TUTORIALS IN OPERATIONS RESEARCH
INFORMS 2011

 | ISBN 978-0-9843378-2-8
doi: http://dx.doi.org/10.1287/educ.1110.0086

Simple Heuristics for Optimal Inventory Policies in Supply Chains

Kevin H. Shang
Fuqua School of Business, Duke University, Durham, North Carolina 27708, khshang@duke.edu

Abstract This tutorial summarizes recent development on simple heuristics for optimal inventory policies in supply chains. The considered policies are reorder point, order quantity policies and base-stock, reorder interval policies. The former (latter) policies are often implemented in a continuous-review (periodic-review) inventory system. There are fixed order costs for inventory replenishment. The proposed heuristics share a common approach that solves a set of independent, single-stage problems, whose parameters are obtained from the original system data. These heuristics not only simplify the computation and implementation but also help gain insights on managing inventory in supply chains.

Keywords inventory control; supply chain management; heuristics

1. Introduction

Matching demand with supply effectively is an objective for supply chain inventory managers. Although this objective sounds straightforward, finding a systemwide optimal solution to achieving this objective is very difficult. One difficulty is computational. In many supply chain systems, the form of the optimal policy is not known. For systems with known optimal policies, finding the systemwide solution often requires solving interrelated, recursive cost functions between stages across time. This is particularly a concern when the supply chain has different levels of information integration. The other difficulty concerns implementation: a supply chain is composed of hundreds or thousands of firms, each with its own interests. These firms may not be willing to implement the optimal solution without appropriate incentives. One way to mitigate these difficulties is to design a simple heuristic. This tutorial aims to provide a review on recent development of single-stage-based heuristics. These single-stage-based heuristics not only simplify the computation and implementation but also lead to simple coordination mechanisms that help the system achieve a near-optimal solution. These heuristics also have great potential for classroom teaching and for solving complicated interdisciplinary issues occurring in supply chains.

Multiechelon inventory models can be classified by using the following attributes: deterministic versus stochastic, and local versus centralized information. In this tutorial, we focus on serial inventory systems with stochastic demand and centralized information. In §5, we discuss a few extensions such as the inventory system with local information, distribution systems, and the system with nonstationary demand. The reader is referred to Maxwell and Muckstadt [28], Roundy [30], and Federgruen and Zheng [21] for deterministic demand models.

Companies usually employ different policies to replenish inventory. For some products, if the responsiveness of inventory replenishment is an important concern, these products are usually replenished according to continuous-review policies. A commonly implemented continuous-review policy is the reorder point, order quantity policy, or (r, Q) policy. More specifically, under the (r, Q) policy, a manager monitors an echelon inventory order position (inventory on order plus inventory on hand plus inventory in transit to and at its downstream stages minus backorders at the most downstream stage) continuously. If the inventory order position reaches r, an order of Q units is placed with its upstream stage. On the other hand, a majority of consumer goods are replenished periodically according to base-stock, reorder interval policies, or (s, T) policies. Under the (s, T) policy, a manager reviews the echelon inventory order position in every T periods and orders up to a base-stock level s. The biggest advantage of the periodic-review policy is simplicity—it is simpler to coordinate orders, shipments, and labor according to some fixed schedules. Nevertheless, the downside of the periodic-review policy is that it may not be as responsive as the continuous-review policy. Shang et al. [39] compared the (r, Q) and the (s, T) policies and assessed the value of the responsiveness.

The local and echelon (r, Q) policies have been extensively studied in the literature. Most of the work on the (r, Q) policy focuses on developing approximate or exact total cost expressions (see, e.g., Axsäter [1, 2], Axsäter and Rosling [5], Badinelli [6], Cachon [8], Chen and Zheng [15], De Bodt and Graves [18]). Axsäter and Rosling [5] showed that the local (r, nQ) policy is a special case of the echelon (r, Q) policy (and that the local policy is therefore suboptimal). Chen and Zheng [16] provided an algorithm to optimize the echelon (r, Q) policy. For fixed batch sizes, Chen [12] constructed an algorithm to find the optimal reorder points for the local (r, Q) policy. Recently, Shang et al. [39] provided a method to find the optimal batch sizes for the local (r, Q) policy. There are papers aiming to providing simple heuristics on the (r, Q) policy, e.g., Shang and Song [34] and Shang [31]. We refer the reader to Axsäter [3], Chen [13], and Simchi-Levi and Zhao [40] for a review.

The literature on the (s, T) policy is relatively sparse. The (s, T) policy was first discussed by Hadley and Whitin [27]. Naddor [29] studied the (s, T) policy for single- and multi-item systems. Cachon and Zipkin [7] studied the reorder interval policy in a distribution system. He showed that the supplier's demand variance will decline as the retailers' reorder interval becomes longer. Graves [25] provided a new approach to evaluate the cost for distribution systems under the so-called virtual allocation rule. Van Houtum et al. [41] studied a serial model and showed that the echelon (s, T) policies are optimal when the reorder intervals are fixed. They also provided an algorithm to obtain the optimal base-stock levels. Chao and Zhou [11] extended these results to batch-ordering systems. Feng and Rao [23] considered a two-stage system with echelon (s, T) policies. They provided a heuristic for setting the policy parameters. Recently, Shang and Zhou [35] provided a heuristic and an approach for obtaining the exact optimal (s, T) policy when fixed order costs are present.

The rest of this tutorial is organized as follows: §2 introduces the serial inventory system and the two policies. Section 3 considers the continuous-review model with the (r, Q) policy. Section 4 considers the periodic-review model with the (s, T) policy. Section 5 extends the basic model to a few generalizations, such as systems with local information, distribution systems, and nonstationary demand models. Section 6 concludes this tutorial and suggests future research directions on multiechelon systems. Throughout this tutorial, for brevity, we do not provide proofs for technical results but cite their sources.

2. Inventory Systems and Policies

Consider a serial inventory system with N stages. Material flows from stage N to stage $N-1$, $N-1$ to $N-2$, etc., until stage 1, where random customer demand occurs. There is a constant lead time L_j for stage j. Define $L_{[i,j]} = \sum_{k=i}^{j} L_k$. There is a linear echelon holding cost with rate h_j for stage j. Define $h_{[i,j]} = \sum_{k=i}^{j} h_k$, so the local holding cost for stage j is $h_{[j,N]}$. Unsatisfied demand is fully backlogged and incurs a linear penalty cost with rate b. The objective is to minimize the long-run average systemwide cost per unit time. We consider continuous-review and periodic-review systems.

2.1. Continuous-Review System

For the continuous-review system, we assume that the demand follows a Poisson process with arrival rate λ. Let $D[\tau)$ denote the total demand during a time interval τ. All results can be extended to compound Poisson demands. To facilitate our discussion, we define inventory variables as follows: At time t, let

$IL'_j(t) =$ local inventory level at stage j;
$IO_j(t) =$ inventory on order for stage j;
$IL_j(t) =$ echelon net inventory level at stage $j = \sum_{i=1}^{j} IL'_i(t)$;
$IOP_j(t) =$ echelon inventory on-order position at stage j;
$IP_j(t) =$ echelon inventory in-transit position at stage j.

The difference between $IOP_j(t)$ and $IP_j(t)$ is the number of outstanding orders for stage j (i.e., orders not yet filled by stage $j+1$). The inventory cost is determined by the $IL_j(t)$ values. We use IL'_j, IO_j, IL_j, and IOP_j to denote these inventory variables in steady state.

The (r,Q) policy is implemented as follows: stage j monitors $IOP_j(t)$ continuously and will place an order of size Q_j when $IOP_j(t)$ reaches the reorder point r_j. (For the compound Poisson demand case, the (r,Q) policy becomes the so-called (r,nQ) policy. The analysis for the (r,nQ) policy is essentially the same as that for the (r,Q) policy. See, e.g., Shang [31].) There is a fixed order cost k_j for each stage j whenever an order is placed. We assume that the batch sizes satisfy integer-ratio relations, i.e., $Q_j = m_j Q_{j-1}$, where $Q_j, m_j \in \mathbb{N}$, and $\mathbb{N}$ is the set of positive integers, $j = 2,\ldots,N$. Clearly, a continuous-review base-stock policy with base-stock levels s_j is a special case of the (r,Q) policy by setting $Q_j = 1$ and $s_j = r_j + 1$ for all j. Note that the policy can be implemented according to the echelon inventory in-transit position $IP_j(t)$; see Chen and Zheng [15].

We first review the existing results for policy evaluation and optimization, drawn from Chen [14]. Denote by I'_j the local on-hand inventory at stage j, and denote by B the number of backorders at stage 1. We have $I'_j = IL_j - IP_{j-1}$, $j \ge 2$, $I'_1 = [IL_1]^+$, and $B = [IL_1]^-$, where $[x]^+ = \max\{0,x\}$, and $[x]^- = \max\{0,-x\}$. The long-run average systemwide cost is

$$\begin{aligned} C(\mathbf{r},\mathbf{Q}) &= \sum_{i=1}^{N} \frac{k_j\lambda}{Q_j} + \mathsf{E}\left[\sum_{i=1}^{N} h_{[i,N]} I'_i + bB + \sum_{i=2}^{N} h_{[i,N]} D[L_{i-1}]\right] \\ &= \sum_{i=1}^{N} \frac{k_j\lambda}{Q_j} + \mathsf{E}\left[\sum_{i=1}^{N} h_i IL_i + (b + h_{[1,N]})B\right]. \end{aligned} \tag{1}$$

The first term in (1) is the average fixed order cost per unit time, and the second term is the average inventory holding cost and backorder cost per unit time. For brevity, we hereafter refer to the latter cost as "inventory-related cost."

To evaluate the cost, we only need to characterize IL_j for all j. This process starts from stage N, $N-1$, sequentially until stage 1. More specifically, $IP_N = IOP_N$, which is uniformly distributed over $\{r_N+1,\ldots,r_N+Q_N\}$. For $j = N,\ldots,1$, $IL_j = IP_j - D_j$, and for $j = N-1,\ldots,1$, $IP_j = O_j[IL_{j+1}]$, where

$$O_j[x] = \begin{cases} x, & \text{if } x \le r_j, \\ x - mQ_j, & \text{otherwise,} \end{cases} \tag{2}$$

with m being the largest integer such that $x - mQ_j > r_j$. After we characterize all IL_j, we can use (1) to find the average cost per unit time.

We provide a simple bottom-up recursion to evaluate the the total cost per period. Define *echelon* j as a subsystem that includes stage i, $i \le j$. The idea behind this recursion is that, at each iteration, we evaluate the average inventory-related costs for echelon j, referred to as $G_j(y)$, provided that stage j's echelon inventory order position $IOP_j(t)$ is equal to y, and its downstream stage $i(<j)$ follows an (r, Q) policy with parameters (r_i, Q_i). For $j = 1, \ldots, N$, let

$$G_1(y) = \mathsf{E}\big[h_1(y - D[L_1)) + (b + h_{[1,N]})(y - D[L_1))^-\big]. \tag{3}$$

For $j = 2, \ldots, N$,

$$G_j(y) = \mathsf{E}\big[h_j(y - D[L_j)) + G_{j-1}(O_{j-1}[y - D[L_j)])\big]. \tag{4}$$

Then, $C(\mathbf{r}, \mathbf{Q}) = \sum_{i=1}^{N}(k_j\lambda)/Q_j + \sum_{x=1}^{Q_N} G_N(r_N + x)]/Q_N$.

