sum_{i=1}^{N-1} C_i + G_N(y)$. By substituting the latter for the former, the original system reduces to a single-stage system with setup cost K and loss function $\sum_{i=1}^{N-1} C_i + G_N(y)$. Since $\sum_{i=1}^{N-1} C_i$ is constant, the optimal policy of this single-stage system is the (s_N, S_N) policy with minimum cost $\sum_{i=1}^N C_i$. Clearly, this minimum cost is a lower bound on the minimum cost of Series (N, K) . \square

We next show that the above lower bound can be achieved by a feasible policy. Consider the following feasible policy for Series (N, K) . Stage N uses the echelon stock (s_N, S_N) policy: if the echelon inventory position at stage N falls to or below s_N then order up to S_N ; and stage $i, i = 1, \dots, N - 1$, uses an echelon stock order-up-to Y_i policy: if the echelon inventory position at stage i is below Y_i then order up to Y_i . This is indeed the optimal policy characterized by Clark and Scarf (1960) and Federgruen and Zipkin (1984b). To prove its optimality, it suffices to show that it actually achieves the lower bound $\sum_{i=1}^N C_i$.

THEOREM 1. For Series (N, K) , the optimal policy uses the echelon stock (s_N, S_N) policy at stage N and the echelon stock order-up-to Y_i policy at stage $i, i = 1, \dots, N - 1$, with minimum cost $C_N^s(K, \mathbf{h}, p, \mathbf{L}, F) = \sum_{i=1}^N C_i$.

PROOF. Suppose that the above policy is in place for Series (N, K) . All we need to show is that, given $IP_N = y$, the expected system-wide holding and backorder costs charged to period $t - M_N$ are exactly equal to $\sum_{i=1}^{N-1} C_i + G_N(y)$.

Take any $i < N$. Notice that $IP_i \leq IL_{i+1}^-$. Since stage i orders up to Y_i , we have $IP_i = \min\{Y_i, IL_{i+1}^-\}$. Thus $G_i^i(IP_i) = C_i$. Therefore it only remains to show that the inequality in Lemma 1 is an equality. To see this, simply notice that $G_i^{i+1}(IP_i) = G_i^{i+1}(IL_{i+1}^-)$ since if $IP_i = Y_i$ then $IL_{i+1}^- \geq Y_i$ and thus $G_i^{i+1}(IP_i) = G_i^{i+1}(IL_{i+1}^-) = 0$. \square

We proceed to interpret the above lower bound and optimality proof under a cost-allocation framework. For ease of presentation, we only consider Series $(2, K)$. Imagine that the end product (at stage 1) consists of two components, 1 and 2. The subsystem that supplies component 2 is exactly the same as the original system. But component 1 is supplied through a truncated sys-

tem: it is directly sourced from an outside vendor with ample supply. The delivery leadtime from the vendor to stage 1 is L_1 , the same as the transportation time from stage 2 to stage 1. In the original system, replenishment is "coordinated" between the two components: each unit of component 2 leaving stage 2 is *matched* by a unit of component 1 leaving the outside vendor. After arriving at stage 1, these components are sold in packages. Since component 2 is the only item at stage 2, it bears the setup cost K at stage 2 and the stage-2 echelon holding cost with rate h_2 . The loss function at stage 1, $G_1(\cdot)$, is shared by the two components: $G_1^1(\cdot)$ by component 1 and the remaining, $G_1^2(\cdot)$, by component 2. (The systems of the components are depicted in Figure 1.) Notice that the above cost allocation fully recovers the total costs of the original system.

Now relax the above coordination constraint by assuming that the components can be replenished and sold independently of each other. Under this relaxation, the original system decomposes to two independent systems, one for each component. The sum of the minimum costs of the components is a lower bound on the minimum cost of the original system.

The system of component 1 is a single-stage system with loss function $G_1^1(\cdot)$ and no setup cost. Thus its minimum cost is $C_1 = G_1^1(Y_1)$, which can be reached by *any* policy that keeps the inventory position of component 1, IP_1 , at or below Y_1 . Now consider the system of component 2. Notice that the loss function at stage

1, $G_1^2(IP_1)$, is a nonincreasing function of IP_1 . Since no setup cost is incurred for shipping stock from stage 2 to stage 1, there is no incentive to hold inventory at stage 2. Therefore an optimal policy would keep $IP_1 = IL_2^-$ and the system collapses to a single-stage system. Since $IL_2 = IP_2 - D[t - M_2, t - M_1]$ and $IL_2^- = IP_2 - D[t - M_2, t - M_1)$, the total holding and back-order cost of component 2 can be expressed as a loss function of its echelon inventory position at stage 2, IP_2 :

$$\begin{aligned} E[h_2 IL_2 + G_1^2(IP_1)] &= E[h_2 IL_2 + G_1^2(IL_2^-)] \\ &= E[h_2(IP_2 - D[t - M_2, t - M_1]) \\ &\quad + G_1^2(IP_2 - D[t - M_2, t - M_1))] \\ &= EG_2(IP_2). \end{aligned}$$


Therefore the resulting single-stage system has loss function $G_2(\cdot)$ and setup cost K . Since $G_2(\cdot)$ is convex, the (s_2, S_2) policy is optimal with minimum cost C_2 . Thus, $C_1 + C_2$ is a (induced-penalty) lower bound on the minimum cost of Series(2, K).

To show Theorem 1, recall that the only relaxation used to generate the lower bound is that the components can be replenished and sold independently of each other. Therefore, it suffices to show that there exist optimal policies for the component systems such that the components arriving at stage 1 are exactly matched. First consider component 2. The optimal policy for component 2 is to use the echelon stock (s_2, S_2) policy at stage 2 and ship all to stage 1 (i.e., ship up to IL_2^-). However, since $G_1^2(y)$ is flat for $y \geq Y_1$, the policy remains optimal if stage 2 ships up to $\min\{Y_1, IL_2^-\}$. Now consider component 1. Since $G_1^1(y)$ is flat for $y \leq Y_1$, any policy that does not raise the inventory position of component 1 to above Y_1 is optimal. One such policy is to ship up to $\min\{Y_1, IL_2^-\}$. Observe that under the above optimal policies in the component systems, the components arriving at stage 1 are exactly matched. This completes the proof for the theorem. The above cost-allocation approach will be used in §4 to derive a similar bound for Series(N).

3.2. Assembly (\mathcal{G}, K)

Suppose that a batch is released to stage i at period t . Conventionally, this batch is not counted as part of the echelon inventory of stage i until its arrival at period t

Figure 1 Decomposition of Series (2, K)



+ L_i . For convenience, we introduce an accounting scheme under which the above batch becomes a part of the echelon inventory of stage i at period $t + l_i$, where $l_i = M_i - M_{i-1}$. (For readers familiar with Rosling 1989, we reverse his definitions of L_i and l_i .) Let $s(i)$ be the immediate successor of stage i with $s(1) = 0$. Since $L_i = M_i - M_{s(i)}$ and $M_{i-1} \geq M_{s(i)}$, we have $l_i \leq L_i$. Therefore, under the new accounting scheme, the batch released to stage i at period t incurs echelon holding costs of stage i while it is still in assembly at period $t + l_i$. Consequently, the new accounting scheme overcharges the echelon holding costs. But observe that, under any policy, the difference in long-run average system-wide holding costs between the new and the original accounting schemes is

$$\sum_{i=1}^N h_i \mu (L_i - l_i)$$

where μ is the expected one-period demand. Since the difference is a constant, an optimal policy under the original accounting scheme is still optimal under the new one, and vice versa. Unless otherwise mentioned, the new accounting scheme will be used for the rest of this section.

