http://pubsonline.informs.org/journal/opre

**OPERATIONS RESEARCH**
*Articles in Advance*, pp. 1–9
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Contextual Areas

# A Simple Heuristic Policy for Stochastic Distribution Inventory Systems with Fixed Shipment Costs

**Han Zhu,[a] Youhua (Frank) Chen,[b] Ming Hu,[c] Yi Yang[d,*]**

[a] School of Management Science and Engineering, Dongbei University of Finance and Economics, Dalian, Liaoning 116025, China; [b] Department of Management Sciences, City University of Hong Kong, Kowloon Tong, Hong Kong; [c] Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada; [d] School of Management, Zhejiang University, Hangzhou, Zhejiang 310027, China

*Corresponding author

**Contact:** hanzhu@dufe.edu.cn, https://orcid.org/0000-0002-5185-3324 (HZ); youhchen@cityu.edu.hk, https://orcid.org/0000-0003-4707-9361 (Y(F)C); ming.hu@rotman.utoronto.ca, https://orcid.org/0000-0003-0900-7631 (MH); yangyizju@zju.edu.cn, https://orcid.org/0000-0001-5161-991X (YY)

**Received:** April 21, 2020
**Revised:** May 20, 2020; September 8, 2020
**Accepted:** September 21, 2020
**Published Online in Articles in Advance:** March 25, 2021

**Subject Classification:** approximations/heuristics; multi-item/echelon/stage; stochastic models
**Area of Review:** Operations and Supply Chains

https://doi.org/10.1287/opre.2021.2101



**Abstract.** We study a continuous-review, two-echelon inventory system with one central warehouse, multiple local facilities, and each facility facing random demand. Local facilities replenish their stock from the central warehouse (or distribution center), which in turn places orders at an outside supplier with ample supply. Inventory replenishment at each location incurs a fixed-plus-variable cost for each shipment. The optimal policy remains unknown, and even if it exists, such a policy must be extremely complicated. Instead, we evaluate a class of easy-to-implement heuristics, called modified echelon $(r, Q)$ policies. The parameters for such a heuristic are obtained by solving a set of independent single-stage systems. We show that the proposed policy is asymptotically optimal, as pairs of system primitives, such as the ratios of the fixed cost of the central facility to those of the local facilities, are scaled up. We also show that as the number of retailers grows, the performance bound of the heuristic converges to a primitive-dependent constant.

**Funding:** H. Zhu received support from National Natural Science Foundation of China [Grants 71902018 and 71831003]. Y. F. Chen received support from Hong Kong Research Grants Council [Project number CityU 190413]. M. Hu received support from the Natural Sciences and Engineering Research Council of Canada [Grant RGPIN-2015-06757]. Y. Yang received support from National Natural Science Foundation of China [Grants 71871198, 71931009, and 71821002].
**Supplemental Material:** The online appendices are available at https://doi.org/10.1287/opre.2021.2101.

**Keywords:** multi-echelon • distribution system • stochastic demand • performance bound • $(r, Q)$ policy

## 1. Introduction

We consider a distribution system that consists of one warehouse and multiple stores, or one central distribution center and multiple local warehouses/last-mile fulfillment centers, which we refer to as one warehouse and multiple retailers (OWMR) to stay consistent with the vast literature. Inventory is continuously reviewed in the warehouse and at each retailer. The warehouse (or distribution center) replenishes the retailers (or last-mile fulfillment centers) and receives stock from an outside supplier. The local retailers face random demands. Regardless of its size, each shipment, either from the supplier to the warehouse or from the warehouse to a retailer, incurs operational frictions in the forms of a positive constant lead time and a positive fixed shipment cost. Holding costs are charged for each unit carried in the warehouse and at the retailers. Excess demand that cannot be immediately satisfied by each retailer is fully backlogged, but incurs a backlogging cost. Shipment decisions are made with the objective of minimizing the long-run average system-wide cost.

The optimal policy of such a system, even if it exists, must be extremely complicated. Our objective is to identify an easy-to-implement policy that has a theoretical foundation, is intuitive, and performs well. For this purpose, we propose a modified echelon $(r, Q)$ policy that operates as follows: if the echelon inventory position at installation $i$ is at or below $r_i$ and the upstream installation has positive on-hand inventory, then a shipment is sent to installation $i$ to raise its echelon inventory position as close as possible to $r_i + Q_i$. Such a modified $(r, Q)$ policy is intuitive and easy to implement. However, the exact system-wide cost of a given modified echelon $(r, Q)$ policy is difficult to evaluate, which incapacitates searching for an optimal one. To overcome this challenge, we adopt a novel approach to derive an easy-to-compute upper bound for the system-wide cost. This upper bound can be expressed as the sum of $N+1$ subsystems, with each location/installation corresponding to a single-stage subsystem. Based on this upper bound, we further identify a specific modified echelon $(r, Q)$ policy with the values of $r$ and $Q$ for each installation optimizing a single-stage subsystem. The computational procedure and efficiency take the virtue of a standard single-stage system. Moreover, by comparing the upper bound with a lower bound of the optimal system-wide cost established by Chen and Zheng (1994), we can explore asymptotic optimality results when some pairs of system parameters, among fixed, holding, and shortage costs, are scaled up. We also show that as the number of retailers goes to infinity, the performance bound of the heuristic converges to a primitive-dependent constant.

## 2. Literature Review

The OWMR distribution system with random demands has been studied extensively. Various heuristics and evaluation methods have been examined; see Simchi-Levi and Zhao (2012) for a survey.

For continuous-review OWMR models with batch-ordering policies, Chen and Zheng (1997) examine an echelon-stock $(r, Q)$ policy and provide an exact cost evaluation method for total inventory holding and shortage costs, excluding fixed costs. Several papers have also considered modified batch-ordering policies that outperform classical $(r, Q)$ polices; see Moinzadeh (2002), Axsäter and Marklund (2008), and Özer (2003). However, none of these studies take into account any fixed cost at a stage. For OWMR models, only a few papers have considered fixed costs. A stream of research aims to provide lower bounds for the optimal system-wide cost. Examples include Clark and Scarf (1960) and Chen and Zheng (1994). Another stream aims to develop heuristic policies for distribution systems with fixed costs. In particular, as an order may be filled in multiple shipments, two different accounting schemes are used to log fixed costs: one is to charge a fixed cost per order (see, e.g., Shang et al. 2015) and the other is to charge a fixed cost per shipment. In this paper, we adopt the shipment-based fixed-cost scheme that is also consistent with the original work of Clark and Scarf (1960). To our knowledge, this work is the first one to study heuristic policies for a stochastic distribution system with shipment-based fixed costs.