For fixed batch sizes $\mathbf{Q}$, the optimal reorder point r_j^* can be obtained recursively from (3) and (4). First, $r_1^* = \arg\min_y \bar{G}_1(y) \triangleq \sum_{x=1}^{Q_1} G_1(y + x)$. For $j = 2, \ldots, N$, suppose r_{j-1}^* is known. Substituting r_{j-1}^* for r_{j-1} in the O_{j-1} function in (4), the optimal reorder point is $r_j^* = \arg\min_y \bar{G}_j(y) \triangleq \sum_{x=1}^{Q_j} G_j(y + x)$.

To find the optimal batch sizes $\mathbf{Q}^*$, Chen and Zheng [16] developed lower and upper bounds to the total cost function. These cost bounds are a sum of N separable cost functions of batch sizes. With these results, they obtained the bounds for the optimal batch size for each stage. The optimal batch sizes can be found via enumeration. See Shang and Zhou [35] for an alternative algorithm.

2.2. Periodic-Review System

For the period-review system, we assume that the demands in different periods are independent, identically distributed, nonnegative, and integer valued. Let λ denote the mean of the one-period demand. Let $D[t, t+\tau)$ and $D[t, t+\tau]$ be the total demand during time periods $t, t+1, \ldots, t+\tau-1$ (denoted by $[t, t+\tau)$) and $t, t+1, \ldots, t+\tau$ (denoted by $[t, t+\tau]$, respectively). For brevity, we use $D[\tau)$ and $D[\tau]$ to represent τ and $\tau+1$ periods of demand (so $\mathsf{E}[D[\tau)] = \tau\lambda$).

We consider the (s, T) policy in the same serial system. The policy is operated as follows: stage j orders at the beginning of every T_j periods. If the echelon inventory order position is less than an echelon base-stock level s_j, the stage orders to bring the inventory order position back to s_j. Clearly, when $T_j = 1$, the (s, T) policy reduces to the period-review base-stock policy. We refer to these T_jth periods as order periods and to T_j as the reorder interval. The reorder intervals follow integer-ratio relations: $T_{j+1} = n_j T_j$, where $T_j, n_j \in \mathbb{N}$, $j = 1, \ldots, N-1$. In addition to linear echelon holding and backorder costs (denoted by h_j and b, respectively), there is a fixed review cost K_j for each inventory reorder (i.e., K_j is incurred every T_j periods). The review cost may include inventory review costs and shipping costs. We assume that all shipments are synchronized; that is, a downstream stage, whenever possible, places an order when its upstream stage receives a shipment. (A synchronized shipping policy dominates a nonsynchronized one; see Chao and Zhou [11].) The objective is to find the (s, T) policy such that the average total cost per period is minimized.

We assume that the replenishment activities in a period occur at the beginning of the period. At stage $j > 1$, they occur in the following sequence: (1) an order, if any, from stage $j-1$ is received; (2) an order is placed with stage $j+1$ if the period is in stage j's order period; (3) a shipment, if any, sent from stage $j+1$ τ_j periods ago is received; and (4) a shipment is sent to stage $j-1$ if the period is in stage $j-1$'s order period. For stage 1, order placement occurs at the beginning of stage 1's order periods, whereas customer demand arrives during a period. We assume that the stages perform these events sequentially, from stage 1, stage 2, etc., until stage N. Costs are evaluated at the end of a period.

Define the following inventory variables:

$IOP_j(t) =$ echelon inventory order position after ordering and before demand at stage j at the beginning of an order period t;

$IP_j(t) =$ echelon inventory in-transit position after ordering and before demand at stage j at the beginning of an order period t;
$IL_j^-(t) =$ echelon inventory level at stage j at the beginning of a period t;
$IL_j(t) =$ echelon inventory level at stage j at the end of a period t.

Here, IOP_j and IP_j are defined for each order period; IL_j^- and IL_j are defined for all periods.

We now discuss how to evaluate the average total cost per period under the (s,T) policy. This total cost includes two parts, the average inventory holding and backorder cost (inventory-related cost) per period and the average review cost per period. We first show how to evaluate the former.

Consider the dynamics of the echelon inventory variables under the (s,T) policies. Suppose that stage N places an order at the beginning of an order period t. Define a *cycle* for stage j, $j = 1,\ldots,N$, with respect to t as a time interval that includes periods $t + L_{[j,N]} + \tau$, $\tau = 0,\ldots,T_N - 1$. As we shall see below, this order will directly or indirectly determine IL_j^- and IL_j within stage j's cycle. Because the system repeats itself when stage N places an order in every T_N periods, it is a regenerative process with a cycle length of T_N periods. Thus, the long-run average inventory-related cost per period is equal to the expected inventory-related cost incurred in the cycle divided by T_N. Because the expected cost is determined by IL_j, we show below how to derive IL_j within the cycle.

We start from stage N. Suppose that stage N orders at the beginning of an order period t, and the echelon inventory order position after ordering is $IOP_N(t)$. Because stage N has ample supply, $IP_N(t) = IOP_N(t)$. This order will arrive at stage N at period $t + L_N$. Because there will be no other order periods until period $t + T_N$, $IP_N(t)$ will determine both $IL_N^-(t + L_N + \tau)$ and $IL_N(t + L_N + \tau)$ for $\tau = 0,\ldots,T_N - 1$; that is,

$$IL_N^-(t + L_N + \tau) = IP_N(t) - D[t, t + L_N + \tau),$$

and

$$IL_N(t + L_N + \tau) = IP_N(t) - D[t, t + L_N + \tau].$$

Now consider stage $j = N-1, N-2,\ldots,1$ sequentially. Define $\lfloor a \rfloor$ as a roundoff operator, which returns the greatest integer less than or equal to a, a real number. Let $\mathbb{M}_x(y)$ be an operator that returns the remainder of y divided by x, $x \in \mathbb{N}$, and $y \in \{0, \mathbb{N}\}$. According to the synchronized replenishment rule, stage j will order in periods $t + L_{[j+1,N]} + \lfloor \tau/T_j \rfloor T_j$, for $\tau = 0,\ldots,T_N - 1$. Because stage j may not have ample supply, IP_j is determined jointly by its echelon inventory order position IOP_j and stage $j+1$'s net echelon inventory level IL_{j+1}^-; that is, for $\tau = 0,\ldots,T_N - 1$,

$$IP_j\left(t + L_{[j+1,N]} + \left\lfloor \frac{\tau}{T_j} \right\rfloor T_j\right) = \min\left\{ IL_{j+1}^-\left(t + L_{[j+1,N]} + \left\lfloor \frac{\tau}{T_j} \right\rfloor T_j\right), s_j \right\}. \tag{5}$$

Equation (5) means that if stage $j+1$ has sufficient stock such that $IL_{j+1}^- > s_j$, IP_j is equal to IOP_j. Otherwise, stage $j+1$ will ship as much as possible, in which case $IP_j = IL_{j+1}^-$.

The IP_j in the order periods will further determine IL_j^- and IL_j within periods $t + L_{[j,N]} + \tau$, $\tau = 0,\ldots,T_N - 1$:

$$\begin{aligned} IL_j^-(t + L_{[j,N]} + \tau) = {} & IP_j\left(t + L_{[j+1,N]} + \left\lfloor \frac{\tau}{T_j} \right\rfloor T_j\right) \\ & - D\left[t + L_{[j+1,N]} + \left\lfloor \frac{\tau}{T_j} \right\rfloor T_j, t + L_{[j,N]} + \left\lfloor \frac{\tau}{T_j} \right\rfloor T_j + \mathbb{M}_{T_j}(\tau)\right); \end{aligned} \tag{6}$$

$$\begin{aligned} IL_j(t + L_{[j,N]} + \tau) = {} & IP_j\left(t + L_{[j+1,N]} + \left\lfloor \frac{\tau}{T_j} \right\rfloor T_j\right) \\ & - D\left[t + L_{[j+1,N]} + \left\lfloor \frac{\tau}{T_j} \right\rfloor T_j, t + L_{[j,N]} + \left\lfloor \frac{\tau}{T_j} \right\rfloor T_j + \mathbb{M}_{T_j}(\tau)\right]. \end{aligned} \tag{7}$$

We write $IL_j(\tau)$ to represent $IL_j(t+L_{[j,N]}+\tau)$ in steady state. The long-run average inventory-related costs per period are equal to

$$G(\mathbf{s},\mathbf{T}) = \frac{1}{T_N}\,\mathsf{E}\left[\sum_{\tau=0}^{T_N-1}\left(\sum_{j=1}^{N} h_j\, IL_j(\tau) + (b+h_{[1,N]})[IL_1(\tau)]^-\right)\right], \tag{8}$$

where $\mathbf{s}=(s_1,\ldots,s_N)$ and $\mathbf{T}=(T_1,\ldots,T_N)$.

Below we provide a similar bottom-up recursion to conveniently evaluate $G(\mathbf{s},\mathbf{T})$. The idea behind this scheme is similar to that for the (r,Q) policy: at each iteration, we evaluate the average inventory-related costs for echelon j, referred to as $G_j(y)$, provided that stage j's echelon inventory order position $IOP_j(t)$ is equal to y and its downstream stage $i(<j)$ follows an (s,T) policy with parameters (s_i,T_i).

Proposition 1 (Shang and Zhou [35]). *Define*

$$G_1(y) = \frac{1}{T_1}\left(\sum_{\tau=0}^{T_1-1}\mathsf{E}\big[h_1(y-D[L_1+\tau]) + (b+h_{[1,N]})(y-D[L_1+\tau])^-\big]\right). \tag{9}$$

For $j=2,\ldots,N$, define recursively

$$G_j(y) = \frac{1}{T_j}\sum_{\tau=0}^{T_j-1}\mathsf{E}\Bigg[h_j(y-D[L_j+\tau]) + G_{j-1}\left(\min\left\{s_{j-1},\left(y-D\left[L_j+\left\lfloor\frac{\tau}{T_{j-1}}\right\rfloor T_{j-1}\right)\right)\right\}\right)\Bigg]. \tag{10}$$

Then, $G(\mathbf{s},\mathbf{T}) = G_N(s_N)$.

We next determine the average fixed costs per period. For stage j, the review cost K_j is incurred for each T_j. So the average cost per period is K_j/T_j. The total cost per period is

$$C(\mathbf{s},\mathbf{T}) = \sum_{j=1}^{N}\left(\frac{K_j}{T_j}\right) + G_N(s_N). \tag{11}$$

Similar to the (r,Q) policy, for fixed reorder intervals $\mathbf{T}$, the optimal base-stock levels can be obtained recursively. First, let $s_1^* = \arg\min_y G_1(y)$. For $j=2,\ldots,N$, suppose s_{j-1}^* is known, and we substitute s_{j-1}^* for s_{j-1} in (10), and let $s_j^* = \arg\min_y G_j(y)$. Then, $(s_1^*,\ldots,s_N^*)$ are the optimal base-stock levels.

