

Modified Echelon (r, Q) Policies with Guaranteed Performance Bounds for Stochastic Serial Inventory Systems

Ming Hu

Rotman School of Management, University of Toronto, Toronto, Ontario M5S 3E6, Canada,
ming.hu@rotman.utoronto.ca

Yi Yang

Department of Management Science and Engineering, Zhejiang University, Hangzhou, 310058, China,
yangyizju@zju.edu.cn

We consider the classic continuous-review N stage serial inventory system with a homogeneous Poisson demand arrival process at the most downstream stage (Stage 1). Any *shipment* to each stage, regardless of its size, incurs a *positive* fixed setup cost and takes a *positive* constant lead time. The optimal policy for this system under the long-run average cost criterion is unknown. Finding a good worst-case performance guarantee remains an open problem. We tackle this problem by introducing a class of *modified echelon* (r, Q) policies that do not require Q_{i+1}/Q_i to be a positive integer: Stage $i + 1$ ships to Stage i based on its observation of the echelon inventory position at Stage i ; if it is at or below r_i and Stage $i + 1$ has positive on-hand inventory, then a shipment is sent to Stage i to raise its echelon inventory position to $r_i + Q_i$ as close as possible. We construct a heuristic policy within this class of policies, which has the following features: First, it has provably primitive-dependent performance bounds. In a two-stage system, the performance of the heuristic policy is guaranteed to be within $(1 + K_1/K_2)$ times the optimal cost, where K_1 is the downstream fixed cost and K_2 is the upstream fixed cost. We also provide an alternative performance bound, which depends on efficiently computable optimal (r, Q) solutions to N single-stage systems but tends to be tighter. Second, the heuristic is simple, it is efficiently computable and it performs well numerically; it is even likely to outperform the optimal integer-ratio echelon (r, Q) policies when K_1 is dominated by K_2 . Third, the heuristic is asymptotically optimal when we take some dominant relationships between the setup or holding cost primitives at an upstream stage and its immediate downstream stage to the extreme, for example, when $h_2/h_1 \rightarrow 0$, where h_1 is the downstream holding cost parameter and h_2 is the upstream holding cost parameter.

Subject classifications: multi-echelon; serial system; stochastic demand; performance evaluation; (r, Q) policy.

Area of review: Operations and Supply Chains

History: Received June 2012; revisions received January 2013, June 2013, September 2013, December 2013, March 2014; accepted April 2014. Published online in *Articles in Advance* June 13, 2014.

1. Introduction

The multi-echelon serial inventory system is widespread in supply chains in the retail, manufacturing and service industries. We consider a classic continuous-review N stage serial inventory system. Customer demand arises only at the most downstream stage, Stage 1, and follows a homogeneous Poisson process. Each stage replenishes inventory only from its immediate upstream, i.e., Stage 1 replenishes from Stage 2, etc., and Stage N replenishes from an outside supplier with unlimited stock. Each *shipment* from one stage to its successor, regardless of its size, takes a *positive* constant lead time and triggers a *positive* fixed setup cost. Excess demand that cannot be satisfied immediately is fully backlogged. The system is stationary, and the objective is to minimize the long-run average system-wide cost.

It is well known that the optimal policy of such a system, even if it exists, must be extremely complex (Clark and Scarf

1962). This observation has caused subsequent research on multi-echelon, stochastic inventory systems with economies of scale to focus on heuristic policies. In particular, attention has been focused on a class of *nested integer-ratio echelon* (r, Q) policies (see, e.g., De Bodt and Graves 1985). “Nested integer-ratio” refers to the fact that the batch size of any upstream stage, Q_{i+1} , is equal to a positive integer multiple of the batch size of an immediate downstream stage, Q_i . (“Nested” here is used to emphasize the difference from the integer-ratio policies in Roundy 1985 that could allow Q_i to be a positive integer multiple of Q_{i+1} ; for simplicity, “nested” is omitted hereafter.) The integer-ratio policies have proved to be cost-effective, with a worst-case performance guarantee in some special settings, e.g., when demand is deterministic (Atkins and Sun 1995) and in a two-echelon system with zero lead times at Stage 2 under demand uncertainty (Chen 1999). However, finding a good guaranteed performance bound for a general setting with fixed shipment costs and positive lead

times at each stage remains an unsolved problem. Zipkin (2000) notes the urgency of addressing this problem:

There are ways to calculate a lower bound on the true optimal cost. By comparing this bound to the cost of a good (or the best) echelon (r, q) policy, we can estimate the policy's performance relative to *all* possible alternatives. In this way researchers [e.g., Chen 1999] have found that, for *some* systems, an echelon (r, q) policy performs quite well; the policy's cost exceeds the lower bound by only a few percent. We still lack strong results [...], however. Specifically, we do not have a policy-construction heuristic that is guaranteed to perform well. This remains a pressing priority for research. (Zipkin 2000, §8.3.7.2)

In some special settings, which are mentioned previously, the optimal policies have been shown to be *nested* in the sense that, whenever Stage $i + 1$ orders, so does Stage i . The reason is that, in those special settings, any order at an upstream stage can be delayed until just before the next shipment to its immediate downstream stage so as to save inventory holding costs at the upstream stage. However, this argument no longer holds for a general serial system under demand uncertainty with positive lead times at each stage, because an upstream stage cannot in general predict, ahead of its lead time, exactly when its immediate downstream stage will reach the reorder point. As a result, the optimal policies in a general system may *not* be nested and the need to confine within the class of integer-ratio echelon (r, Q) policies for good heuristics is less compelling.

We introduce a class of continuous-review *modified echelon (r, Q) policies* that do not require Q_{i+1}/Q_i to be a positive integer: Stage $i + 1$ ships to Stage i on the basis of its observation of the echelon inventory position at Stage i ; if the echelon inventory position is at or below r_i and Stage $i + 1$ has positive on-hand inventory, then a shipment is sent to Stage i to raise its echelon inventory position *as close as possible* to $r_i + Q_i$. (Because the proposed policies do not necessarily satisfy integer-ratio constraints, as assumed in the classic echelon (r, Q) models, we coin the term of *modified echelon (r, Q) policy*.)

The inventory flow processes under the modified echelon (r, Q) policy can be completely different from those under the integer-ratio echelon (r, Q) policy, even with the same set of policy parameters. For example, in a two-stage system, suppose $r_1 = 1$, $Q_1 = 5$ and $Q_2 = 10$. Consider the following scenario: at some time when Stage 2 has been out of stock for a while, the inventory position of Stage 1 drops to -1 ; meanwhile, $Q_2 = 10$ units have just arrived at Stage 2. We compare how this scenario triggers different shipment quantities under the two policies. Under the integer-ratio echelon (r, Q) policy, the shipment quantity to Stage 1 is $Q_1 = 5$ units to raise the inventory position of Stage 1 above $r_1 = 1$. However, under the modified echelon (r, Q) policy, the shipment quantity to Stage 1 is 7 units to raise the inventory position of Stage 1 to $r_1 + Q_1 = 6$. Here we note that the shipment size to Stage i under the modified echelon

(r, Q) policy may not be in an integral multiple of Q_i , unlike under the integer-ratio echelon (r, Q) policy.

In view of the possibly complicated inventory flow processes under a given modified echelon (r, Q) policy, it can be extremely hard to express its exact cost. Nevertheless, we show that a constructed heuristic modified echelon (r, Q) policy, specifically designed for the shipment-based fixed-cost scheme (see later for a comparison with the batch-based fixed-cost scheme), has the following features: First, it has provably primitive-dependent performance bounds. Second, it is simple and free from the clustering step required in other heuristics like those in Shang (2008), it can be computed efficiently and it performs well numerically. Third, it is asymptotically optimal when we take some dominant relationships between the setup or holding cost primitives at an upstream stage and its immediate downstream stage to the extreme.