For a general distribution system with or without fixed costs, although most of the heuristic policies proposed in the literature make intuitive sense, how good their performance is remains unclear in theory, not to mention their (asymptotic) optimality. Gallego et al. (2007) show that one of their heuristic policies is asymptotically optimal in the number of retailers within a class of local base-stock policies for distribution systems without fixed costs. Chu and Shen (2010) study a distribution system with service-level constraints and develop a power-of-two policy, which is shown to be 1.26-optimal within a class of power-of-two policies. Although base-stock and power-of-two policies perform well in single-stage and deterministic systems, respectively, whether they still perform well in stochastic distribution systems remains unclear. Apart from these two papers, we are unaware of other papers on the optimality or even asymptotic results of heuristic policies for a stochastic distribution inventory system.

Methodologically, our paper is closely related to Hu and Yang (2014), who apply a class of modified echelon $(r, Q)$ policies to a serial inventory system. However, there exist fundamental differences between serial and distribution systems. For example, stock allocation among retailers, as a feature of distribution systems, is absent in a serial system. Specifically, in a distribution system, if the on-hand inventory in the warehouse is insufficient to satisfy all retailers in need, then it is the allocation policy that determines who among the retailers gets what or not at all. As a result, the allocation policy determines inventory flows, and hence, directly affects the system-wide cost. In addition, the interaction between the warehouse and a specific retailer in OWMR models may affect the other retailers because all retailers are linked through the common warehouse. For example, after satisfying a retailer's order, the warehouse may not have sufficient stock to fully satisfy the subsequent order requests from other retailers, thereby making system dynamics difficult to analyze. These differences make the detailed analysis of a distribution system fundamentally different from and more challenging than that of a serial system.

### Contributions

Our contribution is threefold. First, our asymptotic optimality results demonstrate the robustness of single-stage $(r, Q)$ inventory policies. That is, $(r, Q)$ policies based on single-stage systems with some adaptations can still perform well even in a distribution system. As the single-stage $(r, Q)$ policy is easy to compute and implement, this result is of managerial importance and relevance. Second, our model studies a stochastic distribution system with fixed shipment costs, which is more consistent with practice in logistics. We fill in a gap in literature by developing an easy-to-implement heuristic policy that is computationally efficient and asymptotically optimal. Third, compared with other heuristics developed for distribution systems, our proposed heuristic policy has an important advantage, as it does not require a nested integer ratio or synchronized ordering property for the replenishment policies.[1] Our policy imposes no such coordination requirements on the batch sizes or ordering time, but still performs efficiently with asymptotic optimality properties.

## 3. Model

We consider a firm that manages a two-echelon distribution inventory system consisting of one warehouse and $N$ retailers. For notation convenience, we use $[N]^+$ and $[N]$ to denote the sets of $\{1,2,\ldots,N\}$ and $\{0,1,\ldots,N\}$, respectively. We use retailer $i \in [N]^+$ to denote a specific retailer and installation $i \in [N]$ to denote a specific installation, which can be either the warehouse $i = 0$ or a specific retailer $i$. Retailers are replenished from the warehouse, which, in turn, is replenished from an outside supplier with unlimited stock. Retailer $i$ faces a steam of random demand following a Poisson process with stationary rate $\lambda_i$. Demands across retailers are assumed to be independent. Moreover, due to independence among different retailers' demands and the superposition property of the Poisson process, the entire system still faces a Poisson process with a demand rate $\lambda_0 \equiv \sum_{i=1}^N \lambda_i$. For installation $i$, we use $D_i(t,t+\tau]$ to denote the total demand of installation $i$ over the time interval $(t,t+\tau]$. For the warehouse, $D_0(t,t+\tau] = \sum_{i=1}^N D_i(t,t+\tau]$ represents the total demand of all retailers over the time interval $(t,t+\tau]$. A constant lead time $L_i > 0$ exists for installation $i$. That is, any shipment sent out to installation $i$ at time $t$ will be received by installation $i$ at time $t+L_i$. Each shipment to installation $i$ incurs a fixed cost $K_i$. Without loss of generality, we assume that the variable ordering cost is zero. Let $h_i > 0$ be the echelon holding cost rate at installation $i$. Whenever retailer $i$ runs out of stock, the unmet demand is fully backlogged with a backlog cost rate of $p_i > 0$. The firm's objective is to determine a replenishment policy that minimizes the long-run average system-wide cost. For simplicity, we adopt the notation with subscript $i$ to denote installation $i$. Given that the optimal policy of such a system, even if it exists, must be extremely complicated, we focus on a class of modified echelon $(r,Q)$ policies as follows.

**Definition 1** (Modified Echelon $(r,Q)$ Policy)**.**

The modified echelon $(r,Q)$ policy consists of replenishment and allocation rules.

**Replenishment.** The upstream installation ships to a downstream installation on the basis of its observation of the echelon inventory position in the downstream installation: If the echelon inventory position in installation $i$ is at or below $r_i$ and the upstream installation has positive on-hand inventory, then a shipment is sent to installation $i$ to raise its echelon inventory position $IP_i$ to $r_i + Q_i$ as close as possible.

It should be noted that due to the possible shortage at the warehouse, it is possible that the warehouse following this replenishment rule has no on-hand inventory to ship to a requesting retailer. In such a case, the retailer keeps on triggering its replenishment request until the warehouse has some on-hand inventory to ship, and such a trigger will be terminated only if a shipment (irrespective of its size) is sent to the retailer to raise its echelon inventory position $IP_i > r_i$. In addition to the replenishment rule, the modified echelon $(r,Q)$ policy also includes the following allocation rule.

**Allocation.** When facing shortages from multiple retailers, the warehouse can determine an arbitrary fulfillment sequence. Once the sequence is determined, the warehouse sequentially fulfills retailers one at a time and raises their inventory position as close as possible to their target.