Write IP_i for $IP_i(t - M_i)$, IL_i for $IL_i(t - M_{i-1})$, IL_i^- for $IL_i^-(t - M_{i-1})$, and B for $B(t)$. Notice that

$$IL_i = IP_i - D[t - M_i, t - M_{i-1}] \quad \text{and} \quad B = (IL_1)^-$$

Define

$$IP_{ij} = IP_i - D[t - M_i, t - M_j]$$

for $i \geq j$. Clearly $IP_{ii} = IP_i$. Note that IP_{ij} is the echelon inventory position of stage i at period $t - M_j$ after excluding orders released after period $t - M_i$. In particular, $IP_{i,i-1} = IL_i^-$, the beginning echelon inventory level at stage i ; while $IP_{i,s(i)}$ is the beginning echelon inventory level at stage i under the original accounting scheme. Write IP^i for $\min_{n \geq i} \{IP_{ni}\}$. Thus $IP^N = IP_N$.

LEMMA 3.

- (i) $IP^1 = IP_1$
- (ii) $IP^i \leq IP_i, i = 2, \dots, N - 1, IP^N = IP_N$
- (iii) $IP^i \leq IP^{i+1} - D[t - M_{i+1}, t - M_i], i = 1, \dots, N - 1$.

PROOF. (i) Take any stage $i \geq 2$. Since stage $s(i)$ is the immediate successor of stage i , the echelon inven-

tory position of stage $s(i)$ is constrained by the beginning echelon inventory level at stage i under the original accounting scheme, i.e.,

$$IP_{i,s(i)} \geq IP_{s(i)}.$$

By subtracting $D[t - M_{s(i)}, t - M_i]$ from both sides of the above inequality, we have

$$IP_{i1} \geq IP_{s(i),1}.$$

Replacing i with $s(i)$, we have $IP_{s(i),1} \geq IP_{s(s(i)),1}$ and thus $IP_{i1} \geq IP_{s(s(i)),1}$. Continuing this process, we have

$$IP_{i1} \geq IP_{11} = IP_1$$

since stage 1 is the root of the assembly tree and thus a successor of any other stage. That $IP^1 = IP_1$ follows from the definition of IP^1 .

(ii) Follows from the definition of IP^i .

(iii)

$$\begin{aligned} IP^i &= \min_{n \geq i} \{IP_{ni}\} \\ &\leq \min_{n \geq i+1} \{IP_{ni}\} \\ &= \min_{n \geq i+1} \{IP_{n,i+1} - D[t - M_{i+1}, t - M_i]\} \\ &= \min_{n \geq i+1} \{IP_{n,i+1}\} - D[t - M_{i+1}, t - M_i] \\ &= IP^{i+1} - D[t - M_{i+1}, t - M_i]. \quad \square \end{aligned}$$

As in the series case, we charge the holding and backorder costs in (2) to period $t - M_N$. The following lemma parallels Lemma 1.

LEMMA 4. For Assembly (G, K) , we have $E[\sum_{i=1}^N h_i IL_i + (p + H_1)B] \geq E[\sum_{i=1}^{N-1} G_i^i(IP^i) + G_N(IP^N)]$.

PROOF. From Lemma 3(i) we have

$$E[h_1 IL_1 + (p + H_1)B] = EG_1(IP^1).$$

Notice that $G_1(IP^1) = G_1^1(IP^1) + G_1^2(IP^1) \geq G_1^1(IP^1) + G_1^2(IP^2) - D[t - M_2, t - M_1]$, where the inequality follows since $G_1^2(\cdot)$ is nonincreasing and $IP^1 \leq IP^2 - D[t - M_2, t - M_1]$ from Lemma 3(iii). Therefore,

$$\begin{aligned} E\left[\sum_{i=1}^2 h_i IL_i + (p + H_1)B\right] &\geq E[G_1^1(IP^1) + h_2 IL_2 \\ &\quad + G_1^2(IP^2 - D[t - M_2, t - M_1])]. \end{aligned}$$

Since $IL_2 = IP_2 - D[t - M_2, t - M_1] \geq IP^2 - D[t - M_2, t - M_1]$ (Lemma 3(ii)), we have

$$\begin{aligned} & E\left[\sum_{i=1}^2 h_i IL_i + (p + H_1)B\right] \\ & \geq E[G_1^1(IP^1) + h_2(IP^2 - D[t - M_2, t - M_1]) \\ & \quad + G_2^2(IP^2 - D[t - M_2, t - M_1])] \\ & = E[G_1^1(IP^1) + G_2^2(IP^2)] \end{aligned}$$

where the last equality follows from the definition of $G_2(\cdot)$. Now we have proved the lemma for $N = 2$. The above proof can be easily adapted for any value of N by induction. \square

Recall that C_i is the minimum value of $G_i^1(\cdot)$ for $i < N$ and that $IP_N = IP^N$. From Lemma 4 we have

$$\begin{aligned} E\left[\sum_{i=1}^N h_i IL_i + (p + H_1)B\right] & \geq \sum_{i=1}^{N-1} C_i + EG_N(IP^N) \\ & = \sum_{i=1}^{N-1} C_i + EG_N(IP_N). \end{aligned}$$

In other words, given $IP_N = y$, the expected system-wide holding and backorder costs charged to period $t - M_N$ under any policy are bounded below by $\sum_{i=1}^{N-1} C_i + G_N(y)$. By substituting the latter for the former and following the same argument as in the serial case, we have:

LEMMA 5. $\sum_{i=1}^N C_i$ is a lower bound on the minimum long-run average cost of Assembly(\mathcal{G}, K).

Now consider the following feasible policy for Assembly(\mathcal{G}, K). Stage N uses the echelon stock (s_N, S_N) policy: if IP_N falls to or below s_N , an order is placed with the outside supplier to raise IP_N up to S_N ; and stage $i, i = 1, \dots, N - 1$, uses the echelon stock order-up-to $\min\{Y_i, IP_{i+1,i}\}$ policy (Y_i is the minimum point of $G_i(\cdot)$): if IP_i is below $\min\{Y_i, IP_{i+1,i}\}$, order up to $\min\{Y_i, IP_{i+1,i}\}$. This policy is exactly the one characterized by Rosling (1989). Therefore, one way to show that the policy is optimal is to check that the cost function given by Rosling coincides with the lower bound. Below we present an alternative proof.

THEOREM 2. For Assembly(\mathcal{G}, K), the optimal policy uses the echelon stock (s_N, S_N) policy at stage N and the echelon stock order-up-to $\min\{Y_i, IP_{i+1,i}\}$ policy at stage

$i, i = 1, \dots, N - 1$, with minimum cost $C_N^a(K, \mathbf{h}, p, \mathbf{L}, F) = \sum_{i=1}^N C_i - \sum_{i=1}^N h_i \mu(L_i - l_i)$ (under the original accounting scheme).

PROOF. Suppose that the above policy is in place for Assembly(\mathcal{G}, K). In the long run, we have

$$IP_n \leq IP_{n+1,n} \quad \text{for } n = 1, \dots, N - 1.$$

By subtracting $D[t - M_n, t - M_i]$ from both sides of the above inequality, we have

$$IP_n \leq IP_{n+1,i} \quad \text{for } n \geq i. \tag{3}$$

With (3), it is clear that

$$IP^i = IP_i \quad \text{for } i = 1, \dots, N. \tag{4}$$

Take any $i < N$. Let $P(i)$ be the set of the immediate predecessors of stage i . Take any $m \in P(i)$. Notice that the echelon inventory position at stage i, IP_i , is constrained by the beginning echelon inventory level at stage m (under the original accounting scheme), IP_{mi} : $IP_i \leq IP_{mi}$. Considering all the immediate predecessors of stage i , we have $IP_i \leq \min\{IP_{mi} | n \in P(i)\}$. This is the only constraint on IP_i . Now consider the order-up-to (target) level for IP_i , $\min\{Y_i, IP_{i+1,i}\}$. Since $m \geq i + 1$, we have from (3) $IP_{i+1,i} \leq IP_{mi}$. Therefore, $\min\{Y_i, IP_{i+1,i}\} \leq \min\{IP_{mi} | n \in P(i)\}$, i.e. the target level is within the constraint. As a result, $IP_i = \min\{Y_i, IP_{i+1,i}\}$. Since $IP_{i+1,i} = IL_{i+1}^-$, we have

$$IP_i = \min\{Y_i, IL_{i+1}^- \} \quad \text{for } i = 1, \dots, N - 1 \tag{5}$$

With (4) and (5), the proof of Theorem 2 is identical to that of Theorem 1. (Recall that $\sum_{i=1}^N h_i \mu(L_i - l_i)$ is the amount by which the new accounting scheme overcharges the holding costs. Therefore, the minimum cost in the theorem is given under the original accounting scheme.) \square

4. Lower Bounds

This section establishes lower bounds on the minimum costs of Series(N), Assembly(\mathcal{G}), and Distribution(N) through cost allocation and physical decomposition. We first establish induced-penalty bounds by allocating loss functions, and then parameter-allocation bounds by allocating cost rates. We then show that these two classes of lower bounds can be integrated to form a class of lower bounds, or integrated bounds, that dominate

both the induced-penalty bounds and the parameter-allocation bounds.