For performance evaluation, we make the crucial assumption that discrete units of inventories can be approximated by continuous quantities. That is a common approach in the inventory literature. For example, the well-known Economic Order Quantity (EOQ) formula is often used, even though the order quantity is usually discrete. Under this assumption, we show guaranteed performance bounds for the constructed modified echelon (r, Q) policy. In a two-stage system, the performance of the heuristic policy is guaranteed to be within $(1 + K_1/K_2)$ times the optimal cost (i.e., $1 + K_1/K_2$ -optimal hereafter), where K_1 is the downstream fixed cost and K_2 is the upstream fixed cost. We also provide an alternative performance bound that depends on efficiently computable optimal (r, Q) solutions to N single-stage systems. Specifically, the effectiveness of the heuristic is shown to be linked to the ratios between Q_{i+1}^* and Q_i^* , $i = 1, 2, \dots, N - 1$, where Q_i^* is the optimal order quantity for Stage i in a series of N decomposed single-stage problems through the *induced-penalty cost-allocation scheme* (see Chen and Zheng 1994b). In a two-stage system for primitives such that $Q_2^*/Q_1^* \geq (1 + 1/\rho - \sqrt{1 + 2/\rho})/2$, we show that the heuristic policy is $(1 + \rho)$ -optimal. For example, if $Q_2^* \geq Q_1^*$, the heuristic is 1.25-optimal. If the upstream stage has a higher setup cost or a lower holding cost than the downstream stage, a situation that will most likely result in a relatively large value of Q_2^*/Q_1^* , the alternative performance bound depending on Q_2^*/Q_1^* tends to be tighter than $(1 + K_1/K_2)$ -optimality.

Because a shipment to Stage i may contain multiple batches of size Q_i , there are at least two schemes in the literature for charging fixed setup costs under echelon (r, Q) policies: a fixed cost per batch (see, e.g., Chen and Zheng 1998, Shang 2008) and a fixed cost per shipment (see, e.g., De Bodt and Graves 1985, Chen and Zheng 1994a). In our setting with a Poisson arrival process, the two schemes can be reconciled under an integer-ratio echelon (r, Q) policy if the system is operated in a centralized way, thus avoiding unnecessary orders at a stage when the immediately upstream stage is out of stock. But under our proposed modified

echelon (r, Q) policy, the two cost schemes lead to different outcomes. We adopt the *shipment-based fixed-cost scheme* that is consistent with the original work by Clark and Scarf (1960). We caution that our theoretical performance bounds do not apply to the batch-based fixed-cost scheme. Heuristics proposed by Chen and Zheng (1998) and Shang (2008) within the class of integer-ratio echelon (r, Q) policies can be more cost-effective for the batch-based fixed-cost scheme, because integer-ratio policies can completely avoid any partial batch that would be charged with a setup cost as a full batch under the batch-based cost scheme.

Our contribution to the literature on stochastic serial inventory systems is fourfold. First, we introduce the class of modified echelon (r, Q) policies that do not require Q_{i+1}/Q_i to be a positive integer. Second, for *any* arbitrarily given modified echelon (r, Q) policy, we establish an upper bound on its system-wide long-run average cost. This upper bound is not only easily computable, but is also convenient for performance bound analysis. Third, we evaluate the performance of a constructed modified echelon (r, Q) heuristic by identifying the gap between the established upper bound and the induced-penalty lower bound on the optimal cost shown in Chen and Zheng (1994b). Lastly, we provide asymptotic optimality results on modified echelon (r, Q) policies.