In case of shortages at multiple retailers, our allocation rule is quite flexible and can be in any arbitrary sequence, including several commonly used rules such as the first-come, first-served rule, the last-come, first-served rule, and a priority scheme based on historic sales. Notably, such a sequence needs not to be prefixed and can vary over time.

## 4. Cost Analysis

To facilitate analysis, at time $t$, we charge the inventory holding cost incurred in the warehouse at time $t+L_0$, and charge the inventory holding and backlog costs incurred in retailer $i$ at time $t+L_0+L_i$. This cost accounting scheme only shifts costs across time points and hence does not affect the long-run average expected inventory holding and backlog costs.

Given that the external supplier has ample inventory, a modified echelon $(r,Q)$ policy for the warehouse operates as a classical echelon $(r,Q)$ policy with parameters $(r_0,Q_0)$. Therefore, the steady state of its inventory position $IP_0(t)$ is uniformly distributed among $\{r_0+1,\ldots,r_0+Q_0\}$. We use $\Omega(IP_0(t))$ to denote the total expected average cost rate of all retailers at time $t$ when the inventory position of the warehouse is $IP_0(t)$. Following Zheng (1992) and Chen and Zheng (1994), we assume that discrete units of inventories can be approximated by continuous variables. Let $D_0$ denote the total demand over $(0,L_0]$. Given an arbitrarily modified echelon $(r,Q)$ policy, the long-run average system-wide cost $C(\mathbf{r},\mathbf{Q})$ can be expressed as

$$C(\mathbf{r},\mathbf{Q}) = \frac{1}{Q_0}\left[\lambda_0 K_0 + \int_{r_0}^{r_0+Q_0} \mathbb{E}\big[h_0(y-D_0)+\Omega(y)\big]dy\right], \tag{1}$$

where $\mathbf{r} = (r_0,r_1,\ldots,r_N) \in \mathbb{N}^{N+1}$ and $\mathbf{Q} = (Q_0,Q_1,\ldots,Q_N) \in \mathbb{N}^{N+1}$. However, considering the possible shortage in the warehouse, directly computing $\Omega(IP_0(t))$ and $C(\mathbf{r},\mathbf{Q})$ is extremely difficult. To overcome this, we propose a novel approach for obtaining an upper bound for $C(\mathbf{r},\mathbf{Q})$. We can construct a specific heuristic by optimizing this upper bound and further derive asymptotic results by comparing this upper bound with an existing lower bound (see Chen and Zheng 1994).

### 4.1. Outline of the Upper Bound Construction

First, we divide the entire time horizon into cycles by tracing how the first shipment of an order of the warehouse is sent to a specific downstream retailer.

**Definition 2** (Cycle). For $j \in \mathbb{N}$, we denote by $[T^j, T^{j+1})$ the $j$th cycle of retailers, where $T^j$ is the time epoch of the first unit in the $j$th order of the warehouse being sent to a retailer. That is, $T^j$ is the time epoch when the first unit in the warehouse's $j$th order leaves the warehouse.

The length of a cycle is random and dependent on the realization of demands. A specific cycle $[T^j, T^{j+1})$ may also be an empty set. In such a case, all units contained in the $j$th order of the warehouse are shipped to a retailer along with one or multiple units in the $(j+1)$th order of the warehouse. Our analysis focuses on nonempty cycles because retailers incur no costs for empty cycles. The following lemma characterizes the expected long-run average cycle length. Somewhat surprisingly, the expected cycle length depends only on the policy parameters of the warehouse. The omitted proofs in this paper can be found in the online appendix.

**Lemma 1** (Long-Run Cycle Length). *Under any modified echelon $(r, Q)$ policy, the long-run average expected cycle length is $Q_0/\lambda_0$, that is, $\lim_{j\to\infty} \mathbb{E}[(T^{j+1} - T^1)]/j = Q_0/\lambda_0$.*

We then investigate the shipments from the warehouse to the retailers over a cycle. To do so, we classify the shipments from the warehouse to the retailers into regular and irregular ones, and further classify irregular shipments into types I and II shipments.

Over any nonempty cycle $[T^j, T^{j+1})$, the $j$th order of the warehouse is shipped to retailers in one or multiple, say $M \in \mathbb{N}$ shipments in total. Among the $M$ shipments, denote by $M_i \in \mathbb{Z}^+$ the number of shipments sent to retailer $i$. Then, $\sum_{i=1}^N M_i = M$. Let $T_i^{j,m}$ be the time of the $m$th shipment sent to retailer $i$ over the cycle $[T^j, T^{j+1})$, where $m = 1, 2, \ldots, M_i$. By definition, we have $T^j \le T_i^{j,1} \le \cdots \le T_i^{j,M_i} < T^{j+1}$ for any $i$. Defining $T_i^{j,M_i+1} \equiv T_i^{j+1,1}$, we call $[T_i^{j,m}, T_i^{j,m+1})$ the $m$th shipment interval of retailer $i$ over the cycle $[T^j, T^{j+1})$. For the case with $M_i = 0$, that is, when no shipment is sent to retailer $i$, we define $T_i^{j,1} \equiv T_i^{j+1,1}$.

Depending on a retailer's inventory position at the beginning and the end of a shipment interval, we categorize the shipments to retailers and their associated shipment intervals as follows.

**Definition 3** (Regular and Irregular Shipment (Interval)). For a shipment interval $[T_i^{j,m}, T_i^{j,m+1})$, if $IP_i(T_i^{j,m}) = r_i + Q_i$ and $IP_i^-(T_i^{j,m+1}) = r_i$, then we call it a regular shipment interval of retailer $i$; otherwise, we call it an irregular shipment interval. Specifically, if $IP_i(T_i^{j,m}) = r_i + Q_i$ and $IP_i^-(T_i^{j,m+1}) < r_i$, then we call it a type I irregular shipment interval; if $IP_i(T_i^{j,m}) < r_i + Q_i$, then we call it a type II irregular shipment interval. The shipment associated with a regular (or type I or II irregular) shipment interval is called a regular (or type I or II irregular) shipment.

Finally, we adopt different cost assessment schemes for these shipment types (see Section 4.2.2). Based on these schemes, we construct a cost upper bound of retailers' total expected cost $\Omega(IP_0(t))$ over a specific cycle. Although the total costs of retailers vary across different cycles, they share the same upper bound (see Section 4.2.3).