4.1. Induced-penalty Bounds

4.1.1. Series(N). The induced-penalty bound for Serial(N) is a generalization of that for Series(N, K) (Lemma 2). To see this, we first need to revise the definitions of $G_i(\cdot)$'s, $G_i^i(\cdot)$'s and $G_i^{i+1}(\cdot)$'s used in §3.1. $G_1(\cdot)$ is the same as before. Suppose that $G_i(\cdot)$, $i < N$, is defined and is convex. Consider a single-stage inventory system with setup cost K_i , loss function $G_i(\cdot)$ and the demand of the original system. Since $G_i(\cdot)$ is convex, (s, S) policies are optimal for the system. Let (s_i, S_i) be the optimal (s, S) policy and C_i the minimum cost. Define

$$G_i^i(y) = \begin{cases} C_i & \text{if } y \leq s_i \\ G_i(y) & \text{otherwise} \end{cases}$$

$$G_i^{i+1}(y) = G_i(y) - G_i^i(y)$$

and

$$G_{i+1}(y) = E[h_{i+1}(y - D[t - M_{i+1}, t - M_i]) + G_i^{i+1}(y - D[t - M_{i+1}, t - M_i])].$$

Notice that $G_i^{i+1}(\cdot)$ is an induced-penalty cost for stage $i + 1$, i.e., the cost increment at stage i when stage $i + 1$ is unable to raise IP_i to above s_i .

From Veinott and Wagner (1965), we know that $s_i < Y_i$ where Y_i is the minimizer of $G_i(\cdot)$. From Zheng (1991), we have $G_i(s_i) \geq C_i \geq G_i(s_i + 1)$. Therefore, the convexity of $G_i(\cdot)$ implies that $G_i^{i+1}(\cdot)$ is convex and nonincreasing. Consequently, $G_{i+1}(\cdot)$ is also convex. By induction, all the $G_i(\cdot)$'s are convex. Figure 2 illustrates the above definitions.

We proceed to show that $\sum_{i=1}^N C_i$ is a lower bound for Series(N). First consider Series(2). Create components 1 and 2 and allocate costs as in §3.1 for Series(2, K) by replacing K with K_2 and using the above new definitions of $G_1^1(\cdot)$, $G_1^2(\cdot)$ and $G_2(\cdot)$. The only addition here is the setup cost K_1 , which is allocated to component 1. Now consider the component systems separately. The system of component 1 is a single-stage system with setup cost K_1 and loss function $G_1^1(\cdot)$. Note that $-G_1^1(\cdot)$ is unimodal (see Figure 2). Therefore, an (s, S) policy is optimal (Veinott 1966). It can be shown that the minimum cost of component 1 is C_1 , the min-


imum cost as if the loss function were $G_1(\cdot)$ (see Appendix A). Now consider component 2. As in §3.1, there is no incentive to hold inventory at stage 2 since $G_1^2(\cdot)$, the induced-penalty cost, is nonincreasing and there is no setup cost in shipping stock from stage 2 to stage 1. As a result, the system of component 2 reduces to a single-stage system with setup cost K_2 and loss function $G_2(\cdot)$. Therefore, the minimum cost of component 2 is C_2 . Consequently, $C_1 + C_2$ is an induced-penalty bound for Series(2). (The idea of using induced-penalty costs to create lower bounds originated in Clark and Scarf 1962. The form of the above lower bound first appeared in Atkins and De 1992.)

For Series(N), we first establish a lower bound on the expected system-wide holding and backorder costs. Similar to Lemma 1, it can be shown that

$$E\left[\sum_{i=1}^N h_i IL_i + (p + H_1)B\right] \geq E\left[\sum_{i=1}^{N-1} G_i^i(IP_i) + G_N(IP_N)\right].$$

Now suppose that the final product consists of N components. The supply system for component i ($= 1, \dots, N$) is a subsystem of the original system: from an outside supplier to stage i , to stage $i - 1, \dots$, to stage 1. Con-

Figure 2 Definition of Functions



sider the system of component i . Its system inventory position is IP_i . Now allocate $G_i^i(IP_i)$ ($G_N(IP_N)$ if $i = N$) to component i . Therefore, the system-wide holding and backorder costs for component i depend solely on its system inventory position. Also allocate K_i to component i . It is clear that the system of component i is effectively a single-stage system with loss function $G_i^i(\cdot)$ and setup cost K_i , whose minimum cost is C_i (Appendix A). Consequently, $\sum_{i=1}^N C_i$ is a lower bound for Series(N).

REMARK. Notice that echelon holding costs are a natural way to allocate installation holding costs among different stages. In the above derivation, we took this allocation as given. This does not have to be the case: the installation holding costs can be re-allocated before allocating loss functions. For example, consider Series(2). The echelon holding costs h_1 and h_2 are an allocation of the installation holding cost at stage 1, H_1 , between the two stages. For the deterministic counterpart of the two-stage system (with zero leadtime and no backlogging), Roundy (1985) observed that when $K_2/h_2 < K_1/h_1$, a better lower bound can be obtained by allocating $h'_2 (< h_2)$ to stage 2 and $h'_1 (= H_1 - h'_2)$ to stage 1. The induced-penalty bound can also be made more general by allowing a similar holding cost re-allocation. Let h'_1 and h'_2 be any such allocation of H_1 between stage 1 and stage 2. Under this allocation, the loss function at stage 1 becomes

$$\begin{aligned} \bar{G}_1(y) &= E[h'_1(y - D[t - M_1, t]) \\ &\quad + (p + H_1)(y - D[t - M_1, t])^-]. \end{aligned}$$

Let (\hat{s}_1, \hat{S}_1) be the optimal (s, S) policy and \bar{C}_1 the minimum cost for the single-stage system with loss function $\bar{G}_1(\cdot)$ and setup cost K_1 . Define

$$\bar{G}_1^1(y) = \begin{cases} \bar{C}_1 & \text{if } y \leq \hat{s}_1 \\ \bar{G}_1(y) & \text{otherwise} \end{cases}$$

$$\bar{G}_1^2(y) = \bar{G}_1(y) - \bar{G}_1^1(y)$$

and

$$\begin{aligned} \bar{G}_2(y) &= E[h'_2(y - D[t - M_2, t - M_1]) \\ &\quad + \bar{G}_1^2(y - D[t - M_2, t - M_1])]. \end{aligned}$$

Let \bar{C}_2 be the minimum cost of a single-stage system with loss function $\bar{G}_2(\cdot)$ and setup cost K_2 . Now we have a new lower bound for Series(2): $\bar{C}_1 + \bar{C}_2$. Nu-

merical examples suggest that the above generalization can lead to better bounds at the expense of additional computational effort. (The above idea of re-allocating installation holding costs before allocating loss functions can be applied to the other two systems as well.)