Finding the optimal reorder intervals is more complicated. Shang and Zhou [35] decomposed the total cost function with an induced penalty cost function, i.e., the penalty cost charged to an upstream stage if the stage cannot fulfill its downstream stage's order. Then, they developed bounds to the induced penalty cost function by regulating downstream reorder intervals. With these steps, they obtained solution bounds for the optimal reorder intervals. An enumeration on the feasible region yields the optimal reorder intervals.

3. Continuous-Review System

This section presents single-stage-based heuristics for the continuous-review (r,Q) system. We consider two model setups for the (r,Q) policy. Section 3.1 assumes that the batch sizes are fixed. We aim to provide approximations for the optimal reorder points. Section 3.2 assumes batch sizes are variables. We provide two simple heuristics for the optimal batch sizes.

3.1. Heuristics for Optimal Reorder Points

From the recursions (3) and (4), we observe that for each stage j, the optimal reorder point r_j^* does not depend on the decisions at upstream stages. To determine r_j^*, the echelon j manager only needs to know b, $h_{[1,N]}$ and the parameters within his echelon: $(r_i^*, Q_i, D[L_i), h_i)$, for $i < j$.

More specifically, conditioning on $IP_j = y$, let $I_i'(y)$ denote the local on-hand inventory at stage i, $i \le j$, and let $B(y)$ denote the number of backorders at stage 1, assuming the echelon (r_i^*, Q_i) policy is employed at stage i, $i < j$. Denote

$$\tilde{D}_j = \sum_{i=1}^{j} D[L_i) = D[L_{[1,j]}).$$

Then, we can obtain the following decomposition for $G_j(y)$.

Proposition 2 (Shang and Song [34]). *For each $j \ge 2$ and conditioning on $IP_j = y$, $G_j(y)$ is the average inventory holding and backorder cost for echelon j assuming the echelon (r_i^*, Q_i) policy at stage i, $i < j$ is employed, and $G_j(y) = G_j'(y) + \tau_j$, where*

$$G_j'(y) = \mathsf{E}\big[h_j I_j'(y) + h_{[j,j-1]} I_{j-1}'(y) + \cdots + h_{[1,j]} I_1'(y) + (b + h_{[j+1,N]}) B(y)\big], \tag{12}$$

$$\begin{aligned} \tau_j &= \textit{the average in-transit holding cost in echelon } j \\ &= \sum_{i=2}^{j} (h_{[i,j]})\, \mathsf{E}[D_{i-1}) = \sum_{i=2}^{j} h_i\, \mathsf{E}[\tilde{D}_{i-1}). \end{aligned} \tag{13}$$

Comparing with (1), Proposition 2 implies that, under the echelon policy (r_i^*, Q_i) for $i < j$, the echelon j manager in effect faces a system with the local holding cost rate $h_{[i,j]}$ for stage i, $i \le j$, and backorder cost rate $(b + h_{[j+1,N]})$ at stage 1. In other words, the system has exactly the same structure as a truncated j-stage system consisting of the stages $1, 2, \ldots, j$ of the original system, but now stage j has ample supply. For convenience, we term this system as echelon j.

Now, for any y, according to (12), if we replace the different holding cost rates at different stages in echelon j by a single common value, then there would be no incentive to carry inventories at stage i, $2 \le i \le j$, and the j-stage system would collapse into a single-stage base-stock system with the total lead time $L_{[1,j]}$ and base-stock level y.

By setting this single common value to the minimum holding cost rate h_j, we then obtain a single-stage, lower-bound base-stock system, whose cost function is

$$\begin{aligned} G_j^\ell(y) &= \mathsf{E}\big[h_j (y - \tilde{D}_j)^+ + (b + h_{[j+1,N]})(y - \tilde{D}_j)^-\big], \\ \underline{\tau}_j &= h_j\, \mathsf{E}[\tilde{D}_{j-1}]. \end{aligned} \tag{14}$$

Similarly, by setting the single holding cost rate value to the maximum holding cost rate $h_{[1,j]}$ in echelon j, we obtain a single-stage, upper-bound base-stock system, whose cost function is

$$\begin{aligned} G_j^u(y) &= \mathsf{E}\big[h_{[1,j]} (y - \tilde{D}_j)^+ + (b + h_{[j+1,N]})(y - \tilde{D}_j)^-\big], \\ \bar{\tau}_j &= h_{[1,j]}\, \mathsf{E}[\tilde{D}_{j-1}]. \end{aligned} \tag{15}$$

In other words, for any given y, $G_j(y)$ is bounded by two single-stage base-stock cost functions as follows:

$$G_j^\ell(y) + \underline{\tau}_j \le G_j(y) \le G_j^u(y) + \bar{\tau}_j. \tag{16}$$

Recall that stage j follows an (r_j, Q_j) policy. With ample supply for echelon j, IP_j is uniformly distributed over $\{r_j+1, \ldots, r_j+Q_j\}$. Taking expectations on all terms in (16) over this uniform distribution, we have

$$\frac{1}{Q_j}\sum_{x=1}^{Q_j} G_j^{\ell}(y+x) + \underline{\tau}_j \le \frac{1}{Q_j}\sum_{x=1}^{Q_j} G_j(y+x) \le \frac{1}{Q_j}\sum_{x=1}^{Q_j} G_j^{u}(y+x) + \bar{\tau}_j; \tag{17}$$

that is, the average echelon j cost is bounded by the costs of the lower-bound (r, Q) system $R(h_j, \tilde{D}_j, b + h_{[j+1,N]})$ and the upper-bound (r, Q) system $R(h_{[1,j]}, \tilde{D}_j, b + h_{[j+1,N]})$.

Observe that the average in-transit holding cost τ_j in the original system is independent of the choice of policy. One may wonder whether $\underline{\tau}_j$ and $\bar{\tau}_j$ can be replaced by τ_j in (16) and (17). Because $\underline{\tau}_j \le \tau_j \le \bar{\tau}_j$, this implies tighter bounds. Theorem 1 below shows that this is indeed true.

Theorem 1 (Shang and Song [34]). *For $j = 1, \ldots, N$, (1) $G_j^u(y) + \tau_j \ge G_j(y) \ge G_j^{\ell}(y) + \tau_j$ for all y, and (2) $r_j^{\ell} \le r_j^* \le r_j^u$. When $j = 1$, the above inequalities reduce to equalities.*

In general, there are two ways of constructing approximations for the optimal reorder points. (i) According to Theorem 1, any convex combination of r_j^{ℓ} and r_j^u can be used to approximate r_j^*. (ii) We can replace the coefficients of $I_i'(y)$ in Proposition 4 by a single convex combination of them to obtain a base-stock system (similar to the bounding systems) and use its solution as the approximate solution. In principle, a numerical experiment can be carried out to identify effective weights. Shang and Song [33, 34] found that using a weight equal to 0.5 leads to an effective solution in most cases.

3.2. Heuristics for Optimal Batch Sizes

Chen and Zheng [16] constructed cost bounds to replace the average total cost in the objective function. These cost bounds are a sum of separable functions of batch sizes. Thus, a standard clustering algorithm (e.g., Maxwell and Muckstadt [28]) can be applied to find the optimal solution for these revised problems. They propose heuristics by converting these optimal solutions to power-of-two batch sizes.

Here, we propose two heuristics that employ the same two steps as the heuristic for the deterministic model. As demonstrated by Shang [31], these heuristics outperform Chen and Zheng's [16] heuristics in a numerical study.

These two steps are *clustering* and *minimization*. In the clustering step, the stages are grouped into disjoint clusters $\{c(1), c(2), \ldots, c(M)\}$ according to cost ratios. (Let $S = \{1, 2, \ldots, N\}$. For any $i, j \in S$ with $i \le j$, the set $\{i, i+1, \ldots, j\}$ is called a cluster. These consecutive stages will use the same batch size.) Specifically, define

$$h[m] = \sum_{i\in c(m)} h_i, \quad h'[m] = \sum_{i\in c(m)} h_{[i,N]}, \quad \text{and} \quad k[m] = \sum_{i\in c(m)} k_i.$$

These clusters satisfy the following two conditions:

(i) $k[1]/h[1] < \cdots < k[M]/h[M]$;

(ii) for each cluster $c(m) = \{l_1, \ldots, l_2\}$, there does not exist an l with $l_1 \le l < l_2$ so that $k[m^-]/h[m^-] < k[m^+]/h[m^+]$, where $c(m^-) = \{l_1, \ldots, l\}$ and $c(m^+) = \{l+1, \ldots, l_2\}$.

Through a two-dimensional diagram suggested by Zipkin [42], the above clusters $\{c(1), c(2), \ldots, c(M)\}$ can be conveniently identified.

In the minimization step, a single-stage problem is solved for each cluster $c(m)$ sequentially, starting with $m = 1$. In each problem, the solution of batch size $Q_{c(m)}$ is restricted to be an integer multiple of $Q_{c(m-1)}$, $m \ge 2$. These two heuristics are different in how the single-stage systems are formed.

Heuristic 1. $Q_{c(m)}$ is the solution of the following problem:

$$\begin{aligned} \min_Q \quad & \left\{ \frac{\lambda\mu k[m] + \sum_{x=1}^{Q} G^{\ell}_{c(m)}(R^{\ell}(Q)+x)}{Q} \right\}, \\ \text{s.t.} \quad & Q = qQ_{c(m-1)}, \quad q \in \Im^{+}, \, m > 1, \end{aligned} \tag{18}$$

where

$$R^{\ell}(Q) = \arg\min_y \left\{ \sum_{x=1}^{Q} G^{\ell}_{c(m)}(y+x) \right\}, \quad G^{\ell}_{c(m)}(y) = \sum_{i \in c(m)} G^{\ell}_i(y),$$

and

$$G^{\ell}_i(y) = \mathsf{E}\big[h_i(y - \tilde{D}_i) + (b + h_{[i,N]})(y - \tilde{D}_i])^{-}\big]. \tag{19}$$

Heuristic 2. Suppose that $c(m)$ contains stage $i \in \{v, v+1, \ldots, v+n(m)-1\}$. Alternatively, $Q_{c(m)}$ is the solution of the following problem:

$$\begin{aligned} \min\nolimits_Q \quad & \left\{ \frac{\lambda\mu k[m] + \sum_{x=1}^{Q} G_{c(m)}(R(Q)+x)}{Q} \right\}, \\ \text{s.t.} \quad & Q = qQ_{c(m-1)}, \quad q \in \Im^{+}, \, m > 1, \end{aligned} \tag{20}$$

where

$$R(Q) = \arg\min_y \left\{ \sum_{x=1}^{Q} G_{c(m)}(y+x) \right\},$$

and

$$G_{c(m)}(y) = \mathsf{E}\big[h[m](y - \tilde{D}_{v+n(m)-1}) + (n(m)b + h'[m])(y - \tilde{D}_{v+n(m)-1})^{-}\big].$$

The $Q_{c(m)}$ is solved sequentially, starting with $m = 1$. Then, set $Q^a_i = Q_{c(m)}$ for $i \in c(m)$; $(Q^a_1, \ldots, Q^a_N)$ are the heuristic batch sizes. One can apply the recursion in §3.1 to find the corresponding heuristic reorder points $(r^a_1, \ldots, r^a_N)$ or the single-stage approximations for the reorder points.