### 4.2. Cost Upper Bound Construction

In this subsection, we discuss in detail our approach of constructing the cost upper bound.

#### 4.2.1. Shipment Intervals.

For $m = 1, \ldots, M_i - 1$ with $M_i \in \mathbb{N}$, the shipment interval $[T_i^{j,m}, T_i^{j,m+1})$ of retailer $i$ is located within the cycle $[T^j, T^{j+1})$. However, the shipment interval associated with the last shipment in this cycle, $[T_i^{j,M_i}, T_i^{j,M_i+1})$, may not be contained within the cycle $[T^j, T^{j+1})$. The beginning of this interval must be within this cycle, that is, $T_i^{j,M_i} \le T^{j+1}$, which is why we associate this interval with the cycle $[T^j, T^{j+1})$. However, the end of this shipment interval, $T_i^{j,M_i+1}$, may be outside the cycle. This occurs when after shipment at $T_i^{j,M_i}$, no shipment is sent to retailer $i$ over $[T_i^{j,M_i}, T^{j+1})$ and the next shipment to retailer $i$ occurs at $T_i^{j+1,1}$, which is strictly later than $T^{j+1}$.

Whether a shipment interval is regular depends on the retailer's inventory positions at the beginning and end of this interval. The type of an irregular shipment interval depends only on the inventory position at the beginning of this interval. To better illustrate how retailers' inventory positions evolve within a cycle with multiple shipment intervals, we display one scenario of two retailers' inventory positions over a cycle $[T^j, T^{j+1})$ in Figure 1.

The following lemma shows that the frequency of both irregular shipments for each retailer and type II irregular shipments for all retailers can be bounded by that of cycles.

**Lemma 2** (Irregular Shipment Frequency). *(i) For any retailer, there exists at most one irregular shipment interval within each cycle, regardless of whether the cycle is empty or not. (ii) Across all retailers, there exists at most one type II irregular shipment interval within each cycle.*

#### 4.2.2. Cost Assessment in Cycles.

In this subsection, we provide a cost assessment over a cycle. For retailer $i \in [N]^+$, following Chen and Zheng (1994), $G_i(IP_i(t))$ denotes its holding and shortage cost rates at time $t + L_i$ in terms of its inventory position at time $t$,

$$G_i(IP_i(t)) = \mathbb{E}\big[h_i(IP_i(t) - D_i(t, t+L_i))^+ . \\ +\big(h_0 + p_i\big)(IP_i(t) - D_i(t, t+L_i))^-\big]. \quad (2)$$

**Figure 1.** (Color online) Illustration of Two Retailers' Inventory Positions over a Cycle $[T^j, T^{j+1})$

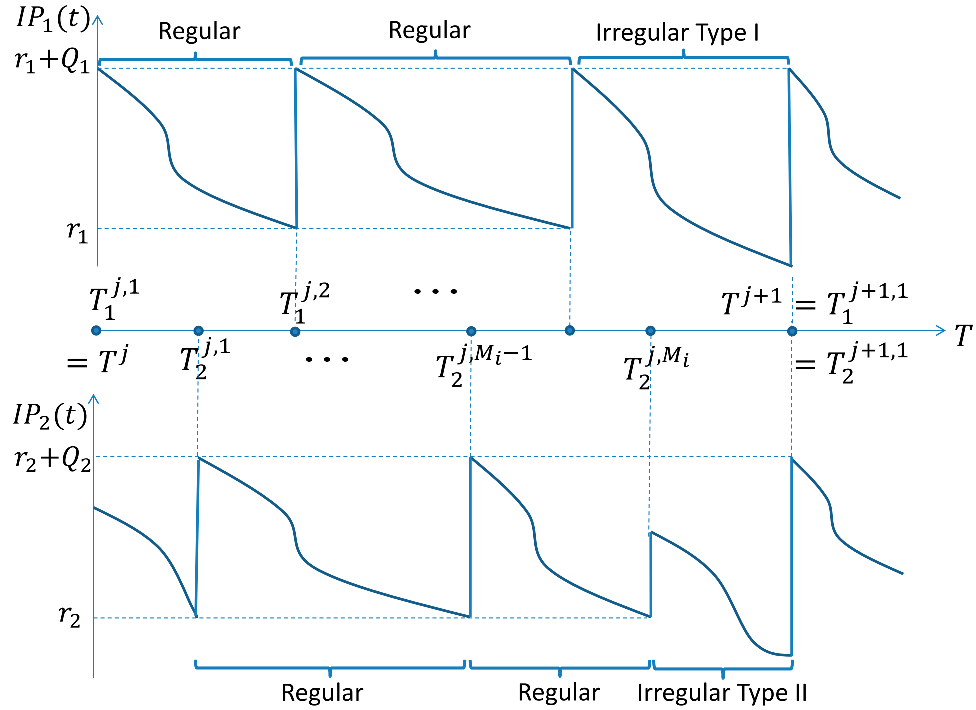


We use the following cost assessment schemes for the three types of shipment within a cycle:

**Cost in regular shipment intervals.** In a regular shipment interval, the inventory position of retailer $i$ gradually drops from $r_i + Q_i$ to $r_i$. At the time when $r_i$ is reached, another shipment is triggered and the regular shipment ends. It behaves the same as a single-stage system with an outside supplier of unlimited supply; see Zheng (1992). Therefore, the total expected cost rate, including the fixed, inventory holding, and backlog costs, of retailer $i$ in a regular interval is

$$C_i(r_i, Q_i) = \frac{\lambda_i K_i + \int_{r_i}^{r_i+Q_i} G_i(y)dy}{Q_i}, \tag{3}$$

where $G_i(y)$ is defined in (2).

**Cost in type I irregular shipment intervals.** In a type I irregular shipment interval, the inventory position of retailer $i$ initially drops from $r_i + Q_i$ to $r_i$ and then drops below $r_i$. In the former subinterval, the expected cost rate is $C_i(r_i, Q_i)$, which is the same as in a regular shipment interval, whereas in the latter, the expected inventory holding and backlog costs are accrued at a rate equal to $G_i(IP_i(t))$. No inventory replenishment occurs and thus, no fixed cost is incurred in the latter subinterval.