4.1.2. Assembly(\mathcal{G}). Following Atkins (1988), we can decompose the assembly system into a number of serial systems, each of Series(N)-type. A lower bound can thus be generated for each of the resulting serial systems (e.g., induced-penalty bounds), and the sum of these lower bounds is a lower bound on the minimum cost of the original assembly system. For example, consider a simple assembly system where two components are assembled into an end product. By allocating the setup cost, the backorder cost as well as the echelon holding cost of the end product between the two components, the assembly system decomposes to two two-stage, serial systems. The summation of the lower bounds for the resulting serial systems is a lower bound for the original system.

4.1.3. Distribution(N). For $i = 1, \dots, N$, define $G_i(y) = E[h_i(y - D_i[t, t + L_i])^+ + (h_0 + p_i)(y - D_i[t, t + L_i])^-]$.

Notice that $G_i(\cdot)$ is convex. Consider a single-stage system with setup cost K_i , loss function $G_i(\cdot)$, and the demand process of retailer i . Since $G_i(\cdot)$ is convex, (s, S) policies are optimal for the system. Let (s_i, S_i) be the optimal (s, S) policy and C_i the minimum cost. Define

$$G_i^i(y) = \begin{cases} C_i & \text{if } y \leq s_i \\ G_i(y) & \text{otherwise} \end{cases}$$

and $G_i^0(y) = G_i(y) - G_i^i(y)$. Note that the above definitions are different from those used earlier in other systems. Here $G_i^0(\cdot)$ represents an induced-penalty cost for the warehouse: the cost increment at retailer i due to the warehouse's inability to raise IP_i to above s_i .

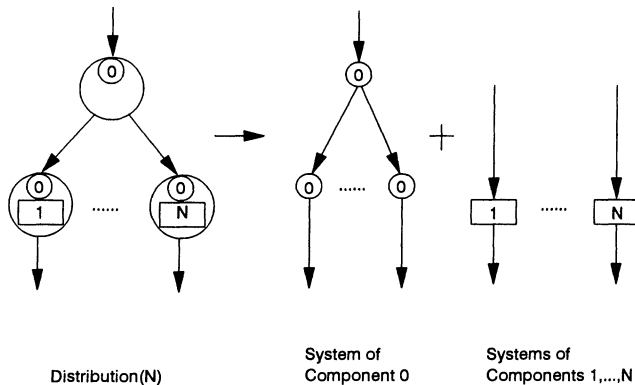
We proceed to derive an induced-penalty bound for Distribution(N). Imagine that the product at retailer i consists of two components: a common component 0 and a retailer-specific component i . Therefore, there are a total of $N + 1$ components. Component 0 is supplied through the original distribution system: from the outside supplier to the warehouse and then to the retailers.

Component i is directly shipped from an outside supplier to retailer i . Allocate setup cost K_i ($i = 0, \dots, N$) to component i and warehouse echelon holding cost rate h_0 to component 0. The loss function at retailer i , $G_i(\cdot)$, is allocated between component 0 and component i : $G_i^0(\cdot)$ to component 0 and $G_i^i(\cdot)$ to component i . In summary, the system of component 0 is $\text{Distribution}(N, K_0)$ with warehouse echelon holding cost rate h_0 (assessed on the warehouse echelon inventory level), loss function $G_i^0(\cdot)$ at retailer i (assessed on the inventory position of retailer i). The system of component $i = 1, \dots, N$ is a single-stage system with setup cost K_i , loss function $G_i^i(\cdot)$. The component systems are depicted in Figure 3.

In the original system, replenishment of the components is "coordinated": when a shipment of component 0 is released at the warehouse to retailer i , a same quantity of component i is released at a supplier to retailer i . Both shipments arrive at retailer i at the same time. After arriving at the retailer, components 0 and i are sold in packages. Now assume that the components can be replenished and sold independently of each other. As a result, the original system decomposes to $N + 1$ component systems. The minimum cost of component $i = (1, \dots, N)$ is C_i (Appendix A). A lower bound on the minimum cost of component 0, C_0 , can be obtained by replacing $G_i(\cdot)$ in (B5) in Appendix B with $G_i^0(\cdot)$. Therefore, a lower bound for $\text{Distribution}(N)$ is

$$C^{ip} \stackrel{\text{def}}{=} \sum_{i=0}^N C_i.$$

Figure 3 Decomposition of $\text{Distribution}(N)$



(Rosling (1977) obtained precisely this bound in a different form, but only under additional technical conditions.)

4.2. Parameter-allocation Bounds

The previous subsection established induced-penalty bounds by allocating loss functions, or "expected" costs. Next we derive a new class of lower bounds for $\text{Series}(N)$, $\text{Assembly}(\mathcal{G})$, and $\text{Distribution}(N)$ by allocating "real" costs, i.e., holding and backorder cost rates.

4.2.1. Series(N). First consider $\text{Series}(2)$. Create fictitious components and their systems as in Figure 1. Allocate setup cost K_i ($i = 1, 2$) to component i and echelon holding cost rate h_2 to component 2. The echelon holding cost rate h_1 and the backorder cost rate p at stage 1 are allocated between components 1 and 2: h_1^1 and p^1 to component 1; h_1^2 and p^2 to component 2. The allocated cost rates satisfy $h_1^1 + h_1^2 = h_1$ and $p^1 + p^2 = p$. Check that the sum of the allocated costs of the components is equal to the total cost of the original system. Since the minimum costs for components 1 and 2 are $C_1^s(K_1, h_1^1, p^1, L_1, F)$ and $C_2^s(K_2, h_1^2, h_2, p^2, L_1, L_2, F)$ respectively, a lower bound on the minimum cost of $\text{Series}(2)$ is

$$C_1^s(K_1, h_1^1, p^1, L_1, F) + C_2^s(K_2, h_1^2, h_2, p^2, L_1, L_2, F).$$

Consider the special cost allocation with $h_1^2 = 0$. Examine the system of component 2. In this case, since there is no setup cost or echelon holding cost at stage 1, there is no incentive to keep any stock at stage 2. In other words, the system of component 2 effectively reduces to a single-stage system with setup cost K_2 , holding cost rate h_2 , backorder cost rate p_2^2 , and transportation leadtime $L_1 + L_2$. Notice that holding costs start to accumulate at rate h_2 as soon as the component enters the system (at stage 2). The time for each unit of the component to go from stage 2 to stage 1 is L_1 periods, during which a holding cost of $h_2 L_1$ is incurred. Consequently, the long-run average holding cost incurred while shipping the component from stage 2 to stage 1 is $\mu h_2 L_1$ per period, where μ is the mean one-period demand. In summary, under the above cost allocation the parameter-allocation bound becomes

$$C_1^s(K_1, h_1, p^1, L_1, F) + C_2^s(K_2, h_2, p^2, L_2 + L_1, F) + \mu h_2 L_1.$$

With the above illustration, constructing a lower bound for Series(N) becomes straightforward. Suppose that there are N fictitious components, denoted by $1, \dots, N$. The system of component j is Series(j, K_j) with echelon holding cost rate h_i^j at stage $i, i = 1, \dots, j$, and backorder cost rate p^j (at stage 1). The cost allocation satisfies $\sum_{j=i}^N h_i^j = h_i$ for $i = 1, \dots, N$ and $\sum_{j=1}^N p^j = p$. Notice that the minimum cost for component j is

$$C_j^s(K_j, \mathbf{h}^j, p^j, \mathbf{L}^j, F),$$

where $\mathbf{h}^j = \{h_i^j\}_{i=1}^j$ and $\mathbf{L}^j = \{L_i\}_{i=1}^j$. A lower bound on the minimum cost of Series(N) is simply the sum of the minimum costs of all the components.

4.2.2. Assembly(\mathcal{G}). From §3, we know that it is straightforward to compute the minimum cost of an assembly system with a setup cost only at its highest stage, the stage with the longest total leadtime. Since there is a setup cost at each stage in Assembly(\mathcal{G}), a natural approach is to decompose the system into subsystems of Assembly(\mathcal{G}, K)-type.