Literature Review

The multi-echelon serial inventory system with stochastic demand has been studied extensively because of the seminal work of Clark and Scarf (1960). In a finite-horizon setting *without* setup costs, the authors show that an echelon base-stock policy is optimal for each stage. Federgruen and Zipkin (1984) extend this result to the infinite-horizon setting. Shang and Song (2003) and Gallego and Özer (2005) develop easily implemented heuristics for the optimal echelon base-stock levels. Gallego and Özer (2003) prove the optimality of state-dependent, echelon base-stock policies for finite and infinite horizon problems with advance demand information, and show that myopic policies are optimal when costs and demands are stationary. Dong and Lee (2003) show that the echelon base-stock policy is also optimal under time-correlated demand processes by using a martingale model of forecast evolution. Chen (2000a) considers batch-ordering constraints and establishes the optimality of an echelon stock (R, nQ) policy. More recently, van Houtum et al. (2007) consider fixed replenishment intervals and show that the echelon base-stock policy is still optimal; Chao and Zhou (2009) further take batch ordering into account. However, all of these results rely on the assumption that *no* explicit setup costs are incurred at any stage except the most upstream stage.

For a serial inventory system with stochastic demand, Clark and Scarf (1962) point out that the presence of setup costs, especially at the downstream stages, results in extreme complexity in the optimal policy. As a result, most of the growing research on stochastic serial inventory systems with economies of scale, focuses on simple, cost-effective heuristic

policies, see, e.g., Graves (1985), Moynzadeh and Lee (1986), Jackson (1988), Svoronos and Zipkin (1988, 1991), Axsäter (1990, 1993), Chen and Zheng (1994a, 1998), Shang (2008), Shang and Zhou (2010), and Yang et al. (2011). To evaluate the performance of a specific heuristic policy, the first step is to provide the exact measure or bounds on the policy's system-wide cost. To the best of our knowledge, De Bodt and Graves (1985) are the first to study the echelon-stock (r, nQ) policy (i.e., the integer-ratio echelon (r, Q) policy) in stochastic serial inventory systems with setup costs at all stages, by analyzing an approximate-cost model. Chen and Zheng (1994a) provide a recursive procedure for computing various performance measures of echelon-stock (r, nQ) policy. For convenience of implementation and coordination across stages, the two papers focus on integer-ratio policies. Yet there is still a lack of rigorous argument showing that the optimal policy for the general serial system with stochastic demand has to be an integer-ratio policy or has to satisfy the nestedness property (i.e., whenever Stage $i + 1$ orders, so does Stage i). As observed from the two papers, even for the integer-ratio policies, the system-wide cost expressions are too complex to be used further for performance bound analysis. We provide an easily computable upper bound on the system-wide cost under any given modified echelon (r, Q) policy, that can be used for further performance bound analysis.

Although most of the heuristic policies proposed in the literature make intuitive sense, it is unclear how wide the gap is between a heuristic and an optimal policy. For a deterministic multi-echelon inventory system, it is well known that the so-called power-of-two policies perform surprisingly well. Roundy (1985) shows that for the one-warehouse multi-retailer distribution system, the performance gap between an optimal power-of-two policy and the optimal cost is guaranteed to be within 2% (98%-effective hereafter). Atkins and Sun (1995) demonstrate that the effectiveness of power-of-two policies does not change in the serial system. More recent works that analyze the heuristic policies with worst-case bounds in deterministic systems include Chen (2000b), Chan et al. (2002) and Levi et al. (2008a). In contrast, fewer papers are studying the worst-case bound of a heuristic policy in the setting with stochastic demand. Chu and Shen (2010) study a two-echelon distribution system under demand uncertainty. With a given set of target service levels, the authors show that a power-of-two policy is guaranteed to be 1.26-optimal. Works closely related to ours are Chen (1999) and Levi et al. (2006). The former proposes a simple and 94%-effective heuristic policy for the two-stage serial inventory system under a restrictive assumption that Stage 2 has zero lead times. The latter provides a 3-optimal heuristic policy for a multi-echelon serial system with correlated, non-stationary demands over a finite horizon. The applied approach was originally proposed by Levi et al. (2008b) that consider a capacitated single-location stochastic inventory system. Other than these two papers, we are not aware of other works on performance bound analysis

of heuristic policies for a stochastic multi-echelon serial inventory system.