**Cost in type II irregular shipment intervals.** In a type II irregular shipment interval, we separately calculate the fixed cost and holding/backlog costs. For any time $t$ in a nonempty type II irregular shipment interval of retailer $i$, the expected inventory holding and backlog costs are accrued at a rate equal to $G_i(IP_i(t))$. In addition, a fixed cost $K_i$, is incurred for the irregular shipment.

The system-wide costs can be decomposed into (i) costs in regular shipment intervals, (ii) costs in type I irregular shipment intervals, (iii) inventory holding and backlog costs in type II irregular shipment intervals, and (iv) fixed costs in type II irregular shipment intervals. We next provide the cost upper bounds for the first three parts and for the last part separately.

**4.2.3. Cost Upper Bound.** We first provide an upper bound on the fixed costs associated with type II irregular shipments. From Lemma 2(i), the frequency of incurring irregular shipment intervals is bounded by the frequency of cycles. Define $\overline{K} = \max_{i\in[N]^+}\{K_i\}$. By Lemma 2(ii), the fixed cost for type II irregular shipments over any cycle $[T^j, T^{j+1}) \neq \emptyset$ is incurred at most once and should not exceed $\overline{K}$. Consequently, the corresponding fixed costs for type II irregular shipments are accrued at a rate that is no more than $\overline{K}/(T^{j+1} - T^j)$. Then, by Lemma 1, the long-run average fixed cost for type II irregular shipments has an upper bound $\lambda_0\overline{K}/Q_0$. Then we can bound the expected cost rate of all retailers, excluding the fixed costs for type II irregular shipments. The bound is expressed in terms of the echelon inventory level of the warehouse. Let $IL_0(t)$ denote the echelon inventory level at the warehouse, that is, the echelon inventory at the warehouse (the on-hand inventory at the warehouse plus the inventories at or in transit to all retailers) minus the total number of customers back-ordered at all retailers.

**Lemma 3.** *For any time $t \in [T^j, T^{j+1}) \neq \emptyset$, the expected cost rate of all retailers, excluding the fixed costs for type II irregular shipments, denoted by $\hat{\Gamma}(IL_0(t))$, is upper bounded as*

$$\hat{\Gamma}(IL_0(t)) \leq \bar{\Gamma}(IL_0(t)) \equiv \max_{i \in [N]^+} \bar{\Gamma}_i(IL_0(t)), \tag{4}$$

*where*

$$\bar{\Gamma}_i(IL_0(t)) \equiv \sum_{j \neq i} \max\{G_j(w_j), C_j(r_j, Q_j)\} + \begin{cases} \max\{G_i(w_i), C_i(r_i, Q_i)\} \\ \qquad \text{if } IL_0(t) - \sum_{j \neq i}(r_j + Q_j) > r_i, \\ \max\left\{G_i\left(IL_0(t) - \sum_{j \neq i}(r_j + Q_j)\right), G_i(w_i), C_i(r_i, Q_i)\right\} \\ \qquad \text{otherwise,} \end{cases}$$

*and $w_i \equiv \arg\max_{r_i < z \leq r_i + Q_i}\{G_i(z)\}$.*

As mentioned, multiple retailers with an inventory position that is less than their reorder level may exist. One salient feature of Lemma 3 is that we construct a cost upper bound for a given retailer $i$ in terms of $IL_0(t)$, instead of $IP_i(t)$. To this end, we charge the cost at a time for retailer $i$ in terms of $IL_0(t)$ in the worst-case scenario, in which the inventory position of all other $N-1$ retailers than retailer $i$ reaches their highest possible inventory position, that is, $r_j + Q_j$. It implies that retailer $i$ has the lowest possible inventory position, that is, $IL_0(t) - \sum_{j \neq i}(r_j + Q_j)$, which leads to a cost upper bound through the convexity of $G_i(y)$.

Although the upper bound in Lemma 3 is given in terms of $IL_0(t)$, the system-wide cost can be bounded in terms of $IP_0(t)$ based on the relationship $IL_0(t+L_0) = IP_0(t) - D_0(t, t+L_0]$. By combining the aforementioned cost bounds, we are ready to present an upper bound for the total expected costs by noting that $\Omega(y) \leq \lambda_0\bar{K}/Q_0 + \mathbb{E}[\hat{\Gamma}(y - D_0)]$, and derive an upper bound for $C(\mathbf{r}, \mathbf{Q})$ in (1).

We define

$$\hat{G}(y) \equiv \bar{\Gamma}(y) - \sum_{i=1}^{N} C_i(r_i, Q_i), \tag{5}$$

$$\Lambda_0(y) \equiv \mathbb{E}[h_0(y - D_0) + \hat{G}(y - D_0)], \tag{6}$$

$$\hat{C}_0(r_0, Q_0) \equiv \frac{1}{Q_0}\left[\lambda_0 K_0 + \int_{r_0}^{r_0+Q_0} \Lambda_0(y)dy\right]. \tag{7}$$

Recall that $\bar{\Gamma}(\cdot)$, defined in Lemma 3, is the upper bound for the expected cost rate of all retailers, excluding the fixed costs of type II irregular shipments. Then, $\hat{C}_0(r_0, Q_0)$ is the upper bound for the expected cost rate of all installations, excluding $\sum_{i=1}^{N} C_i(r_i, Q_i)$ and the fixed costs of type II irregular shipments. As mentioned, $\lambda_0\bar{K}/Q_0$ is an upper bound on the long-run average fixed cost for type II irregular shipments. Then we have the following system-wide cost upper bound.

**Theorem 1** (An Upper Bound)**.** *For a given modified echelon $(r, Q)$ policy, the long-run average system-wide cost has an upper bound:*

$$C(\boldsymbol{r}, \mathbf{Q}) \leq C^U(\boldsymbol{r}, \mathbf{Q}) \equiv \sum_{i=1}^{N} C_i(r_i, Q_i) + \hat{C}_0(r_0, Q_0) + \frac{\lambda_0\bar{K}}{Q_0}.$$

## 5. Heuristic Policy: Performance Bounds and Asymptotic Optimality

In this section, we propose a heuristic modified echelon $(r, Q)$ policy by minimizing the upper bound $C^U(\mathbf{r}, \mathbf{Q})$. Directly optimizing $C^U(\mathbf{r}, \mathbf{Q})$ appears difficult, given that $\hat{C}_0(r_0, Q_0)$ depends on $r_i$ and $Q_i$, $i \in [N]^+$, and may not be a convex function.