Recall that there are N stages in Assembly(\mathcal{G}). The stages are numbered in increasing order of total leadtimes. (Thus stage N has the longest total leadtime.) Imagine that the end product consists of N fictitious components. Consider component j ($=1, \dots, N$). Delete stages $j + 1, \dots, N$ of \mathcal{G} . The resulting graph is still a tree, denoted by \mathcal{G}_j with $\mathcal{G}_N = \mathcal{G}$. Component j is produced through \mathcal{G}_j . Allocate K_j to stage j , echelon holding cost h_i^j to stage i ($=1, \dots, j$), and backorder cost p^j to stage 1. The cost allocation satisfies $\sum_{j=i}^N h_i^j = h_i$ for $i = 1, \dots, N$ and $\sum_{j=1}^N p^j = p$. The system of component j is of Assembly(\mathcal{G}_j, K_j)-type, with minimum cost

$$C_j^a(K_j, \mathbf{h}^j, p^j, \mathbf{L}^j, F)$$

where $\mathbf{h}^j = \{h_i^j\}_{i=1}^j$ and $\mathbf{L}^j = \{L_i\}_{i=1}^j$. The sum of these minimum costs over all components is a lower bound on the minimum cost of Assembly(\mathcal{G}).

4.2.3. Distribution(N). Create fictitious components as in §4.1.3, but allocate costs differently. The system of component 0 is Distribution(N, K_0) with echelon holding cost rate h_0 at the warehouse, and echelon holding cost rate h_i^0 and backorder cost rate p_i^0 at retailer i . The system of component $i = 1, \dots, N$ is a single-stage system with setup cost K_i , holding cost

rate h_i^i , backorder cost rate p_i^i . The cost allocation satisfies $h_i^i + h_i^0 = h_i$ and $p_i^i + p_i^0 = p_i$ for $i = 1, \dots, N$. Recall that the minimum cost of component 0 is unknown, but a lower bound on it is

$$C_N^d(K_0, h_0, \mathbf{h}^0, \mathbf{p}^0, \mathbf{L}, F),$$

where $\mathbf{h}^0 = \{h_i^0\}_1^N$ and $\mathbf{p}^0 = \{p_i^0\}_1^N$. The minimum cost of component $i = 1, \dots, N$ is

$$C_i^s(K_i, h_i^i, p_i^i, L_i, F_i)$$

where F_i is the cdf of one-period demand at retailer i . Consequently, the sum of the minimum costs of components $1, \dots, N$ and the lower bound of component 0,

$$C_N^d(K_0, h_0, \mathbf{h}^0, \mathbf{p}^0, \mathbf{L}, F) + \sum_{i=1}^N C_i^s(K_i, h_i^i, p_i^i, L_i, F_i),$$

is a lower bound on the minimum cost of Distribution(N). For easy reference, write \underline{C}^{pa} for the maximum value of the above expression over all possible cost allocations. (In the above cost allocation, we should restrict $p_i^0 \geq 0$ for all i in order to prevent the lower bound for component 0 from becoming negative infinity. But h_i^0 is allowed to be negative so long as $h_i^0 + h_0 \geq 0$.)

It is interesting to contrast our approach here with the bounding scheme of Atkins and Iyogun (1987) for one-warehouse multi-retailer systems with deterministic demand. In their approach, the warehouse setup cost is allocated to decompose the system into a number of two-stage, serial systems. However, as recognized by Atkins (1990), this method fails for not capturing the "risk pooling" function of the warehouse when demand is stochastic. Our method overcomes this difficulty: we preserve the warehouse node leaving its risk-pooling function intact but decompose the retailer nodes.

4.3. Integrated Bounds

We have obtained induced-penalty bounds by allocating loss functions and parameter-allocation bounds by allocating cost rates. These allocation schemes can be integrated to generate a class of dominant bounds. We use Distribution(N) as an illustration.

First, create fictitious components and allocate cost rates as in §4.2.3. Let \mathcal{A} denote an allocation of the holding and backorder cost rates, i.e., h_i^i, h_i^0, p_i^i and p_i^0 with $h_i^i + h_i^0 = h_i$ and $p_i^i + p_i^0 = p_i$ for $1 \leq i \leq N$. Under \mathcal{A} , the loss function for component i is

$$G_i(y|\mathcal{A}) \stackrel{\text{def}}{=} E\{h_i^i(y - D_i[t, t + L_i])^+ + p_i^i(y - D_i[t, t + L_i])^-\}$$

and the loss function for retailer i in the system of component 0 is

$$G_i^0(y|\mathcal{A}) \stackrel{\text{def}}{=} E\{h_i^0(y - D_i[t, t + L_i])^+ + (h_0 + p_i^0)(y - D_i[t, t + L_i])^-\}.$$

The system of component $i = 1, \dots, N$ is a single-stage system with setup cost K_i , loss function $G_i(\cdot|\mathcal{A})$, and the demand process of retailer i . Since $G_i(\cdot|\mathcal{A})$ is convex, (s, S) policies are optimal for the system. Let $s_i(\mathcal{A})$ be the optimal reorder point and $C_i(\mathcal{A})$ the minimum cost. Note that $C_i(\mathcal{A}) = C_1^i(K_i, h_i^i, p_i^i, L_i, F_i)$, where F_i is the cdf of one-period demand at retailer i .

Now allocate loss functions. For $i = 1, \dots, N$, define

$$G_{ii}(y|\mathcal{A}) = \begin{cases} C_i(\mathcal{A}) & \text{if } y \leq s_i(\mathcal{A}) \\ G_i(y|\mathcal{A}) & \text{otherwise} \end{cases}$$

and $G_{i0}(y|\mathcal{A}) = G_i(y|\mathcal{A}) - G_{ii}(y|\mathcal{A})$. Allocate $G_{i0}(\cdot|\mathcal{A})$ back to component 0. With this allocation, component i 's loss function ($1 \leq i \leq N$) reduces from $G_i(\cdot|\mathcal{A})$ to $G_{ii}(\cdot|\mathcal{A})$, but its minimum cost remains unchanged at $C_i(\mathcal{A})$ (Appendix A). On the other hand, the loss-function allocation increases the loss function of component 0 at retailer i from $G_i^0(\cdot|\mathcal{A})$ to $G_i^0(\cdot|\mathcal{A}) + G_{i0}(\cdot|\mathcal{A})$. As a result, the lower bound for component 0, denoted by $C_0(\mathcal{A})$, is greater than $C_N^d(K_0, h_0, \mathbf{h}^0, \mathbf{p}^0, \mathbf{L}, F)$, the lower bound for component 0 before allocating loss functions. ($C_0(\mathcal{A})$ is obtained by replacing $G_i(\cdot)$ in (B5) of Appendix B with $G_i^0(\cdot|\mathcal{A}) + G_{i0}(\cdot|\mathcal{A})$.) Therefore, by allocating cost rates as well as loss functions, we have a new lower bound for Distribution(N):

$$\sum_{i=0}^N C_i(\mathcal{A}).$$

The integrated bound, denoted by \underline{C} , is the maximum value of the above expression over all possible allocations \mathcal{A} . (In §4.2.3, we mentioned that p_i^0 should be restricted to be nonnegative. This restriction should not be used in computing the integrated bound. To see this, consider the system of component 0. Recall that under the above integrated cost allocation, the loss function at retailer i is $[G_i^0(y|\mathcal{A}) + G_{i0}(y|\mathcal{A})]$, whose slope is always $-(p_i + h_0)$ as $y \downarrow -\infty$ regardless what p_i^0 is. In

other words, a unit of backorder at retailer i ultimately increases cost by $(p_i + h_0)$. On the other hand, since the echelon holding cost at the warehouse is assessed on the warehouse echelon inventory level, a unit of backorder at retailer i decreases cost by h_0 . Thus the net penalty of a backorder at retailer i is p_i . Therefore, no matter what the value of p_i^0 is, a positive penalty is incurred for each additional unit of backorder.)