4. Periodic-Review System

This section considers the periodic-review (s, T) policy. We also consider two model setups. Section 4.1 assumes that the reorder intervals are fixed, and we present heuristics for the optimal base-stock levels. Section 4.2 assumes that the reorder intervals are decision variables, and we provide heuristics for the optimal reorder intervals.

4.1. Heuristics for Optimal Base-Stock Levels

This section develops bounds and approximations for the optimal base-stock levels $(s^*_1, \ldots, s^*_N)$ when $\mathbf{T}$ is fixed. We first develop cost bounds for each echelon. More specifically, the cost bounds for echelon j are obtained from a single-stage (s, T) system that has the original system parameters. The rationale for constructing these echelon cost bounds is the same as that for the system with (r, Q) policies described in §3.1. Briefly, consider echelon j that includes stages $1, 2, \ldots, j$. If we restrict the local holding cost rate for each stage in this echelon to the same value, there will be no incentive to stock inventories in stages $i = 2, \ldots, j$. Thus, the echelon system collapses into a single-stage system with lead time $L_{[1,j]}$. The lower (upper) bound is obtained by undercharging (overcharging) the holding cost rate h_j (respectively, $h_{[1,j]}$) to each stage.

The resulting echelon cost bound functions are defined below. For $j = 1, \ldots, N$, let

$$G_j^\ell(y) = \frac{1}{T_j} \sum_{\tau=0}^{T_j-1} \mathsf{E}[h_j(y - D[L_{[1,j]} + \tau]) + (b + h_{[j,N]})(y - D[L_{[1,j]} + \tau])^-], \tag{21}$$

$$G_j^u(y) = \frac{1}{T_j} \sum_{\tau=0}^{T_j-1} \mathsf{E}[h_{[1,j]}(y - D[L_{[1,j]} + \tau]) + (b + h_{[1,N]})(y - D[L_{[1,j]} + \tau])^-]. \tag{22}$$

Let

$$s_j^u = \arg\min_y G_j^\ell(y), \qquad s_j^\ell = \arg\min_y G_j^u(y).$$

Also, define the average pipeline inventory cost (or the average inventory in-transit cost) as $\pi_j = \sum_{i=2}^j (h_i \mathsf{E}[D[L_{[1,i-1]})])$. We have the following proposition:

Proposition 3 (Shang and Zhou [35]). *For $j = 1, \ldots, N$,*

(1) $G_j^\ell(y) + \pi_j \le G_j(y) \le G_j^u(y) + \pi_j$;
(2) $s_j^\ell \le s_j^* \le s_j^u$.

Similar to the (r, Q) policy, an effective heuristic for the base-stock levels can be found by weighted averaging s_j^ℓ and s_j^u or h_j and $h_{[1,j]}$.

4.2. Heuristics for Optimal Reorder Intervals

We provide two heuristics to find effective reorder intervals. Both heuristics solve a sum of N separable cost functions of reorder intervals subject to the integer-ratio constraints. The separable cost functions in Heuristic 1 are developed by constructing the bounds for induced-penalty cost functions; those in Heuristic 2 are the single-stage cost functions developed in §4.1.

Heuristic 1. We develop lower and upper bounds on $C(\mathbf{T})$. These cost bounds are a sum of N separable cost functions. Constructing the cost bound functions includes three steps: First, we decompose $C(\mathbf{T})$ into costs associated with each stage by using an induced-penalty cost function. Second, we construct bounds for the induced-penalty function. By substituting these bounds for the exact induced-penalty function, we effectively establish bounds for $C(\mathbf{T})$. Last, these total cost bounds will be used to derive a heuristic.

Decomposition of the Total Cost Function. The decomposition of $C(\mathbf{T})$ is based on the construction of an induced-penalty cost function, which is the penalty cost charged to an upstream stage if the upstream stage cannot fulfill an order from its immediate downstream stage. Thus, the induced-penalty cost is incurred in each downstream stage's order period.

Let us start by computing the induced-penalty cost charged to stage 2. Define $\mathbf{T_j} = (T_1, T_2, \ldots, T_j)$. Consider an order period t for stage 1. Conditioning on $IP_1(t)$, stage 1's inventory holding and backorder cost per period is

$$g_1(IP_1(t), T_1) = \frac{1}{T_1}\left(\sum_{\tau=0}^{T_1-1} \mathsf{E}[h_1(IP_1(t) - D[L_1 + \tau]) + (b + h_{[1,N]})(IP_1(t) - D[L_1 + \tau])^-]\right).$$

Because stage 2 will be charged for unfilled orders placed by stage 1, the optimal reorder point for stage 1 is a solution that minimizes stage 1's cost, assuming that it has ample supply from stage 2. In such case, $IP_1(t)$ is equal to $IOP_1(t) = s_1$. The average inventory holding and backorder cost per period for stage 1 is

$$\mathsf{E}[g_1(IOP_1(t), T_1)] = g_1(s_1, T_1).$$

Thus, stage 1 will choose the optimal reorder point $s_1(\mathbf{T_1})(= s_1^*)$ as its base-stock level.

However, stage 1 may not be able to obtain ample supply from stage 2, namely, $IP_1(t)$ is constrained by $IL_2^-(t)$. Specifically, $IP_1(t)=\min\{s_1(\mathbf{T_1}), IL_2^-(t)\}$. The induced-penalty cost charged to stage 2 is

$$g_{1,2}(IL_2^-(t),\mathbf{T_1}) = g_1\big(\min\{s_1(\mathbf{T_1}), IL_2^-(t)\},\mathbf{T_1}\big) - g_1\big(IOP_1(t),\mathbf{T_1}\big).$$

We explain why $g_{1,2}$ is the penalty cost charged to stage 2 for holding inadequate stock. If stage 2 has sufficient stock such that $IL_2^-(t) > s_1(\mathbf{T_1})$, $IP_1(t)=IOP_1(t)=s_1(\mathbf{T_1})$. In this case, there is no induced-penalty cost charged to stage 2. On the other hand, if $IL_2^-(t)\le s_1(\mathbf{T_1})$, $IP_1(t)=IL_2^-(t)$, and $g_1(IL_2^-(t),\,\mathbf{T_1}) > g_1(IOP_1(t),\,\mathbf{T_1})$. The difference will be the induced-penalty cost charged to stage 2.

Now we can compute stage 2's average inventory and penalty cost per period. Stage 2 orders every T_2 periods. Consider an order period t for stage 2. Conditioning on $IP_2(t)$, the inventory and penalty cost per period for stage 2 is

$$\begin{aligned} &g_2(IP_2(t),\mathbf{T_2})\\ &\quad=\frac{1}{T_2}\left(\sum_{\tau=0}^{T_2-1}\mathsf{E}\left[h_2(IP_2(t)-D[L_2+\tau]) + g_{1,2}\left(IP_2(t)-D\left[L_2+\left\lfloor\frac{\tau}{T_1}\right\rfloor T_1\right),\mathbf{T_1}\right)\right]\right). \end{aligned}$$

Here, $IP_2(t) - D[L_2 + \lfloor\tau/T_1\rfloor T_1)$ is IL_2^- at the beginning of stage 1's order period $t + L_2 + \lfloor\tau/T_1\rfloor$, $\tau = 0,\ldots,T_2-1$. If stage 2 has ample supply from stage 3, $IP_2(t)$ will be $IOP_2(t)=s_2$. Stage 2's average inventory and penalty cost per period is

$$\mathsf{E}[g_2(IP_2(t),\mathbf{T_2})] = g_2(s_2,\mathbf{T_2}).$$

Following the same logic, because $IP_2(t)$ is in fact constrained by $IL_3^-(t)$, we can derive the induced-penalty cost charged to stage 3 in each stage 2 order period t and obtain the average inventory holding and penalty cost per period for stage 3. The same procedure can be carried out for the rest of the chain.

We generalize the above procedure below. For stage $j=2,3,\ldots,N$, let $IL_j^-(t)=y$, where t is stage $j-1$'s order period. Suppose that $s_{j-1}(\mathbf{T_{j-1}})$ (which can be proven to be equal to s_{j-1}^*) is known. The induced-penalty cost charged to stage j is

$$g_{j-1,j}(y,\mathbf{T_{j-1}}) = g_{j-1}\big(\min\{y, s_{j-1}(\mathbf{T_{j-1}})\}, T_{j-1}\big) - g_{j-1}\big(s_{j-1}(\mathbf{T_{j-1}}),\mathbf{T_{j-1}}\big). \tag{23}$$

Conditioning on $IP_j(t)=y$, the inventory and penalty cost per period for stage j is

$$\begin{aligned} &g_j(y,\mathbf{T_j})\\ &\quad=\frac{1}{T_j}\left(\sum_{\tau=0}^{T_j-1}\mathsf{E}\left[h_j\big(y-D[L_j+\tau]\big) + g_{j-1,j}\left(y-D\left[L_j+\left\lfloor\frac{\tau}{T_{j-1}}\right\rfloor T_{j-1}\right),\mathbf{T_{j-1}}\right)\right]\right). \end{aligned} \tag{24}$$

If stage j has ample supply from stage $j+1$, its average inventory and penalty cost per period is

$$g_j(s_j,\mathbf{T_j}) = \frac{1}{T_j}\sum_{\tau=0}^{T_j-1}\mathsf{E}\big[h_j\big(s_j - D[L_j+\tau]\big)\big] + P_j(s_j,\mathbf{T_j}), \tag{25}$$

where the average penalty cost per period is

$$P_j(y,\mathbf{T_j}) = \frac{1}{T_j}\sum_{\tau=0}^{T_j-1}\mathsf{E}\left[g_{j-1,j}\left(y-D\left[L_j+\left\lfloor\frac{\tau}{T_{j-1}}\right\rfloor T_{j-1}\right),\mathbf{T_{j-1}}\right)\right]. \tag{26}$$

Define the cost for stage j in which the optimal base-stock level $s_j(\mathbf{T_j})$ is implemented as

$$g_j(\mathbf{T_j}) = g_j(s_j(\mathbf{T_j}), \mathbf{T_j}).$$

Recall that in §4.1 the recursion in (9) and (10) can be used for solving s_j^* or $s_j(\mathbf{T_j})$. Let $G_j(\mathbf{T_j})$ denote the echelon inventory-related cost function when the optimal base-stock level $s_i(\mathbf{T_i})$ is implemented at stage i, $i = 1, \ldots, j$. Proposition 4 states that the inventory-related cost for echelon j can be decomposed into costs associated with each stage within the echelon.

Proposition 4 (Shang and Zhou [35]). *$G_j(\mathbf{T_j}) = \sum_{i=1}^{j} g_i(\mathbf{T_i})$, for $j = 1, \ldots, N$.*
Let $c_j(\mathbf{T_j}) \triangleq K_j/T_j + g_j(\mathbf{T_j})$, the average total cost per period for stage j, $j = 1, \ldots, N$. We have

$$C(\mathbf{T}) = \sum_{j=1}^{N} \frac{K_j}{T_j} + G_N(\mathbf{T_N}) = \sum_{j=1}^{N} \frac{K_j}{T_j} + \sum_{j=1}^{N} g_j(\mathbf{T_j}) = \sum_{j=1}^{N} c_j(\mathbf{T_j}).$$

This completes the decomposition of $C(\mathbf{T})$.