We define $(r_i^*, Q_i^*)$ as the minimizer of $C_i(r_i, Q_i)$, $i \in [N]^+$. In our proposed heuristic policy, we choose parameters $(r_i, Q_i)$ by optimizing $C_i(r_i, Q_i)$, without considering $\hat{C}_0(r_0, Q_0)$. That is, we select $(r_i, Q_i) = (r_i^*, Q_i^*)$ for retailer $i$, which means that the retailers just use the optimal decisions of their single-stage inventory problems. With this selection, we have $C_i(r_i^*, Q_i^*) = C_i^*$ and

$$C(\mathbf{r}, \mathbf{Q})\Big|_{(r_i, Q_i) = (r_i^*, Q_i^*)} \leq \sum_{i=1}^{N} C_i^* + \hat{C}_0(r_0, Q_0) + \frac{\lambda_0\bar{K}}{Q_0}. \tag{8}$$

Plugging (7) into (8), we have

$$C(\mathbf{r}, \mathbf{Q})\Big|_{(r_i, Q_i) = (r_i^*, Q_i^*)} \leq \sum_{i=1}^{N} C_i^* + \frac{1}{Q_0}\left[\lambda_0 K_0 + \int_{r_0}^{r_0+Q_0} \Lambda_0(y)dy\right] + \frac{\lambda_0\bar{K}}{Q_0}. \tag{9}$$

We can further tighten the upper bound in (9) by minimizing it over $r_0$ and $Q_0$. Specifically, this can be achieved by solving the following optimization problem:

$$\min_{r_0, Q_0} \tilde{C}_0(r_0, Q_0) \equiv \min_{r_0, Q_0} \frac{1}{Q_0}\left[\lambda_0(K_0 + \bar{K}) + \int_{r_0}^{r_0+Q_0} \Lambda_0(y)dy\right]. \tag{10}$$

Note that $\bar{\Gamma}_i(y)$ (see Lemma 3) may not be convex for any given $r_i$ and $Q_i$ ($i = 1, 2, \ldots, N$), but it is convex for $(r_i, Q_i) = (r_i^*, Q_i^*)$. Consequently, with $(r_i, Q_i) = (r_i^*, Q_i^*)$ for all $i$, $\hat{G}(y)$ in (5), and hence, $\Lambda_0(y)$ in (6), are convex functions as well.

Problem (10) is a single-stage inventory system with a fixed cost equal to $K_0 + \bar{K}$, and its objective function $\tilde{C}_0(r_0, Q_0)$ is jointly convex in $r_0$ and $Q_0$. Therefore, the optimal solution can be efficiently computed. Define $(\tilde{r}_0^*, \tilde{Q}_0^*) \equiv \arg\min_{r_0, Q_0} \tilde{C}_0(r_0, Q_0)$. We construct a heuristic modified echelon $(r, Q)$ policy as

$$\begin{aligned}(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) &= (\hat{r}_0, \hat{r}_1, \ldots, \hat{r}_N, \hat{Q}_0, \hat{Q}_1, \ldots, \hat{Q}_N) \\ &= (\tilde{r}_0^*, r_1^*, \ldots, r_N^*, \tilde{Q}_0^*, Q_1^*, \ldots, Q_N^*),\end{aligned} \tag{11}$$

and refer to it as the Multi-Echelon $(r, Q)$ Distribution (MERQD) policy.

From theorem 1 in Zheng (1992), we have $\tilde{C}_0^* \equiv \tilde{C}_0(\tilde{r}_0^*, \tilde{Q}_0^*) = \Lambda_0(\tilde{r}_0^*)$. Compared with the bound in (9), a tighter upper bound on the system-wide cost can be expressed as

$$C_{\mathcal{B}}^* \leq C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) \equiv C(\mathbf{r}, \mathbf{Q})\Big|_{(\mathbf{r},\mathbf{Q})=(\hat{\mathbf{r}},\hat{\mathbf{Q}})} \leq \sum_{i=1}^{N} C_i^* + \tilde{C}_0^*. \quad (12)$$

We can derive guaranteed bounds on the effectiveness of our heuristic policy by comparing the upper bound in (12) with the induced-penalty lower bound $\underline{C}^* = \sum_{i=0}^{N} C_i^*$ established by Chen and Zheng (1994). Notably, $C_0^*$ may take negative values. Our main results are based on the assumption that $C_0^* > 0$. The following assumption is a sufficient condition for $C_0^* > 0$; see lemma 5 of Zhu et al. (2020).

**Assumption 1.** *Assume* $h_0 \sum_{i=1}^{N} (\lambda_i L_i - \frac{C_i^*}{h_0+p_i}) + \sqrt{\frac{2\lambda_0 K_0 h_0 \underline{p}}{h_0+\underline{p}}} > 0$, *where* $\underline{p} = \min_{i=1,2,\ldots,N} p_i$.

For notational convenience, let $(\hat{r}_0^*, \hat{Q}_0^*) = \arg\min_{r_0,Q_0} \hat{C}_0(r_0, Q_0)$ and $\hat{C}_0^* = \hat{C}_0(\hat{r}_0^*, \hat{Q}_0^*)$.