From the above derivation, it is clear that $\underline{C} \geq \underline{C}^{pa}$. It is also true that $\underline{C} \geq \underline{C}^{ip}$. To see this, consider the following allocation of cost rates \mathcal{A}_0 : $h_i^i = h_i$, $h_i^0 = 0$, $p_i^i = p_i + h_0$, and $p_i^0 = -h_0$. Under \mathcal{A}_0 , check that $\sum_{i=0}^N C_i(\mathcal{A}_0) = \underline{C}^{ip}$. Therefore,

THEOREM 3. $\underline{C} \geq \max\{\underline{C}^{pa}, \underline{C}^{ip}\}$.


5. Extension to Continuous-review Systems

We have so far focused on periodic-review systems. Parallel results can be obtained for their continuous-review counterparts with compound Poisson demand. These include the optimality results of §3 and the lower bounds of §4.

The definitions of the periodic-review systems in §2 can be easily adapted for their continuous-review counterparts. The constant transportation leadtimes can now take any nonnegative values. In a continuous-review system, the holding and backorder costs accrue continuously over time at rates proportional to inventory levels. Therefore, the cost rates h_i 's and p_i 's should be re-defined as costs per unit time. Furthermore, since a continuous-review system monitors its inventory status continuously and can make replenishment decisions at any time, the time index t in the state variables, e.g. $IP_i(t)$, can be any time epoch.

From the previous development, we see that a key building block for our results is that (s, S) policies are optimal for single-location, periodic-review systems. Fortunately, a similar result also holds for single-location, continuous-review systems with compound Poisson demand (Beckmann 1961). To extend the previous results for periodic-review systems to their continuous-review counterparts with compound Poisson demand, we only need to notice that: (a) the compound Poisson demand process is memoryless (this property is parallel to the iid-demand assumption for periodic-review systems); (b) the loss functions for

Figure 4 Frequency Diagrams of the Ratios of Parameter-Allocation Bounds to Induced-Penalty Bounds for Distribution (N). The Horizontal Axes Represent the Ratios in Percentage



periodic-review systems are now loss rate functions; and (c) the idea of cost allocation / physical decomposition does not depend on the inventory review system in place.

6. Numerical Comparisons

This section makes numerical comparisons between different lower bounds and between lower bounds and the costs of heuristic policies for Distribution(*N*) with identical retailers and simple Poisson demand.

In two separate papers (Chen and Zheng 1994, 1993), we studied a class of echelon stock (*R*, *Q*) policies for Distribution(*N*). An echelon stock (*R*, *Q*) policy is specified by *N* + 1 pairs of control parameters: (*R*_{*i*}, *Q*_{*i*})_{*i*=0}^{*N*}. The warehouse orders *Q*₀ units from the outside supplier as soon as its echelon inventory position reaches *R*₀. Retailer *i* orders *Q*_{*i*} units from the warehouse once its inventory position (= on-hand inventory + all outstanding orders – customer backorders) declines to *R*_{*i*}. In case of insufficient on-hand stock at the warehouse to satisfy a retailer order, the order can be partially filled with the remaining backordered at the warehouse. Backorders (of retailer orders) at the warehouse are satisfied on a first-come, first-served basis. Setup costs are incurred only when shipments are sent, not when orders are placed. For Distribution(*N*) with identical retailers, we assume that the retailers use identical policies, i.e. *R*₁ = . . . = *R*_{*N*} and *Q*₁ = . . . = *Q*_{*N*}, and that *Q*₀ is an integer multiple of *Q*₁. We have developed procedures to identify optimal (for *N* = 1) or close-to-optimal (for *N* > 1) echelon stock (*R*, *Q*) policies for Distribution(*N*) with identical retailers and simple Poisson demand. These procedures were used in our numerical study.

The numerical examples used in our study are summarized in Table 1, where λ₁ is the customer arrival rate at retailer 1 or any other retailer. From the table, we see that there is a total of 1,280 different combinations of parameters, representing 1,280 different examples. These examples were divided into five groups according to the value of *N*, the number of retailers. For each example, we computed the induced-penalty bound (\underline{C}^{ip}), the parameter-allocation bound (\underline{C}^{pa}), the integrated bound (\underline{C}), and the cost of the "optimal" echelon stock (*R*, *Q*) policy (\bar{C}). These values were then converted into two ratios: $\underline{C}^{pa} / \underline{C}^{ip}$ and \underline{C} / \bar{C} . A frequency diagram of $\underline{C}^{pa} / \underline{C}^{ip}$ (\underline{C} / \bar{C}) was created for each group, and depicted in Figure 4 (5). From Figure 4, we see that as the number of retailers increases, the parameter-allocation bounds go from smaller than to greater than the induced-penalty bounds. Figure 5 indicates that the gap between \underline{C} and \bar{C} is small for the first group (with only one retailer), and widens somewhat as *N* increases.

Note. In the numerical study, \underline{C}^{pa} and \underline{C} were computed by searching over α and β which parameterize the allocation of cost rates: $h_i^i = \alpha h_1$, $h_i^0 = (1 - \alpha) / h_1$, $p_i^i = \beta p_1$ and $p_i^0 = (1 - \beta) p_1$ for *i* = 1, . . . , *N*. A grid size of 0.1 was used. The search range over (α, β) was [0, 1]² ([0, 1] × [0, (p₁ + h₀) / p₁]) for parameter-allocation (integrated) bounds.