Bounds for the Stage Cost Function. This section derives lower and upper bounds for the stage cost $c_j(\mathbf{T_j})$. At the end, we shall see that these cost bounds are a function of stage j's control variable T_j and are independent of stage is, $i \neq j$.

Consider any given feasible solution $(\mathbf{T})$ with $\mathbf{T_j}$. For $j = 1, \ldots, N$, define the following:

$\mathbf{T_j^j}$ = a vector with j components whose values are T_j;
$\mathbf{T_j^1}$ = a vector with j components whose jth component is T_j and the other components are 1.

For example, $\mathbf{T_3^3} = (T_3, T_3, T_3)$ and $\mathbf{T_4^1} = (1, 1, 1, T_4)$. Recall the average induced-penalty cost function $P_j(y, \mathbf{T_j})$ in (26). We first state the main result of this section.

Proposition 5 (Shang and Zhou [35]). *$P_j(y, \mathbf{T_j^j}) \leq P_j(y, \mathbf{T_j}) \leq P_j(y, \mathbf{T_j^1})$, for all y and $j = 2, \ldots, N$.*

Proposition 5 states that for stage j, when its downstream stage i uses the same reorder interval length as stage j (i.e., $T_i = T_j$, $i < j$), the resulting average induced-penalty cost function $P_j(y, \mathbf{T_j^j})$ is a lower bound to $P_j(y, \mathbf{T_j})$ for all y. On the other hand, when a downstream stage i uses the smallest reorder interval length (i.e., $T_i = 1$, $i < j$), the resulting induced-penalty cost function $P_j(y, \mathbf{T_j^1})$ is an upper bound.

With Proposition 5, we can construct cost bounds to $g_j(y, \mathbf{T_j})$ for any y by first regulating downstream reorder intervals. More specifically, because

$$g_j(y, \mathbf{T_j^j}) \leq g_j(y, \mathbf{T_j}) \leq g_j(y, \mathbf{T_j^1}),$$

we have

$$g_j(s_j(\mathbf{T_j^j}), \mathbf{T_j^j}) \leq g_j(s_j(\mathbf{T_j}), \mathbf{T_j}) \leq g_j(s_j(\mathbf{T_j^1}), \mathbf{T_j^1}),$$

or equivalently, $g_j(\mathbf{T_j^j}) \leq g_j(\mathbf{T_j}) \leq g_j(\mathbf{T_j^1})$, which implies

$$c_j(\mathbf{T_j^j}) \leq c_j(\mathbf{T_j}) \leq c_j(\mathbf{T_j^1}). \tag{27}$$

This completes the construction of the lower- and upper-bound functions. Note that $c_j(\mathbf{T_j^j})$ and $c_j(\mathbf{T_j^1})$ are functions of stage j's control variable T_j. They are independent of downstream stages i's decision variables, $i < j$.

The Heuristic. The heuristic is generated by solving the following problem:

$$\text{(TP)} \quad \min_{\mathbf{T}} \sum_{j=1}^{N} c_j(\mathbf{T_j}), \quad \text{s.t. } T_{j+1} = n_j T_j, \quad n_j, T_j \in \mathbb{N}, \, j = 1, \ldots, N-1,$$

where $c_j(\mathbf{T_j}) = (K_j/T_j) + g_j(\mathbf{T_j})$.

Now consider the lower- and upper-bound functions $g_j(\mathbf{T_j^j})$ and $g_j(\mathbf{T_j^1})$. Define $\underline{g}_j(T_j) = g_j(\mathbf{T_j^j})$, $\bar{g}_j(T_j) = g_j(\mathbf{T_j^1})$, and $\underline{c}_j(T_j) = (K_j/T_j) + \underline{g}_j(T_j)$, $\bar{c}_j(T_j) = (K_j/T_j) + \bar{g}_j(T_j)$.

Proposition 6 (Shang and Zhou [35]). *$\bar{c}_j(T_j)$ is convex in T_j.*

We propose a simple heuristic for (TP) by solving the sum of the stage cost bound functions, subject to relaxed constraints. This heuristic yields a solution by conducting the same two steps, clustering and minimization, as discussed in §3.2. We use the upper-bound cost functions to illustrate the idea. We aim to solve the following problem:

$$\begin{aligned} \min_{\mathbf{T}} \quad & \sum_{j=1}^{N} \bar{c}_j(T_j) \\ \text{s.t.} \quad & T_{j+1} \geq T_j, \quad j = 1, \ldots, N-1. \end{aligned}$$

Clustering. Because the objective function is the sum of N separable, convex functions, a clustering algorithm (e.g., Maxwell and Muckstadt [28], pp. 1325–1334) can solve the relaxed problem efficiently. The output of the algorithm is an optimal partition that includes disjoint clusters, such as $\{c(1), c(2), \ldots, c(M)\}$, where M is the number of clusters in the optimal partition. (A cluster is a set of consecutive stages that use the same reorder interval.)[1]

Minimization. After the optimal partition is identified, we find a solution that satisfies the integer-ratio constraints for each cluster. More specifically, let $T_{c(1)} = \arg\min_T \sum_{i \in c(1)} \bar{c}_i(T)$. For $m = 2, \ldots, M$, we solve the following problem sequentially:

$$T_{c(m)} = \arg\min_{T} \sum_{i \in c(m)} \bar{c}_i(T), \quad \text{s.t. } T = nT_{c(m-1)}, \, n \in \mathbb{N}. \tag{28}$$

In other words, we restrict $T_{c(m)}$ to be an integer multiple of $T_{c(m-1)}$. Let $T'_j = T_{c(m)}$ for $j \in c(m)$, $m = 1, \ldots, M$. We then obtain one feasible reorder interval solution $(T'_1, \ldots, T'_N)$ for (TP). With these reorder intervals, we can find the best reorder points through the procedure in §4.1. Let the resulting total cost be C'.

Similarly, we can apply the same procedure to minimize $\sum_{j=1}^{N} \underline{c}_j(T_j)$. However, because $\underline{c}_j(T_j)$ may not be a convex function, we cannot apply the clustering algorithm directly. Thus, we use the same partition found in the upper bound problem and then find the reorder intervals in the same fashion as in (28) except by replacing $\bar{c}_j(T_j)$ with $\underline{c}_j(T_j)$. In this case, we use the first local minimizer as the reorder interval solution for a cluster. Define the resulting feasible reorder intervals as $(T''_1, \ldots, T''_N)$ and the resulting optimal cost as C''.

The heuristic solution for (TP) is either T'_j or T''_j, $j = 1, \ldots, N$, whichever yields a smaller total cost.

Heuristic 2. Shang and Zhou [37] suggested another simpler heuristic that seems to work well. This heuristic is similar to those suggested for the (r, Q) policy in §3.2. The key idea of the heuristic is that we approximate the exact stage cost function $g_j(\mathbf{T_j})$ by a cost function $g_j^h(T_j)$ generated from a single-stage (s, T) system with the original problem data. Below we demonstrate how to construct this single-stage cost function.

[1] One might be able to follow Shang's [31] analysis to show that the cost ratio K_j/h_j can be used for determining an effective partition. It is conceivable that such a cost-ratio-based partition should work well for the (s, T) model.

Consider the information perceived by the echelon j manager under the scenario where centralized information is available. The echelon j manager knows the echelon holding cost h_j and the real-time customer demand information. Also, because each stage manager shares a common goal of minimizing the supply chain cost, a stage manager is responsible for fulfilling the customer demand. Thus, when making the inventory decision, manager j should consider his *effective* lead time $L_{[1,j]}$, i.e., the duration between when an order is placed by stage j and when the order arrives at stage 1, assuming that stage $j+1$ has ample supply. Following the same logic, manager j should be charged a *perceived* backorder cost rate $b+h_{[j+1,N]}$ if he cannot fulfill the customer demand. (We refer the reader to Shang and Song [33] for an explanation of how the perceived backorder cost rate is derived.) With these perceived parameters, manager j would expect an average inventory holding and backorder cost per period of

$$g_j^h(s_j,T_j) = \frac{1}{T_j}\sum_{\tau=0}^{T_j-1} \mathsf{E}\big[h_j\big(s_j - D[L_{[1,j]}+\tau]\big) + \big(b+h_{[j,N]}\big)\big(y - D[L_{[1,j]}+\tau]\big)^-\big].$$

Note that this is exactly the lower-bound cost function $G_j^\ell(\cdot)$ introduced in §4.1. With fixed T_j, $g_j^h(s_j,T_j)$ is convex in s_j, and manager j will choose the base-stock level as $s_j(T_j) = \arg\min_{s_j} g_j^h(s_j,T_j)$. Let $g_j^h(T_j) = g_j^h(s_j(T_j),T_j)$, and the average total perceived cost per period for stage j is $C_j^h(T_j) = K_j/T_j + g_j^h(T_j)$. Our heuristic is to approximate $g_j(\mathbf{T_j})$ by $g_j^h(T_j)$, which results in the following approximate problem:

$$\begin{aligned} &\min_{\mathbf{T}} \; \sum_{j=1}^{N} C_j^h(T_j) \\ &\text{s.t.} \;\; T_{j+1} = n_j T_j, \quad j = 1,\ldots,N-1. \end{aligned}$$

We can relax the integer-ratio constraints into inequality constraints and apply the same clustering-minimization procedure in Heuristic 1 to generate heuristic reorder intervals. Shang and Zhou [37] showed that on average the heuristic works slightly better than Heuristic 1 in a numerical study.

5. Extensions

We consider a few extensions for the basic serial model.

5.1. Local Information

So far, we assumed that the demand information is centralized. However, in practice, a supply chain may have different levels of information integration. We now consider a serial system where each location only knows its local demand information. The local demand for a stage is the order placed by its immediate downstream stage. Below is a summary of the main findings.

5.1.1. Local (r,Q) Policies. Axsäter and Rosling [5] showed that the local (r,Q) policy is a special case of the echelon (r,Q) policy (and that the local policy is therefore suboptimal). Chen [12] provided an algorithm to search for the optimal local reorder points when batch sizes are fixed. It is, however, not clear how to obtain the optimal local batch sizes. Shang et al. [39] provided an approach to obtain the optimal batch sizes for the local (r,Q) policy.

5.1.2. Local (s,T) Policies. Shang et al. [39] characterized the dynamics of key inventory variables under the local (s,T) policy. These dynamics lead to a simple, bottom-up recursion that can evaluate a given local (s,T) policy. More specifically, by converting the local inventory variables into echelon terms, the local policy can be evaluated as if it were an echelon (s,T) policy with modified system lead times. This evaluation procedure can be used further to find the optimal local base-stock levels for given reorder intervals. Unlike the (r,Q) policy, the local (s,T) policy is not a special case of the echelon (s,T) policy. Nevertheless, we show that the optimal echelon (s,T) policy always dominates the optimal local one.[2] With this result, Shang et al. [39] provided a method to search for the optimal reorder intervals for the local (s,T) policy.