**Theorem 2** (Performance Bounds for General Cases)**.** *We have:*

i. *The MERQD policy is* $(1+\frac{\tilde{C}_0^*-C_0^*}{\sum_{i=1}^{N} C_i^*+C_0^*})$*-optimal, that is,* $\frac{\sum_{i=1}^{N} C_i^*+\tilde{C}_0^*}{\sum_{i=1}^{N} C_i^*+C_0^*}$*-optimal.*

ii. *The MERQD is* $\max\left\{\sqrt{\frac{\lambda_0}{2\beta_1\beta_2\lambda_m} + \frac{1}{4}} + \frac{1}{2}, \frac{1}{\beta_2}\right\}$*-optimal, where* $m \in \arg\max_{i\in[N]^+}\{K_i\}$, $\beta_1 \equiv \frac{\hat{Q}_0^*}{Q_m^*}$, *and* $\beta_2 \equiv \frac{C_0^*}{\hat{C}_0^*} \leq 1$.[2]

Theorem 2(i) directly follows from comparing the induced-penalty lower bound with the upper bound established in (12). The performance bound is expressed in terms of the optimal costs for a set of single-stage, single-period inventory systems that are simple and easy to compute. Theorem 2(ii) provides a looser performance bound that depends on the order size ratio $\beta_1 \equiv \hat{Q}_0^*/Q_m^*$ and the cost ratio $\beta_2 \equiv C_0^*/\hat{C}_0^*$. In a supplement, figure 1 in Zhu et al. (2020) displays the contour plot of the performance bound as a function of $\beta_1$ and $\beta_2$. The identified heuristic performs well when the two ratios are large. Denote $f_0 \equiv \lambda_0/\hat{Q}_0^*$ and $f_i \equiv \lambda_i/Q_i^*$ for $i \in [N]^+$. The measure $f_i$ represents the replenishment frequency of installation $i \in [N]$ under our heuristic assuming that each installation's replenishment can always be fulfilled. Then, the performance guarantee in Theorem 2(ii) can be rewritten as $\max\left\{\sqrt{\frac{\beta_0}{2\beta_2} + \frac{1}{4}} + \frac{1}{2}, \frac{1}{\beta_2}\right\}$, where $\beta_0 \equiv f_0/f_m$ is the ratio of two replenishment frequencies for single-stage inventory systems. This implies that the theoretical performance bound can be expressed by a replenishment frequency ratio ($\beta_0$) and a cost ratio ($\beta_2$) for single-stage inventory systems, both of which can be efficiently computed. It is easy to see that the proposed heuristic tends to perform better when $\beta_2$ is larger and/or $\beta_0$ is smaller. For the case with identical retailers, we can derive a sharper performance bound.

**Corollary 1** (Performance Bounds for Identical Retailers)**.** *Suppose that all retailers are identical. The MERQD policy is* $\max\left\{\sqrt{\frac{1}{2\beta_1\beta_2} + \frac{1}{4}} + \frac{1}{2}, \frac{1}{\beta_2}\right\}$*-optimal, where* $\beta_1 \equiv \hat{Q}_0^*/Q_i^*$ *is identical for all retailers.*

It can be observed from Corollary 1 that when $\beta_1$ and $\beta_2$ become larger, the bound becomes tighter. Suppose that $\beta_2$ is not much smaller than 1, then the bound in Corollary 1 highly depends on the value of $\beta_1$. When $N$ increases, $\beta_1$ tends to become larger. As we show in Theorem 4 and Zhu et al. (2020), the performance of the heuristic converges when $N \to +\infty$.

The following theorem further demonstrates the asymptotic optimality of our heuristic if we scale the dominant relationships of pairs of cost primitives to the extreme.

**Theorem 3** (Asymptotic Optimality)**.** *The MERQD policy is asymptotically optimal if one of the following conditions holds: (i)* $K_m > 0$ *and* $K_0/K_m \to \infty$*; (ii)* $h_0/h_m \to 0$*; and (iii)* $h_0/p_m \to 0$*, where* $m \in \arg\max_{i\in[N]^+}\{K_i\}$.

Theorem 3 shows the asymptotic optimality of our heuristic when some dominant relationships among system primitives are taken to the extreme. It implies the conditions under which the heuristic policy is more likely to perform well. First, a distribution system typically has larger economies of scales at the warehouse (or central distribution center) than at retailers (or last-mile fulfillment centers). Second, for an e-retailer, the last-mile fulfillment centers (or local distribution centers), particularly those in dense urban areas with limited land availability (and hence high rents), incur a higher inventory holding cost than upstream distribution centers that are typically located in suburban areas. Third, the shortage cost of each retailer tends to be considerably higher than holding costs. In pursuit of high service levels and consumer satisfaction, distribution inventory systems usually adopt a small value of $h_0/p_m$; see, for example, 0.05 in Chen and Zheng (1997) and 0.01–0.1 in Shang et al. (2015).

We next investigate another asymptotic result with respect to the number of retailers.

**Theorem 4.** *Suppose that all retailers are identical. As* $N \to +\infty$*, the MERQD policy is* $(1 + (1 + \frac{p_i+h_0}{h_i})(1 + \frac{h_0}{p_i}))$*-optimal.*

We refer readers to Zhu et al. (2020) for other explicit performance bounds.

## 6. Numerical Study

We report our numerical study in this section. Through a set of comprehensive numerical experiments, we verify the effectiveness of our heuristic policy and test its sensitivity to system parameters. Moreover, we compare our heuristic policy with the echelon-stock $(r, Q)$ policy studied in the literature. Lastly, we also study the impact of the allocation rule on the performance of the MERQD policy. Here we summarize main observations from the numerical experiments:

1. We test the performance of the MERQD policy for 6,561 instances. In general, the MERQD policy performs well with an average gap of about 5% above the cost lower bound and a gap of less than 10% for about 90% of all instances; see online Appendix B.1.

2. We then study the sensitivity of the MERQD policy with respect to $N$, $K_0$, $h_0$, $p_i$, and $\lambda_i$ separately. The MERQD policy performs quite well in most instances, especially when the asymptotic conditions tend to hold, for example, when $h_0/h_i \leq 1$. Furthermore, the MERQD policy still performs well when $L_0$ is much larger than $L_i$; see section 5.1.1 of Zhu et al. (2020).

3. When $h_0/h_i$ is large, the performance of our heuristic may not be as good as those in other cases. To address this issue, we particularly develop an alternative heuristic policy (referred to as heuristic CD under which the warehouse only acts as a Cross-Docking distributor without holding any inventory); see also Marklund (2002, 2006), Berling and Marklund (2006), and Axsäter et al. (2013) for other possible effective policies. Together with the MERQD policy, heuristic CD can form a new hybrid policy that chooses one between the two. For the cases where the MERQD policy may not perform well, heuristic CD can greatly improve the performance of the hybrid policy; see section 5.1.2 of Zhu et al. (2020).