7. Concluding Remarks

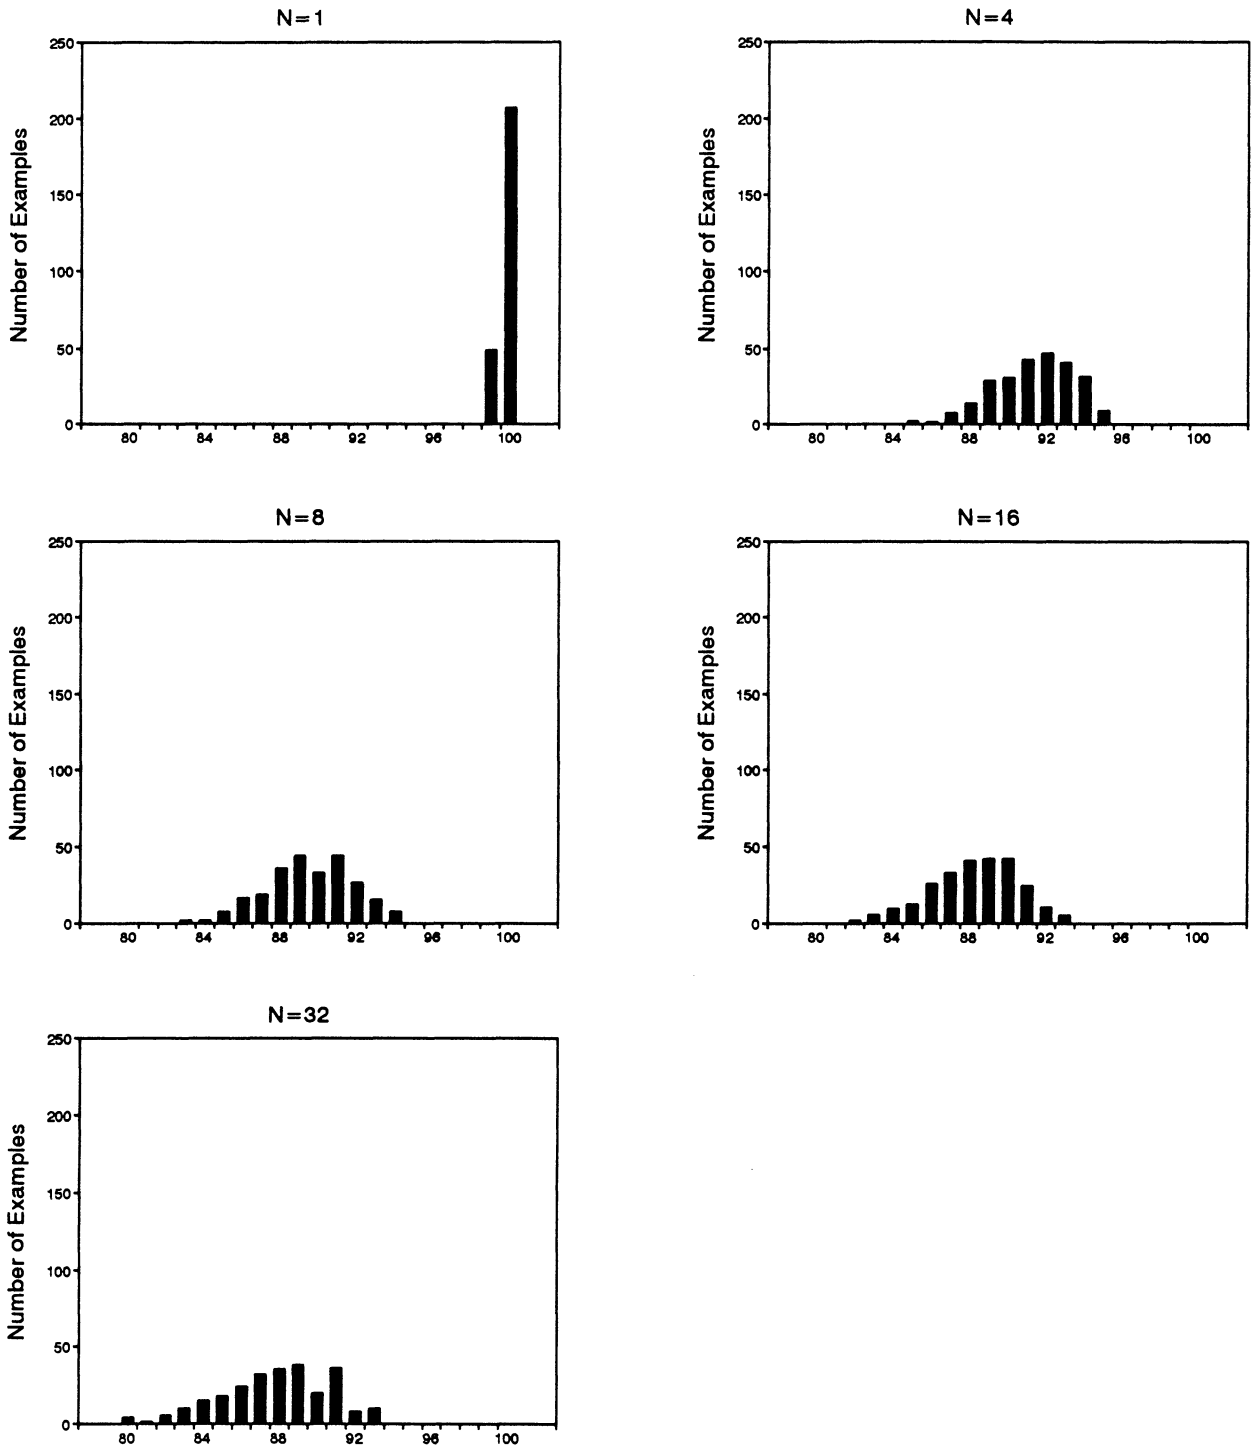
In this paper, we have provided simple proofs of known optimality results for Series(*N*, *K*) and Assembly(\mathcal{G} , *K*) and established parallel results for their continuous-review counterparts with compound Poisson demand. We have created lower bounds for Series(*N*), Assembly(\mathcal{G}), and Distribution(*N*) under a cost-allocation, physical-decomposition framework. We expect that this framework can be used to generate lower bounds for more general systems.

Remaining research issues abound: (a) an efficient algorithm is desired to find optimal cost allocations, e.g., the α and β in the numerical study; (b) a further study is needed on the bounds' effectiveness; and (c) simple, cost-effective heuristic policies are to be proposed and studied for general systems. The ultimate result is heuristic policies with guaranteed cost effectiveness. It is

Table 1 Numerical Examples

Parameter	Values
<i>N</i>	1, 4, 8, 16, 32
λ ₁ , <i>L</i> ₁ , <i>L</i> ₀	1, 2
<i>K</i> ₁	5, 10
<i>K</i> ₀	50, 100
<i>h</i> ₁	1
<i>p</i> ₁	5, 10
<i>h</i> ₀	0.5, 0.75, 1.0, 1.25

Figure 5 Frequency Diagrams of the Ratios of Integrated-Bounds to the Minimum Costs of (R, nQ) Policies for Distribution (N) . The Horizontal Axes Represent the Ratios in Percentage



our hope that the results presented in this paper will serve as a stepping stone in this endeavor.¹

¹ The authors would like to thank the Associate Editor and two referees for their constructive comments that have led to many improvements in the exposition of this paper. This research was supported in part by NSF grant DDM 9111183.

Appendix A

Consider a single-location, periodic-review system with infinite planning horizon. Demands in different periods are independent and identically distributed with one-period probability mass function p_j , $j = 0, 1, \dots$. It has setup cost K and loss function $G(\cdot)$ (a function of its inventory position). The long-run average cost of an (s, S) policy in this system is

$$c(s, S) = \frac{K + \sum_{j=0}^{S-s-1} G(S-j)m(j)}{\sum_{j=0}^{S-s-1} m(j)}$$

where

$$m(0) = \frac{1}{1-p_0} \quad \text{and} \quad m(i) = \sum_{j=0}^i p_j m(i-j) \quad i = 1, 2, \dots$$

Let Y be the minimizer of $G(\cdot)$. It follows from Zheng (1991) that there exists an (s^*, S^*) policy that satisfies

- (i) $g^* \stackrel{\text{def}}{=} c(s^*, S^*) = \min_{s < S} c(s, S)$
- (ii) $s^* < Y \leq S^*$
- (iii) $G(s^* + 1) \leq g^* \leq G(s^*)$
- (iv) $g^* \geq G(S^*)$.

Define

$$G^-(i) = \begin{cases} G(i) & i > s^* \\ g^* & i \leq s^* \end{cases}$$

Now modify the above system by replacing the loss function $G(\cdot)$ with $G^-(\cdot)$. The long-run average cost of an (s, S) policy in this modified system is

$$c^-(s, S) = \frac{K + \sum_{j=0}^{S-s-1} G^-(S-j)m(j)}{\sum_{j=0}^{S-s-1} m(j)}$$

LEMMA. $g^* = \min_{s < S} c^-(s, S)$.

PROOF. First notice that $c^-(s^*, S^*) = g^*$. Thus we only need to show that $c^-(s, S) \geq g^*$ for all $s < S$. Take any $s < S$. Consider the following cases: (1) If $s \geq s^*$ then $c^-(s, S) = c(s, S) \geq g^*$; (2) if $s < s^*$ and $S \leq s^*$ then

$$c^-(s, S) = \frac{K + g^* \sum_{j=0}^{S-s-1} m(j)}{\sum_{j=0}^{S-s-1} m(j)} > g^*;$$

and (3) if $s < s^*$ and $S > s^*$ then

$$c^-(s, S) = wc(s^*, S) + (1-w)g^* \geq g^*$$

where $w = \sum_{j=0}^{S-s-1} m(j) / \sum_{j=0}^{S-s-1} m(j) \in (0, 1]$, and the inequality follows since $c(s^*, S) \geq g^*$. \square

The above lemma shows that the (s^*, S^*) policy satisfying conditions (i)–(iv) is the optimal (s, S) policy for the above modified system (with loss function $G^-(\cdot)$). In fact, this policy is optimal among all policies.

THEOREM. The (s^*, S^*) policy is optimal for the modified system.