These analytical results motivated Shang et al. [39] to investigate how to improve a supply chain operated by local (s,T) policies. It is arguably true that most supply chains belong to this type. More specifically, we assume that the supply chain partners (who possibly belong to a single firm) share a common goal of optimizing systemwide performance, but they only have local information at the operational planning level. There is a central planner who determines all inventory control parameters (e.g., based on summary data in a spreadsheet), and each stage simply generates replenishment orders based on these control parameters and the local demand and inventory information. Shang et al. [39] considered two improvement strategies—expanding the information flow by acquiring real-time demand information or accelerating the material flow by providing flexible deliveries. Notice that the first improvement strategy corresponds to the echelon (s,T) policy, and the second corresponds to the continuous-review local (r,Q) policy. The question is, which strategy should be chosen and under what conditions?

In a numerical study, Shang et al. [39] found that the average (maximum) cost reduction when the system switches from the local periodic-review (s,T) policy to the local continuous-review (r,Q) policy is 11.27% (28.32%), which is significantly larger than the 5.51% (16.67%) achieved when switching from the local (s,T) policy to the echelon (s,T) policy. This suggests that increasing the flexibility of delivery is more beneficial than acquiring real-time demand information. In addition, the difference between the benefits of both improvement strategies increases when demand becomes more variable. This result provides an important insight to avoid a management pitfall: when demand becomes more variable, it may not be effective to invest in IT systems to acquire real-time demand information; an agile logistics system is key to achieving supply chain efficiency.

Shang et al. [39] further used the model to examine the issue of the *value of demand information* (*VOI*). The VOI is equivalent to the cost reduction by switching from the local policy to the echelon policy. In a numerical study, they found that there is no significant benefit achieved by switching from the local (r,Q) policy to the echelon (r,Q) policy. This observation suggests that adding demand information to a system with flexible deliveries provides little value. This also suggests that, in general, the optimal batch sizes obtained from the echelon (r,Q) policy are effective approximations for those of the local (r,Q) policy. For the (s,T) policies, the VOI is higher when the fixed costs are larger and lead times are shorter. This implies that the demand information is most beneficial when the supply chain has long reorder intervals (higher fixed costs lead to longer reorder intervals) and short lead times. In other words, when the fixed costs are small, one can use the optimal reorder intervals obtained from the echelon (s,T) policy as approximations for those of the local (s,T) policy.

5.2. Distribution Systems

The single-stage-based heuristic can be extended to one-warehouse, multiretailer distribution systems. More specifically, we consider a periodic-review, two-echelon distribution system in which a single warehouse supplies N nonidentical retailers. Time periods are indexed

[2] An echelon policy may not always dominate a local one. For example, Axsäter and Juntti [4] showed that a local policy may dominate an echelon policy under some conditions for a two-echelon distribution system.

$0, 1, 2, \ldots$. We use j as the stage index, where 0 represents the warehouse, and j represents the retailer j, $j = 1, \ldots, N$. Retailer j faces Poisson demand with stationary rate λ_j. The retailers' demands are independent of each other. Let $D_j[t, t+\tau)$ and $D_j[t, t+\tau]$ denote the cumulative demand over time in periods $[t, t+\tau)$ and $[t, t+\tau]$ for stage j, respectively. There is a constant lead time L_j for stage j, $j = 0, \ldots, N$. Let $L_{[0,j]} = L_0 + L_j$. Each stage implements a stationary echelon (s, T) policy, where s_j is the echelon base-stock level, and T_j is the reorder interval for stage j. The policy is operated as follows: At the beginning of every T_j periods, stage j orders up to s_j if the echelon inventory order position is less than s_j. For the retailer, the echelon inventory order position is equal to its local inventory order position, i.e., inventory on order plus inventory on hand minus backorders. For the warehouse, the echelon inventory order position is the inventory on order plus inventory on hand plus inventory in transit to and held at the retailers minus backorders at the retailers. We term these T_jth periods as *order periods* and the moment of placing an order as an *order epoch*. Let h_j be the echelon holding cost rate for stage j, $j = 0, \ldots, N$, and the local holding cost rate $H_j = h_0 + h_j$ for $j = 1, \ldots, N$. Unmet demand is fully backlogged at all stages. Let b_j be the backorder cost rate for stage j, $j = 1, \ldots, N$. Finally, there is a fixed cost K_j associated with each inventory reorder. The objective is to obtain the policy such that the total system cost per period is minimized.

We assume that the reorder intervals are an integer multiple of some base period. Thus, T_0 may be smaller than T_j, $j = 1, \ldots, N$. Without loss of generality, let the base period be one. The ordering activities between the warehouse and the retailers are coordinated in a synchronized manner: whenever possible, retailer j places an order when the warehouse receives a shipment. For example, consider a one-warehouse, two-retailer system with $L_0 = 3$, $T_0 = 2$, $T_1 = 1$, and $T_2 = 3$. Suppose that the warehouse places an order at the beginning of period t, $t+2$, $t+4, \ldots$. Consider the order epoch t. This order placed at t will arrive at the beginning of period $t + L_0 = t + 3$. This is the moment that both retailers place an order. Thus, the order periods for retailer 1 are $t+3$, $t+4$, $t+5, \ldots$, and for retailer 2 they are $t+3$, $t+6$, $t+9, \ldots$. We term such an order coordination scheme as a *synchronized ordering* rule. Let

$$T = lcm\{T_0, T_1, \ldots, T_N\},$$

where *lcm* is an operator that generates the least common multiplier of T_js. If we define t and $t + L_0$ as the starting time of a *cycle*; clearly, the next cycle will start T periods later under the synchronized ordering rule (that is, the next time when the warehouse receives a shipment and all retailers place an order occurs T periods later).

Because of the centralized information available, we assume that the warehouse commits its inventory, if available, to fill the incoming orders according to the sequence of the demands occurring at the retailer site. If the warehouse does not have an uncommitted unit to fill an arrived demand, the warehouse creates a backorder and adds this to the current list of outstanding orders. When inventory becomes available, the outstanding orders are filled in the sequence in which they are created. This is the so-called *virtual allocation* rule (Graves [25]).

Shang and Zhou [36] provided a bottom-up recursion to evaluate the total cost per period $C(\mathbf{s}, \mathbf{T})$. This procedure is to first evaluate retailer j's cost by assuming that the retailer has ample supply. Then, they evaluated the echelon cost of the warehouse, which is equivalent to the total system cost. More specifically, let $\mathbb{M}_x(y)$ be an operator that returns the remainder of y divided by x, where x is a positive integer and y is a nonnegative integer.

Let $D_j[\tau)$ and $D_j[\tau]$ denote the total demand in τ and $\tau + 1$ periods, respectively, for retailer j. For $j = 1, \ldots, N$, define

$$G_j(y, T_j, r) = \mathsf{E}\big[h_j\big(y - D_j[L_j + \mathbb{M}_{T_j}(r)]\big) + (b_j + H_j)\big(y - D_j[L_j + \mathbb{M}_{T_j}(r)]\big)^-\big], \qquad (29)$$

and

$$G_0(\mathbf{s}, \mathbf{T}, r) = \mathsf{E}\left[h_0\big(s_0 - D_0[L_0 + \mathbb{M}_{T_0}(r)]\big) + \sum_{j=1}^{N} G_j\big(s_j - B_{0j}(L_0 + r_j(r)), T_j, r\big)\right], \quad (30)$$

where $B_{0j}(L_0 + r_j(r))$ is the steady-state distribution for $B_{0j}(t + L_0 + r_j(r))$, which is the number of backorders at the warehouse because of retailer j's order. We refer to Shang and Zhou [36] for the derivation of this variable. Let

$$G(\mathbf{s}, \mathbf{T}) \triangleq \frac{1}{T} \sum_{r=0}^{T-1} G_0(\mathbf{s}, \mathbf{T}, r).$$

Proposition 7 (Shang and Zhou [36]). *For given echelon base-stock policies with parameters* $(\mathbf{s}, \mathbf{T})$*, the average total cost per period is*

$$C(\mathbf{s}, \mathbf{T}) = \sum_{j=0}^{N} \frac{K_j}{T_j} + G(\mathbf{s}, \mathbf{T}),$$

where $\sum_{j=0}^{N}(K_j/T_j)$ *is the average total fixed order cost per period, and* $G(\mathbf{s}, \mathbf{T})$ *is the average inventory holding and backorder cost per period.*

Shang and Zhou [36] suggested a complete enumeration to find the optimal reorder intervals. To facilitate the search, they constructed bounds for the optimal reorder interval T_j^*, $j = 0, \ldots, N$. These solution bounds can be obtained by constructing a lower-bound function for each retailer's cost and for the entire system cost.

They further proposed a heuristic based on solving a set of single-stage cost functions. The heuristic is composed of two steps. The first step is clustering, i.e., clustering all stages into three groups $G(T_0)$, $E(T_0)$, and $L(T_0)$, where $G(T_0)$ contains the stages whose reorder intervals are larger than T_0, $E(T_0)$ contains the stages whose reorder intervals are equal to T_0, and $L(T_0)$ contains the stages whose reorder intervals are less than T_0. This step is essentially the same as that in the deterministic model. We refer the reader to Roundy [30, pp. 1419–1420] for a detailed analysis on finding these three clusters.

The key difference between Shang and Zhou's [36] heuristic and Roundy's [30] is in the second step, where an integer-ratio policy is generated. After the above clustering step, each stage belongs to one of these three clusters. They then use the lower-bound cost function $\underline{c}_j(T_j)$ to generate an initial solution $(\hat{T}_0, \hat{T}_1, \ldots, \hat{T}_N)$. More specifically, for stage $j \in E(T_0)$, let $c^E(T) = \sum_{j \in E(T_0)} \underline{c}_j(T)$, and $\hat{T}_j = \arg\min_T c^E(T)$. For stage $j \notin E(T_0)$, $\hat{T}_j = \arg\min_{T_j} \underline{c}_j(T_j)$.[3]

The initial solution $(\hat{T}_0, \ldots, \hat{T}_N)$ will be used to generate the heuristic solution. The approach is as follows: First, find the smallest reorder interval among $\hat{T}_j$, i.e., $T^o = \min\{\hat{T}_0, \hat{T}_1, \ldots, \hat{T}_N\}$. Then,

$$\tilde{T}_j = \begin{cases} \displaystyle \mathop{\arg\min}_{T \in \{n_j T^o,\, (n_j+1) T^o\}} \sum_{j \in E(T_0)} \underline{c}_j(T), & j \in E(T_0), \\ \displaystyle \mathop{\arg\min}_{T \in \{n_j T^o,\, (n_j+1) T^o\}} \underline{c}_j(T), & j \notin E(T_0), \end{cases}$$

where $n_j \in \mathbb{N}$. The integer-ratio solution $(\tilde{T}_0, \tilde{T}_1, \ldots, \tilde{T}_N)$ will be our heuristic solution.

A common approach for obtaining reorder intervals is through a power-of-two solution generated from the corresponding deterministic model. Shang and Zhou [36] found that the deterministic power-of-two solution in general performs well. However, the power-of-two solution can perform poorly under certain conditions.