4. We compare our MERQD policy with the so-called echelon-stock $(r, Q)$ policy used in Chen and Zheng (1997). Our proposed policy outperforms the echelon-stock $(r, Q)$ policy in all tested instances, and the gap can be as large as 9%. Compared with their policy, our policy generates slightly higher inventory costs, but a considerably lower fixed cost; see section 5.2 of Zhu et al. (2020).

5. We finally test the robustness of the MERQD policy with respect to several alternative allocation rules by the warehouse. The performance of the MERQD policy seems quite robust to allocation rules for those instances in which it performs well. For the instances in which our policy does not perform well, its performance can be improved by a better allocation rule. For example, in the scenario with identical retailers, the inventory position priority rule (in which retailers are replenished based on the reverse order of their inventory position) outperforms other rules and can save more than 4% of the total cost; see section 5.3 of Zhu et al. (2020).

## 7. Conclusion

We study a classical distribution inventory system with fixed shipment costs. We develop a class of modified echelon $(r, Q)$ policies that do not require a nested integer ratio or synchronized ordering property. We examine the effectiveness of the proposed policy by comparing its performance with a lower bound of the optimal cost, and identify conditions under which the heuristic is asymptotically optimal. Despite the daunting challenges in analyzing an OWMR system, our work represents the first attempt to identify easy-to-compute policies with provable performance guarantees and to derive their asymptotic properties for a distribution system with fixed shipment costs.

Our approach has limitations. First, although our proposed heuristic achieves superior performance under those scenarios with system primitives likely to satisfy asymptotic optimality conditions, it may perform poorly in other cases. It is highly desirable to design an effective heuristic for all scenarios. Second, our model considers a continuous-review system with a simple Poisson demand process. However, our current analysis cannot be readily extended to account for a periodic-review system or a continuous-review system with a compound Poisson arrival process. It is of interest to establish theoretical performance bounds for certain heuristic policies in these two systems. This work hopefully paves the way for future endeavors.

## Acknowledgments

The authors thank the department editor, Tava Olsen, the anonymous associate editor, and three reviewers for their constructive and insightful suggestions, which significantly improved the paper.

## Endnotes

[1] *Nested integer ratio* indicates that the batch size of the warehouse is equal to a positive multiple of that of retailers. *Synchronized ordering* means that whenever the warehouse receives a shipment from the supplier, all retailers must simultaneously place an order. Either or both of the requirements are prevalent in the heuristics proposed in the literature.

[2] In addition, if $(\beta_2^2 - 1)(\beta_1 \lambda_m) + \beta_2 \lambda_0 \geq 0$, then the MERQD policy can be alternatively shown to be at least $1 + \lambda_0 / (2(\beta_1 \beta_2 \lambda_m + \sqrt{(\beta_2^2 - 1)(\beta_1 \lambda_m)^2 + \beta_1 \beta_2 \lambda_m \lambda_0}))$-optimal; see the proof of Theorem 2 in the online Appendix. The numerical results show that the bound in Theorem 2(ii) tends to perform better but not always.

## References

Axsäter S, Marklund J (2008) Optimal position-based warehouse ordering in divergent two-echelon inventory systems. *Oper. Res.* 56(4):976–991.

Axsäter S, Howard C, Marklund J (2013) A distribution inventory model with transshipments from a support warehouse. *IISE Trans.* 45(3):309–322.

Berling P, Marklund J (2006) Heuristic coordination of decentralized inventory systems using induced backorder costs. *Production Oper. Management* 15(2):294–310.

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

Chen F, Zheng YS (1997) One-warehouse multiretailer systems with centralized stock information. *Oper. Res.* 45(2):275–287.

Chu LY, Shen ZJM (2010) A power-of-two ordering policy for one-warehouse multiretailer systems with stochastic demand. *Oper. Res.* 58(2):492–502.

Clark AJ, Scarf H (1960) Optimal policies for a multi-echelon inventory problem. *Management Sci.* 6(4):475–490.

Gallego G, Özer Ö, Zipkin P (2007) Bounds, heuristics, and approximations for distribution systems. *Oper. Res.* 55(3):503–517.

Hu M, Yang Y (2014) Modified echelon ($r$, $Q$) policies with guaranteed performance bounds for stochastic serial inventory systems. *Oper. Res.* 62(4):812–828.

Marklund J (2002) Centralized inventory control in a two-level distribution system with Poisson demand. *Naval Res. Logist.* 49(8): 798–822.

Marklund J (2006) Controlling inventories in divergent supply chains with advance-order information. *Oper. Res.* 54(5):988–1010.

Moinzadeh K (2002) A multi-echelon inventory system with information exchange. *Management Sci.* 48(3):414–426.

Özer Ö (2003) Replenishment strategies for distribution systems under advance demand information. *Management Sci.* 49(3): 255–272.

Shang KH, Tao Z, Zhou SX (2015) Optimizing reorder intervals for two-echelon distribution systems with stochastic demand. *Oper. Res.* 63(2):458–475.

Simchi-Levi D, Zhao Y (2012) Performance evaluation of stochastic multi-echelon inventory systems: A survey. *Adv. Oper. Res.* 2012: 126254, https://doi.org/10.1155/2012/126254.

Zheng YS (1992) On properties of stochastic inventory systems. *Management Sci.* 38(1):87–103.

Zhu H, Chen F, Hu M, Yang Y (2020) Supplement to "A simple heuristic policy for stochastic distribution inventory systems with fixed shipment costs." Preprint, submitted September 8, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3606084.

**Han Zhu** is an assistant professor in School of Management Sciences and Engineering at Dongbei University of Finance and Economics. His research interests include inventory control, healthcare management, sharing economy, and the operations management–marketing interface.

**Youhua (Frank) Chen** is chair professor and dean of College of Business, City University of Hong Kong. His current research projects cover healthcare management, data analytics in business operations, and applications of machine learning in supply chain management. He is also coordinating a team of registered nurses and social workers to innovate community-based elderly care.

**Ming Hu** is the University of Toronto Distinguished Professor of Business Operations and Analytics and a professor of operations management at Rotman School of Management, University of Toronto. His recent research focus is on operations management in the context of sharing economy, social buying, crowdfunding, crowdsourcing, and two-sided markets, with the goal of exploiting operational decisions to the benefit of the society.

**Yi Yang** is a professor in the School of Management at Zhejiang University. His main research interests include inventory management and data-driven operations management.