The rest of the paper is organized as follows. In §2, we describe the model in detail. For the main body of the paper, we focus on the two-stage system; the results for the multi-stage system are in the Online Appendix (available as supplemental material at <http://dx.doi.org/10.1287/opre.2014.1291>). In §3, we build an upper bound on the system-wide cost of any given modified echelon (r, Q) policy. On the basis of that upper bound, we provide performance bounds of a constructed modified echelon (r, Q) policy. We also explore the asymptotic optimality of the constructed heuristic with respect to system cost primitives. In §4, we present numerical results. We finish with §5, in which we discuss extensions and limitations. Throughout the paper, \mathbb{N} is defined as the positive integer set, and \mathbb{E} as the expectation operator. All proofs are given in the appendix.

2. Model and Preliminaries

In this section, we first describe our model in detail and then review some relevant results of Zheng (1992) and Chen and Zheng (1994b), that will be used in our subsequent analysis.

2.1. Notation and Formulation

As mentioned previously, we consider a firm that manages an N stage *continuous-review* serial inventory system, where Stage 1 is replenished from Stage 2, etc., and Stage N in turn is replenished from an outside supplier (denoted by Stage $N + 1$) with *unlimited* stock. External customer demand arrives only at Stage 1 according to a Poisson process with a constant rate λ . We denote by $D(t, t + \tau]$ the total demand over the time interval $(t, t + \tau]$. The transportation lead times from Stage $i + 1$, $i = 1, 2, \dots, N$, to its successor, Stage i , are constant $L_i > 0$. In other words, any shipment sent out by Stage $i + 1$ at time t will be received by Stage i at time $t + L_i$. Each *shipment* from Stage $i + 1$ to Stage i incurs a fixed setup cost $K_i > 0$. Without loss of generality, we assume the variable order cost to be zero. Whenever Stage 1 runs out of stock, the unmet demand is fully backlogged with a backlog cost rate $p > 0$. Let $h_i > 0$ be the echelon holding cost rate at Stage i . The firm's objective is to determine a shipment policy that minimizes the long-run average system-wide cost.

We proceed to review some concepts that are commonly used in the literature. *Echelon inventory* at Stage i is the inventory on hand at Stage i plus inventories at or in transit to all its downstream stages. (Note that for Stage 1, echelon inventory is merely its on-hand inventory.) *Echelon inventory level* at Stage i is the echelon inventory at Stage i minus the total number of customers back-ordered at Stage 1. *Echelon inventory position* at Stage i is the sum of echelon inventory level at Stage i and the inventories in transit to Stage i . For $i = 1, 2, \dots, N$, define the following inventory variables at time t :

$I_i(t)$ = echelon inventory at Stage i ,
 $IL_i(t)$ = echelon inventory level at Stage i ,
 $IP_i^-(t)$ = echelon inventory position before a shipment is sent to Stage i at time t ,

$IP_i(t)$ = echelon inventory position after a shipment is sent to Stage i at time t ,
 $q_i(t)$ = shipping quantity to Stage i ,
 $B(t)$ = backorder inventory level at Stage 1,
 $OI_i(t)$ = on-hand inventory at Stage i before a shipment is sent to its successor at time t .

The system dynamics are as follows:

$$I_i(t) = IL_i(t) + B(t),$$

$$IL_i(t + L_i) = IP_i(t) - D(t, t + L_i], \quad (1)$$

$$IP_i^-(t) = IL_{i+1}(t) - OI_{i+1}(t), \quad (2)$$

$$IP_i(t) = IP_i^-(t) + q_i(t). \quad (3)$$

The first equation follows directly from the definitions. The second equation follows from the fact that the inventory position of Stage i at time t can be decomposed into two components: the echelon inventory level of Stage i at time $t + L_i$ and the total demand over $(t, t + L_i]$. The third equation follows from the fact that the echelon inventory level at Stage $i + 1$ is the sum of the on-hand inventory at Stage $i + 1$ and the inventory position at Stage i before a replenishment is sent. The last equation specifies how an ordering decision changes the inventory position.