[3] Notice that $\underline{c}_j(\cdot)$ may not be convex. Thus, it is possible that we are unable to find the minimizer when searching for $\hat{T}_j$. In our numerical study, the cost functions in question are unimodal, and we therefore set $\hat{T}_j$ to be the first minimizer in the code.

5.3. Nonstationary Demand Model

Customer demand is often nonstationary in practice. Causes of nonstationary demand include product life cycles, seasonality, trends, and economic conditions. Nonstationary demand makes it difficult for managers to determine optimal stocking levels in a supply chain because finding the optimal systemwide solution often requires solving several interrelated, recursive cost functions between stages across time. The other difficulty concerns implementation: if a supply chain is composed of independent firms, these firms may not be willing to implement the optimal solution without appropriate incentives. In a nonstationary demand environment, the optimal solution is often time varying. Thus, designing an incentive scheme that can induce each stage to choose the optimal stocking level in each time period poses a difficult challenge.

Shang [32] provided a heuristic that can simplify the computation and help facilitate the implementation of the optimal solution. Consider the same N-stage serial inventory system. Define p_j as the unit order cost for stage j. The other system parameters are the same as defined before. We use t to index the time period and count the time backward: that is, if t is the current period, $t-1$ will be the next period, etc. Let T be the planning horizon length. Denote $D(t)$ as the demand in period t. The demands are independent between periods, but the demand distributions may differ from period to period. Let $D[t,s] = \sum_{i=s}^{t} D(i)$ represent the total demand in period $t, t-1, t-2, \ldots, s$, where $t \geq s$. The sequence of events in a period is as follows: (1) each stage receives a shipment sent one period ago from its upstream stage at the beginning of the period; (2) each stage places an order at the beginning of the period; (3) each stage sends a shipment to its downstream stage; (4) demand occurs at stage 1 during the period; and (5) inventory cost is evaluated at the end of the period.

Clark and Scarf [17] showed that time-varying, echelon base-stock policies are optimal. Let $s_j(t)$ be the optimal echelon base-stock level for stage j. The echelon base-stock policy is executed as follows: stage j reviews x_j at the beginning of each period, where x_j is the echelon inventory position for stage j (equal to inventory on order plus inventory at stage j plus inventory in transit to and at stage $i(<j)$ minus backorders). The stage orders up to $s_j(t)$ if $x_j < s_j(t)$ and does not order otherwise. It is well known that $s_j(t)$ depends on the downstream base-stock levels and can be found by solving j sets of dynamic programs sequentially. More specifically, finding $s_1(t)$ is equivalent to solving a single-stage system. With the known $s_1(t)$, one can form a dynamic program to compute the functional equation for stage 2, assuming that stage 2 has ample supply. The optimal base-stock level $s_2(t)$ is the optimal solution obtained from the functional equation. Continuing this procedure, with the known $s_i(t)$, $i < j$, one can compute the functional equation for stage j and the corresponding optimal solution $s_j(t)$. In other words, when finding the optimal solution for stage j, we can only focus on *echelon* j that includes stage 1 to stage j by viewing the echelon as having ample supply from its upstream stage. Clearly, the complexity will grow quickly when the chain and the planning horizon become longer because the optimal value function of an upstream stage depends on all of its downstream solutions across time.

In the literature, there is a stream of research aiming to simplify the computation, e.g., Federgruen and Zipkin [22], De Bodt and Graves [18], Shang and Song [33], Dong and Lee [19], Gallego and Özer [24], and Chao and Zhou [10]. These papers assume an independent and identically distributed demand stream, and therefore the induced penalty cost function (the cost charged to an upstream stage if the stage cannot fulfill the order from its immediate downstream stage) becomes stationary, and the analysis becomes a single-cycle problem. Although this single-period approach simplifies the analysis, it cannot reflect how the base-stock level changes based on the nonstationary demand over time.

Shang's [32] heuristic solution is time varying. It can generate an approximation for stage j $(1 \leq j \leq N)$ without knowing stage i's base-stock level, $i < j$. More specifically, he showed that the optimal base-stock level for stage j is bounded by the optimal solutions of two single-stage systems with the original problem data. This result is established in three steps.

First, the optimal value function for stage j is bounded above and below by that of a revised j-stage system. These revised systems are referred to as the upper-bound system and the lower-bound system, respectively. The upper-bound system is constructed by requiring stage i ($<j$) to always order up to stage $i+1$'s echelon inventory level in each period. On the other hand, the lower-bound system is constructed by regulating stage i's holding and order cost parameters. Second, Shang [32] showed that the optimal base-stock level for stage j in the upper-bound (lower-bound) system is a lower (upper) bound to that in the original system. Finally, Shang [32] proved that solving these revised j-stage systems is equivalent to solving a single-stage system with parameters obtained from the original problem data. A numerical study suggests that the gap between these two solution bounds is generally quite small. This result suggests a heuristic solution for each stage by solving a single-stage system with a weighted average of the cost parameters obtained from the upper- and the lower-bound systems. In a numerical study, Shang [32] found that the heuristic generates 58% of the optimal solutions, and that 98% of the heuristic solutions are within the ± 1 unit range of the optimal solutions.

Below we characterize the solution bounds as well as the single-stage systems that yield these bounds. We first introduce two single-stage systems. Let p_j^u, h_j^u, b_j, and $\mathfrak{T}_j$ represent the purchase cost, holding cost, backorder cost, and lead time, respectively, for the upper-bound system; a similar definition applies to the lower-bound system:

$$\text{upper-bound system: } \mathbb{S}_j^u[p_j^u, h_j^u, b_j, \mathfrak{T}_j],$$
$$\text{lower-bound system: } \mathbb{S}_j^\ell[p_j^\ell, h_j^\ell, b_j, \mathfrak{T}_j],$$

where

$$p_j^u = p_j + \sum_{i=2}^{j} \alpha^{L_{[i,\,j]}}(p_{i-1} + h_i),$$
$$h_j^u = h_{[1,\,j]},$$
$$p_j^\ell = p_j + \alpha^{L_{[j,\,j]}} h_j,$$
$$h_j^\ell = h_j,$$
$$b_j = b + h_{[j+1,\,N]},$$
$$\mathfrak{T}_j = L_{[1,\,j]}.$$

Theorem 2 (Shang [32]). (1) *If the initial echelon in-transit position for stage j-1 is less than* $s_{j-1}(t)$, $s_j(t) \geq s_j^\ell(t)$ *for* $t > L_{[1,\,j]}$, *where* $s_j^\ell(t)$ *is the solution obtained from* $\mathbb{S}_j^u[p_j^u, h_j^u, b_j, \mathfrak{T}_j]$;

(2) $s_j^u(t) \geq s_j(t)$, *for* $t > L[1, j]$, *where* $s_j^u(t)$ *is the solution obtained from* $\mathbb{S}_j^\ell[p_j^\ell, h_j^\ell, b_j, \mathfrak{T}_j]$.

5.3.1. Heuristic and Myopic Solution. Let $p_j^a = wp_j^u + (1-w)p_j^\ell$, and let $h_j^a = wh_j^u + (1-w)h_j^\ell$, where $0 \leq w \leq 1$. We call the resulting single-stage system *heuristic system j*, denoted by $\mathbb{S}_j^a(p_j^a, h_j^a, b_j, \mathfrak{T}_j)$, and we define the resulting optimal solution as $s_j^a(t)$. Shang [32] showed that $s_j^a(t)$ is an effective approximation to the optimal solution $s_j(t)$. Furthermore, it is well known that the myopic solution is an upper bound to the optimal solution of a single-stage system (Zipkin [42], pp. 378–379). These two results together motivate us to derive a closed-form approximation, referred to as $s_j^m(t)$, for $s_j^a(t)$. In our numerical study, we find that in all cases $s_j^m(t)$ moves in the same direction as the optimal solution $s_j(t)$. Thus, we can use the myopic solution to derive an approximation for the optimal local base-stock level to examine the system behavior. Below we lay out the detailed steps.

Let the myopic solution for $\mathbb{S}_j^a(p_j^a, h_j^a, b_j, \mathfrak{T}_j)$ be $s_j^m(t)$.

Proposition 8 (Shang [32]). *For* $t > \mathfrak{T}_j$,

$$s_j^m(t) = \arg\min_{s_j} \{\mathsf{P}(D[t, t-\mathfrak{T}_j] \le s_j) > \beta_j\},$$

where

$$\beta_j = \frac{\alpha^{\mathfrak{T}_j} b_j - p_j^a(1-\alpha)}{\alpha^{\mathfrak{T}_j}(b_j + h_j^a)}.$$

Note that when $t < \mathfrak{T}_j$, stage j will not order. When $t = \mathfrak{T}_j + 1$, $s_j^m(t)$ is equal to the solution obtained from the above equation except for the removal of the term $(1-\alpha)$ in the numerator because of the termination value being equal to zero.

To obtain a closed-form expression, Shang [32] applied a normal approximation on $D[t, t-\mathfrak{T}_j]$. Let the mean of $D[t, t-\mathfrak{T}_j]$ be $\lambda[t, t-\mathfrak{T}_j]$, and let the standard deviation be $\sigma[t, t-\mathfrak{T}_j] = \sqrt{\mathsf{Var}[D[t, t-\mathfrak{T}_j]]}$. We can form a closed-form expression for $s_j^m(t)$:

$$s_j^m(t) = \lambda\big[t, t-\mathfrak{T}_j\big] + \sigma\big[t, t-\mathfrak{T}_j\big]\Phi^{-1}(\beta_j).$$

Thus, the local base-stock level is $s_1^{m'}(t) = s_1^m(t)$, and for $j = 2, \ldots, N$,

$$\begin{aligned} s_j^{m'}(t) &= s_j^m(t) - s_{j-1}^m(t) \\ &= \lambda\big[t-\mathfrak{T}_{j-1}-1, t-\mathfrak{T}_j\big] + \sigma\big[t, t-\mathfrak{T}_j\big]\Phi^{-1}(\beta_j) - \sigma\big[t, t-\mathfrak{T}_{j-1}\big]\Phi^{-1}(\beta_{j-1}). \quad (31) \end{aligned}$$

The first term in Equation (31) is the average pipeline inventory in period t, which depends on the average τ_j periods of *future* demand in period $t-\mathfrak{T}_{j-1}-1, t-\mathfrak{T}_{j-1}-2, \ldots, t-\mathfrak{T}_j$. The second term is the safety stock for stage j in period t, denoted as $ss_j^m(t)$, which depends on the cost ratios β_j and β_{j-1} and the variability of the demand in period $[t, t-\mathfrak{T}_j]$.