PROOF. Analogous to Theorem 2 of Zheng (1991). \square

Appendix B

Based on Federgruen and Zipkin (1984a, b, c), here we briefly derive $C_N^d(K, \mathbf{h}, \mathbf{p}, \mathbf{L}, F)$. Consider any feasible policy in Distribution (N, K) . The total systemwide holding and backorder costs at period t are

$$h_0 I_0(t) + \sum_{i=1}^N [h_i I_i(t) + p_i B_i(t)].$$

Since $IL_0(t) = I_0(t) - \sum_{i=1}^N B_i(t)$, the above expression can be written as

$$h_0 IL_0(t) + \sum_{i=1}^N [h_i I_i(t) + (h_0 + p_i) B_i(t)].$$

For convenience, we charge the following to period t :

$$h_0 IL_0(t) + \sum_{i=1}^N [h_i I_i(t + L_i) + (h_0 + p_i) B_i(t + L_i)]. \quad (B1)$$

It is clear that this cost accounting scheme will not affect the long-run average systemwide holding and backorder costs. Recall that $IL_0^-(t)$ is the beginning echelon inventory level at the warehouse at period t (before demand occurrence). Thus $IL_0(t) = IL_0^-(t) - D_0[t, t]$. Consequently, (B1) becomes

$$h_0 \{ IL_0^-(t) - D_0[t, t] \} + \sum_{i=1}^N [h_i I_i(t + L_i) + (h_0 + p_i) B_i(t + L_i)]. \quad (B2)$$

Notice that $I_i(t + L_i) = \{IP_i(t) - D_i[t, t + L_i]\}^+$ and $B_i(t + L_i) = \{IP_i(t) - D_i[t, t + L_i]\}^-$. Given $IL_0^-(t) = y$ and $IP_i(t) = y_i$ for $i = 1, \dots, N$, the expected value of (B2) is

$$h_0(y - \mu_0) + \sum_{i=1}^N G_i(y_i) \quad (B3)$$

where μ_0 is the mean one-period system demand and

$$G_i(y) \stackrel{\text{def}}{=} E[h_i \{y - D_i[t, t + L_i]\}^+ + (h_0 + p_i) \{y - D_i[t, t + L_i]\}^-], \quad (B4)$$

the loss function for retailer i .

Notice that, by definition, the total inventory position of the retailers cannot exceed the beginning echelon inventory level at the warehouse, i.e., $\sum_{i=1}^N y_i \leq y$. Thus for any given value of y , a lower bound on (B3) is

$$R(y) \stackrel{\text{def}}{=} h_0(y - \mu_0) + \min_{y_i: \sum_{i=1}^N y_i \leq y} \sum_{i=1}^N G_i(y_i). \quad (B5)$$

In other words, given that the beginning echelon inventory level at the warehouse is y , the expected one-period holding and backorder costs are at least $R(y)$. Notice that $R(\cdot)$ is convex and easy to compute (cf. Zipkin 1984 for further references). The minimization in (B5) is effectively a free inventory position rebalance, a technique used by Eppen and Schrage (1981) and Federgruen and Zipkin (1984a, b, c) in solving what they called the *myopic allocation problem*.

Notice that $IL_0^-(t) = IP_0(t - L_0) - D_0[t - L_0, t)$. Hence, given $IP_0(t - L_0) = z$, the expected system-wide holding and backorder costs charged to period t are at least

$$G_0(z) \stackrel{\text{def}}{=} ER(z - D_0[t - L_0, t)).$$

Since $R(\cdot)$ is convex, $G_0(\cdot)$ is convex. By substituting $G_0(z)$ for the expected system-wide holding and backorder costs at period t , the original system reduces to a single-location system with setup cost K , loss function $G_0(\cdot)$, and the aggregate demand of the original system. Since $G_0(\cdot)$ is convex, (s, S) policies are optimal for the single-location system. The cost of the optimal (s, S) policy is $C_N^*(K, \mathbf{h}, \mathbf{L}, F)$.

References

- Atkins, D., "A Lower Bound to the Dynamic Assembly Problem," Working Paper No. 1259, Faculty of Commerce and Business Administration, The University of British Columbia, Vancouver, B.C., Canada, 1988.
- , "A Survey of Lower Bounding Methodologies for Production/Inventory Models," *Annals of Oper. Res.*, 26 (1990), 9–28.
- , and S. De, "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, 1992.
- and P. Iyogun, "A Lower Bound on a Class of Coordinated Inventory/Production Problems," *O. R. Letters*, 6 (1987), 63–67.
- Beckmann, M., "An Inventory Model for Arbitrary Interval and Quantity Distributions of Demand," *Management Sci.*, 8 (1961), 35–57.
- Chen, F. and Y. S. Zheng, "Evaluating Echelon Stock (R, nQ) Policies in Production/Inventory Systems with Stochastic Demand," *Management Sci.*, 40 (1994), 1262–1275.
- and —, "One-warehouse Multi-retailer Systems with Centralized Stock Information," Working Paper, Department of Decision Sciences, The Wharton School, University of Pennsylvania, 1993.
- Clark, A. J. and H. Scarf, "Optimal Policies for a Multi-Echelon Inventory Problem," *Management Sci.*, 6 (1960), 475–490.
- and —, "Approximate Solutions to a Simple Multi-Echelon Inventory Problem," in *Studies in Applied Probability and Management Science*, pp. 88–110. K. J. Arrow, S. Karlin and H. Scarf (Eds.), Stanford University Press, Stanford, CA, 1962.
- Eppen, G. and L. Schrage, "Centralized Ordering Policies in a Multi-warehouse System with Lead Times and Random Demand," in *Multi-Level Production/Inventory Control Systems: Theory and Practice*, L. Schwarz (Ed.), North-Holland Publishing Company, Amsterdam, 1981.
- Federgruen, A. and P. Zipkin, "Approximation of Dynamic, Multi-Location Production and Inventory Problems," *Management Sci.*, 30 (1984a), 69–84.
- and —, "Computational Issues in an Infinite Horizon, Multi-Echelon Inventory Model," *Oper. Res.*, 32 (1984b), 818–836.
- and —, "Allocation Policies and Cost Approximation for Multi-Location Inventory Systems," *Naval Res. Logist. Quart.*, 31 (1984c), 97–131.
- Iglehardt, D., "Optimalities of (s, S) Inventory Policies in the Infinite Horizon Dynamic Inventory Problem," *Management Sci.*, 9 (1963), 259–267.
- Rosling, K., "A Generalization of Clark and Scarf's Approach to Multi-Echelon Inventory Control," in *Three Essays on Batch Production and Optimization*, Linköpings Tryckeri AB, Linköping, Sweden, 1977.
- , "Optimal Inventory Policies for Assembly Systems Under Random Demands," *Oper. Res.*, 37 (1989), 565–579.
- Roundy, R. O., "98% Effective Integer-Ratio Lot-Sizing for One-Warehouse Multi-Retailer Systems," *Management Sci.*, 31 (1985), 1416–1430.
- , "A 98% Effective Lot-Sizing Rule for a Multi-Product, Multi-Stage Production Inventory System," *Math. Oper. Res.*, 11 (1986), 699–727.
- Scarf, H. E., "The Optimalities of (s, S) Policies in the Dynamic Inventory Problem," in *Mathematical Methods in the Social Sciences*, K. Arrow, S. Karlin and P. Suppes (Eds.), Stanford University Press, Stanford, CA, 1960.
- Veinott, A., Jr., "On The Optimality of (s, S) Inventory Policies: New Conditions and A New Proof," *SIAM J. Appl. Math.*, 14 (1966), 1067–1083.
- and H. Wagner, "Computing Optimal (s, S) Inventory Policies," *Management Sci.*, 11 (1965), 525–552.
- Zheng, Y. S., "Replenishment Strategies for Production/Distribution Networks with General Joint Setup Costs," Ph.D. Dissertation, Columbia University, New York, 1987.
- , "A Simple Proof for the Optimality of (s, S) Policies for Infinite Horizon Inventory Problems," *J. Appl. Prob.*, 28 (1991), 802–810.
- and A. Federgruen, "Finding Optimal (s, S) Policies is About as Simple as Evaluating a Single Policy," *Oper. Res.*, 39 (1991), 654–665.
- Zipkin, P., "On the Imbalance of Inventories in Multi-Echelon Systems," *Math. Oper. Res.*, 9 (1984), 402–423.

Accepted by Awi Federgruen; received August 23, 1991. This paper has been with the authors 7 months for 3 revisions.