At any time t , the shipping quantity to Stage i , $q_i(t)$, is the decision variable. Equivalently, by Equation (3), we can think of the inventory position $IP_i(t)$ as the alternative decision variable. Note that the shipping quantity $q_i(t)$ to Stage i , $i = 1, 2, \dots, N - 1$, should be non-negative and less than the on-hand inventory level at its upstream Stage $i + 1$, $OI_{i+1}(t)$. That is, $0 \leq q_i(t) \leq OI_{i+1}(t)$. By Equations (2) and (3), the inventory position at Stage i , $i = 1, 2, \dots, N - 1$, should satisfy the following constraint: $IP_i^-(t) \leq IP_i(t) \leq IL_{i+1}(t)$. Because of the unlimited stock at the outside supplier, the inventory position at Stage N only needs to satisfy the constraint $IP_N^-(t) \leq IP_N(t)$.

The total controllable long-run average cost includes the fixed setup cost, inventory holding cost and backlog cost. The fixed setup cost incurred for Stage i at time t is $K_i \delta(q_i(t) > 0) = K_i \delta(IP_i(t) > IP_i^-(t))$, where $\delta(\cdot)$ is the indicator function. The system-wide inventory holding and backlog cost rate at time t is given by

$$\sum_{i=1}^N h_i I_i(t) + pB(t) = \sum_{i=1}^N h_i IL_i(t) + (p + H_1)B(t), \quad (4)$$

where $H_1 \equiv \sum_{i=1}^N h_i$ represents the installation holding cost rate at Stage 1. Then the original problem can be formulated as follows:

$$(\mathcal{B}): \quad C_{\mathcal{B}}^* \equiv \min_{\{IP_i(t), \forall i, t\}} \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E} \left[\int_{t=0}^T \sum_{i=1}^N K_i \delta(IP_i(t) > IP_i^-(t)) + h_i IL_i(t) + (p + H_1)B(t) dt \right]$$

s.t. $IP_i^-(t) \leq IP_i(t) \leq IL_{i+1}(t), i = 1, 2, \dots, N - 1$, and $IP_N^-(t) \leq IP_N(t)$.

It is well known that if a positive setup cost is incurred for each shipment at the downstream stages, the optimal policy, if one exists, must be extremely complex. In this paper, we focus on a class of modified echelon (r, Q) policies.

DEFINITION 1 (MODIFIED ECHELON (\mathbf{r}, \mathbf{Q}) POLICY). Stage $i + 1$ ships to Stage i on the basis of its observation of the echelon inventory position at Stage i . In particular, if the echelon inventory position at Stage i is at or below r_i and Stage $i + 1$ has positive on-hand inventory, then a shipment is sent to Stage i to raise its echelon inventory position as close as possible to $r_i + Q_i$.

2.2. Relevant Existing Results

2.2.1. Single-Stage Model. Zheng (1992) studies a continuous-review, single-stage inventory model with positive lead time and a fixed setup cost. Let L be the shipment lead time and $G(y)$ the expected cost rate accruing at time $t + L$, when the inventory position equals y at time t .

ASSUMPTION 1 (REGULARITY). (i) $G(y)$ is a convex function and $\lim_{y \rightarrow \pm\infty} G(y) = \infty$.

(ii) There exist $a > 0$ and $b < 0$, such that $\lim_{y \rightarrow \infty} G'(y) = a$ and $\lim_{y \rightarrow -\infty} G'(y) = b$.

Assume that each shipment incurs a fixed setup cost K . For such a single-stage problem, it is well known that the optimal control policy is an (r, Q) policy (Zheng 1992), under which the inventory position in steady state is uniformly distributed on $\{r + 1, \dots, r + Q\}$ and is independent of the lead-time demand (Zipkin 2000, Theorem 6.2.3). Then the long-run average cost function under any (r, Q) policy should take the following form: $C(r, Q) = [\lambda K + \sum_{y=r+1}^{r+Q} G(y)]/Q$. We adopt the convention that discrete units of inventories can be approximated by continuous variables. As mentioned in the Introduction, such an approximation is widely used in the inventory literature (see, e.g., Chen 1999, Shang et al. 2009; see also Zheng 1992 for more discussion). We will follow that convention in the rest of this paper. Consequently, the long-run average cost of the single-stage problem can be approximated as follows:

$$C(r, Q) = \frac{\lambda K + \int_r^{r+Q} G(y) dy}{Q}. \tag{5}$$

Clearly, the approximation is adequate when Q is large enough. The objective is to determine the values of r and Q that minimize $C(r, Q)$. For any fixed Q , define $r(Q) \equiv \arg \min_r C(r, Q)$. If the optimal solution is not unique, we choose the largest one; this convention will be used throughout the paper. Also, define $C(Q) \equiv C(r(Q), Q)$, $(r^*, Q^*) \equiv \arg \min_{r, Q} C(r, Q)$ and $C^* \equiv C(r^*, Q^*)$. Finally, let $H(Q) \equiv G(r(Q))$ and $A(Q) \equiv QH(Q) - \int_0^Q H(y) dy$.

LEMMA 1 (ZHENG 1992). Under Assumption 1, the following results hold: (i) $C(r, Q)$ is jointly convex in r and Q . (ii) $G(r(Q)) = G(r(Q) + Q)$ and $G(r^*) = G(r^* + Q^*) = C^*$.

- (iii) $C(Q)/C^* \leq \epsilon(Q/Q^*)$, where $\epsilon(q) = (q + q^{-1})/2$.
- (iv) $H(Q)$ is an increasing convex function, $\int_0^Q H(y) dy \geq \frac{1}{2}QH(Q)$ and $H(Q^*) = C^*$.
- (v) $A(Q)$ is an increasing function and $A(Q^*) = \lambda K$.

By Lemma 1, we further obtain the following properties.

LEMMA 2. (i) For any $y \in [r^*, r^* + Q^*]$, $G(y) \leq C^*$. (ii) $\lambda K \leq \frac{1}{2}C^*Q^*$.

2.2.2. Lower Bound for Multi-Stage Model. We adopt the so-called *induced-penalty bound* on the optimal system-wide cost $C_{\mathcal{B}}^*$, that has been established by Chen and Zheng (1994b). The authors decompose the original system into N single-stage systems linked by the induced-penalty cost functions, which are essentially some penalty costs charged to an upstream stage if it cannot immediately fulfill an order request from its successive stage.

Let D_i denote the total demand over $(0, L_i]$. Define

$$G_1(y) \equiv \mathbb{E}[h_1(y - D_1) + (p + H_1)(y - D_1)^-],$$

$$C_1(r_1, Q_1) \equiv \frac{1}{Q_1} \left[\lambda K_1 + \int_{r_1}^{r_1+Q_1} G_1(y) dy \right], \tag{6}$$

where $x^- \equiv \max\{-x, 0\}$. Then we can define recursively as follows: for $i = 2, \dots, N$,

$$\bar{G}_{i-1}(y) \equiv \begin{cases} G_{i-1}(y) - C_{i-1}^* & \text{if } y \leq r_{i-1}^*, \\ 0 & \text{otherwise,} \end{cases}$$

$$G_i(y) \equiv \mathbb{E}[h_i(y - D_i) + \bar{G}_{i-1}(y - D_i)],$$

$$C_i(r_i, Q_i) \equiv \frac{1}{Q_i} \left[\lambda K_i + \int_{r_i}^{r_i+Q_i} G_i(y) dy \right], \tag{7}$$

where C_i^* is the minimum of $C_i(r_i, Q_i)$ achieved at (r_i^*, Q_i^*) , i.e., $(r_i^*, Q_i^*) \equiv \arg \min_{r, Q} C_i(r, Q)$ and $C_i^* \equiv C_i(r_i^*, Q_i^*)$. Because it is easily verified that $G_i(y)$ and $C_i(r, Q)$ are convex, the traditional optimization methods can efficiently compute those critical points, r_i^* and Q_i^* , that will be used to construct our heuristics. Chen and Zheng (1994b) show that the sum of the optimal costs from these single-stage problems is a lower bound on $C_{\mathcal{B}}^*$.