Equation (31) allows us to analytically investigate how the system parameters affect the optimal base-stock level and the safety stock at each stage. For example, if we are interested in the change to the amount of safety stock of the upstream stage in a two-stage system, we can define the change to stage 2's safety stock in period t as

$$\begin{aligned} \Delta ss_2^m(t) &= ss_2^m(t-1) - ss_2^m(t) \\ &= \big(\sigma[t-1, t-1-\mathfrak{T}_2] - \sigma[t, t-\mathfrak{T}_2]\big)\Phi^{-1}(\beta_2) \\ &\quad - \big(\sigma[t-1, t-1-\mathfrak{T}_1] - \sigma[t, t-\mathfrak{T}_1]\big)\Phi^{-1}(\beta_1). \quad (32) \end{aligned}$$

From the above equation, we can see that $\Delta ss_2^m(t)$ will be fairly small unless there is a significant difference between $\mathsf{Var}[D(t-1-\mathfrak{T}_2)]$ and $\mathsf{Var}[D(t-1-\mathfrak{T}_1)]$. This implies that the safety stock at the upstream stage should be fairly stable. (A similar conclusion is observed by Graves and Willems [26] in their numerical study.) In addition, $\Delta ss_2^m(t)$ may not be positive even if $\mathsf{Var}[D(t)] < \mathsf{Var}[D(t-1)]$, $\forall t$. More specifically, when either p_2 is large or h_2 is large, or when b is small, $\Phi^{-1}(\beta_2)$ tends to be smaller than $\Phi^{-1}(\beta_1)$, causing the difference in (32) to become negative even when the demand variance increases over time. This suggests that the safety stock at an upstream stage may not increase with the demand variability.

Example 1. We consider a two-stage system with $\tau_1 = \tau_2 = 1$, $p_2 = 6$, $h_2 = 1$, $p_1 = 4$, $h_1 = 1$, and $b = 15$. The demand follows a Poisson distribution with mean rate shown in Table 1. We report the optimal echelon, heuristic, and myopic base-stock levels in each period. We also report the optimal local base-stock level as well as the corresponding safety stock for stage 2. As can be seen, the safety stock may decrease although the demand rate increases (e.g., from $t = 10$ to $t = 9$).

It is conceivable that Clark and Scarf's [17] nonstationary demand model can be generalized into a batch-ordering model, where the batch sizes are fixed and satisfy the integer-ratio

TABLE 1. A two-stage example with the optimal, heuristic, and myopic solutions, as well as the local base-stock levels and the resulting safety stocks.

Period (t)	10	9	8	7	6	5	4	3	2	1
Demand rate	2	4	6	8	10	9	7	5	3	1
$s_1(t)$	10	15	20	24	26	22	16	10	5	—
$s_2(t)$	16	22	29	33	31	24	16	6	—	—
$s_2^a(t)$	16	23	29	33	31	24	16	7	—	—
$s_1^m(t)$	10	15	20	24	26	23	18	12	7	—
$s_2^m(t)$	16	23	29	33	32	26	19	7	—	—
$s_2'(t) = s_2(t) - s_1(t)$	6	7	9	9	5	2	0	−4	—	—
$ss_2(t) = s_2'(t) - \mathsf{E}[D[t-2, t-2]]$	0	−1	−1	0	−2	−3	−3	−5	—	—

constraints, i.e., $Q_j = n_j Q_{j-1}$, $n_j \in \mathbb{N}$. Thus, the above single-stage bounds can be carried over to the batch-ordering model, and a simple single-stage based heuristic can be obtained for the (time-varying) optimal reorder points.

Finally, all of the results can be applied to forecast demand models and Markov-modulated demand models.

6. Summary and Future Research

This chapter summarizes recent development on single-stage-based heuristics for multiechelon inventory models. Two commonly implemented inventory control policies are considered: continuous-review (r, Q) policies and periodic-review (s, T) policies. For both policies, we demonstrate that the single-stage-based heuristics share a very similar procedure to yield effective solutions. These single-stage heuristics not only simplify the computation but also help implementation.

We lay out three important future research directions. First, the current single-stage-based heuristics have been successfully applied to serial systems and distribution systems with (s, T) policies. However, a supply chain is an intricate network with complicated system structures. For example, some parts of a supply chain may be serial types of networks and some distribution types. In the deterministic demand network model, the power-of-two solution has been proven to guarantee a worst-case cost bound. It will be important to develop a simple heuristic for general supply chain networks with stochastic demand.

Second, multiechelon models lay the foundation for studying supply chain problems. However, there are relevant business activities related to operations that are not specifically incorporated into the models. For example, a supply chain includes material, information, and financial flows, and these flows are often entangled with each other. The multiechelon literature has focused on the first two flows. Nevertheless, the recent financial crisis demonstrates a need for jointly studying these three flows together. It is conceivable that such a joint problem will be extremely difficult to analyze. Thus, one would hope that the single-stage-based heuristics can help simplify the analysis for these joint problems.

Finally, joint inventory and dynamic pricing models in the literature are for single-stage systems, e.g., Federgruen and Heching [20]. The impact of pricing policy on the inventory decision has been well investigated for the single-stage system. It will be interesting to study the impact of the pricing policy on the inventory decisions in the supply chain. The idea of constructing these single-stage heuristics reviewed in this chapter may be useful for constructing simple joint inventory and pricing policies for a multiechelon system.

References

[1] S. Axsäter. Exact and approximate evaluation of batch-ordering policies for two-level inventory systems. *Operations Research* 41:777–785, 1993.
[2] S. Axsäter. Exact analysis of continuous review (r, q) policies in two-echelon inventory systems with compound poisson demand. *Operations Research* 48:686–696, 2000.

[3] S. Axsäter. Supply chain operations: Serial and distribution inventory systems. A. G. de Kok, S. C. Graves, eds. *Supply Chain Management: Design, Coordination, and Operation*. Handbooks in Operations Research and Management Science, Vol. 11. Elsevier, Amsterdam, 525–559, 2003.

[4] S. Axsäter and L. Juntti. Comparison of echelon stock and installation stock policies for two-level inventory systems. *International Journal of Production Economics* 45:303–310, 1996.

[5] S. Axsäter and K. Rosling. Installation vs. echelon stock policies for multi-level inventory control. *Operations Research* 39:1274–1280, 1993.

[6] R. Badinelli. A model for continuous-review pull policies in serial inventory systems. *Operations Research* 40:142–156, 1992.

[7] G. Cachon. Managing supply chain demand variability with scheduled ordering policies. *Management Science* 95:843–856, 1999.

[8] G. Cachon. Exact evaluation of batch-ordering policies in two-echelon supply chains with periodic review. *Operations Research* 49:79–98, 2001.

[9] G. Cachon and P. Zipkin. Competitive and cooperative inventory policies in a two stage supply chain. *Management Science* 45:936–953, 1999.

[10] X. Chao and S. Zhou. Probabilistic solution and bounds for serial inventory system with discounted and average cost. *Naval Research Logistics* 54:623–631, 2007.

[11] X. Chao and S. X. Zhou. Optimal policy for a multiechelon inventory system with batch ordering and fixed replenishment intervals. *Operations Research* 57:377–390, 2009.

[12] F. Chen. Echelon reorder points, installation reorder points, and the value of centralized demand information. *Management Science* 44:S221–S234, 1998.

[13] F. Chen. On (R, nQ) policies in serial inventory systems. S. Tayur, R. Ganeshar, M. Magazine, eds. *Quantitative Models for Supply Chain Management*. Kluwer Academic Publishers, Norwell, MA, 71–109, 1998.

[14] F. Chen. Optimal policies for multi-echelon inventory problems with batch ordering. *Operations Research* 48:376–389, 2000.

[15] F. Chen and Y.-S. Zheng. Lower bounds for multi-echelon stochastic inventory systems. *Management Science* 40:1426–1443, 1994.

[16] F. Chen and Y.-S. Zheng. Near-optimal echelon-stock (r, nq) policies in multi-stage serial systems. *Operations Research* 46:592–602, 1998.

[17] A. Clark and H. Scarf. Optimal policies for a multi-echelon inventory problem. *Management Science* 6:475–490, 1960.

[18] M. De Bodt and S. Graves. Continuous review policies for a multi-echelon inventory problem with Stochastic demand. *Management Science* 31:1286–1295, 1985.

[19] L. Dong and H. Lee. Optimal policies and approximations for a serial multi-echelon inventory system with time-correlated demand. *Operations Research* 51:969–980, 2003.

[20] A. Federgruen and A. Heching. Combined pricing and inventory control under uncertainty. *Operations Research* 47:454–475, 1999.

[21] A. Federgruen and Y. Zheng. Optimal power-of-two replenishment strategies in capacitated general production/distribution networks. *Management Science* 39:710–727, 1993.

[22] A. Federgruen and P. Zipkin. Computational issues in an infinite horizon multi-echelon inventory problem with stochastic demand. *Operations Research* 32:818–836, 1984.

[23] K. Feng and U. S. Rao. Echelon-stock (R, nT) control in two-stage serial stochastic inventory systems. *Operations Research Letters* 35:95–104, 2007.

[24] G. Gallego and Ö. Özer. A new algorithm and a new heuristic for serial supply systems. *Operations Research Letters* 33:349–362, 2005.

[25] S. Graves. A multiechelon inventory model with fixed replenishment intervals. *Management Science* 42:1–18, 1996.

[26] S. Graves and S. Willems. Strategic inventory placement in supply chains: Nonstationary demand. *Manufacturing & Service Operations Management* 10:278–287, 2008.

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

[28] W. Maxwell and J. Muckstadt. Establishing consistent and realistic reorder intervals in production-distribution systems. *Operations Research* 33:1316–1341, 1985.

[29] E. Naddor. Optimal and heuristic decisions in single- and multi-item inventory systems. *Management Science* 21:1234–1249, 1975.

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

[31] K. Shang. Note: A simple heuristic for serial inventory systems with fixed order costs. *Operations Research* 56:1039–1043, 2008.

[32] K. Shang. Single-stage approximations for optimal policies in serial inventory systems with non-stationary demand. Working paper, Fuqua School of Business, Duke University, Durham, NC, 2010.

[33] K. Shang and J.-S. Song. Newsvendor bounds and heuristic for optimal policies in serial supply chains. *Management Science* 49:618–638, 2003.

[34] K. Shang and J.-S. Song. Supply chains with economies of scale: Bounds and approximations. *Operations Research* 55:843–853, 2007.

[35] K. Shang and S. Zhou. Optimal and heuristic echelon (r, nQ, T) policies in serial inventory systems with fixed costs. *Operations Research* 58:414–427, 2010.

[36] K. Shang and S. Zhou. Optimizing reorder intervals for two-echelon distribution systems with stochastic demand. Working paper, Fuqua School of Business, Duke University, Durham, NC, 2010.

[37] K. Shang and S. Zhou. A simle heuristic for echelon (r,nQ,T) policies in serial supply chain. *Operations Research Letters* 37:433–437.

[38] K. Shang, J.-S. Song, and P. Zipkin. Coordination mechanisms in decentralized serial inventory systems with batch ordering. *Management Science* 55:685–695, 2009.

[39] K. Shang, S. Zhou, and G.-J. van Houtum. Improving supply chain performance: Real-time demand information and flexible deliveries. *Manufacturing & Service Operations Management* 12:430–448, 2010.

[40] D. Simchi-Levi and Y. Zhao. Three generic methods for evaluating stochastic multi-echelon inventory systems. Working paper, Massachusetts Institute of Technology, Cambridge, 2007.

[41] G. J. van Houtum, A. Scheller-Wolf, and J. Yi. Optimal control of serial, multi-echelon inventory/production systems with fixed replenishment intervals. *Operations Research* 55:674–687, 2007.

[42] P. Zipkin. *Foundations of Inventory Management*. McGraw-Hill, New York, 2000.