LEMMA 3 (INDUCED-PENALTY LOWER BOUND; LEMMA 2 AND SECTION 5 OF CHEN AND ZHENG 1994B). (i) For any $i = 1, 2, \dots, N$, $G_i(y)$ satisfies Assumption 1 with $a = h_i$ and $b = -(\sum_{j=i+1}^N h_j + p)$, if $i < N$; $a = h_i$ and $b = -p$, if $i = N$. In addition, $C_i(r_i, Q_i)$ is a joint convex function of r_i and Q_i .

(ii) $\sum_{i=1}^N C_i^* \leq C_{\mathcal{B}}^*$.

3. Two-Stage Model

Consider that the two-stage system operates under an arbitrarily given modified echelon (\mathbf{r}, \mathbf{Q}) policy, where $\mathbf{r} = (r_1, r_2) \in \mathbb{N}^2$ and $\mathbf{Q} = (Q_1, Q_2) \in \mathbb{N}^2$. Note that the size of the shipment to Stage 1 is physically constrained by the on-hand inventory of Stage 2. Therefore, Stage 1 may not

be able to raise its echelon inventory position to the desired level, $r_1 + Q_1$. In that case, under the modified echelon (\mathbf{r}, \mathbf{Q}) policy, Stage 2 will feed Stage 1 as much as possible, and the size of the shipment to Stage 1 can be larger or smaller than Q_1 . For instance, if Stage 2 does not have enough on-hand inventory when Stage 1 reaches the reorder point, the shipment can be smaller than Q_1 ; on the other hand, if Stage 2 has run out of stock for a while, the size of the shipment to Stage 1 after Stage 2 is back in stock can be larger than Q_1 . We denote by $C(r_1, Q_1, r_2, Q_2)$, or $C(\mathbf{r}, \mathbf{Q})$, the long-run average system-wide cost under the modified echelon (\mathbf{r}, \mathbf{Q}) policy. It is clear that such a policy, as a feasible solution to Problem (\mathcal{B}), provides a cost upper bound on $C_{\mathcal{B}}^*$.

OBSERVATION 1. For any \mathbf{r} and \mathbf{Q} , $C_{\mathcal{B}}^* \leq C(\mathbf{r}, \mathbf{Q})$.

Our main objective is to evaluate the effectiveness of a given modified echelon (\mathbf{r}, \mathbf{Q}) policy. Specifically, we intend to identify the relative gap between the long-run average cost under a given modified echelon (\mathbf{r}, \mathbf{Q}) policy, $C(\mathbf{r}, \mathbf{Q})$, and the optimal cost $C_{\mathcal{B}}^*$, which can be expressed as $RG(\mathbf{r}, \mathbf{Q}) \equiv (C(\mathbf{r}, \mathbf{Q}) - C_{\mathcal{B}}^*)/C_{\mathcal{B}}^*$. As mentioned previously, the optimal cost $C_{\mathcal{B}}^*$ is difficult to be calculated. Thus, we replace the optimal cost by the induced-penalty lower bound presented in Lemma 3(ii) to obtain an upper bound on the relative gap: $RG(\mathbf{r}, \mathbf{Q}) \leq (C(\mathbf{r}, \mathbf{Q}) - \sum_{i=1}^N C_i^*)/\sum_{i=1}^N C_i^*$. The next step is to provide an upper bound on $C(\mathbf{r}, \mathbf{Q})$ itself.

3.1. An Upper Bound

A critical criterion on a good upper bound for the purpose of performance bound analysis is whether the upper bound is convenient for comparison with the lower bound. Next we will propose a novel approach to obtaining an upper bound on $C(\mathbf{r}, \mathbf{Q})$ that is convenient for comparison with the induced-penalty lower bound $\sum_{i=1}^N C_i^*$. Like Chen and Zheng (1994b), we will use the following cost-accounting scheme.

DEFINITION 2 (COST ACCOUNTING SCHEME). At time t , we charge the inventory holding cost of Stage 2 incurred at time $t + L_2$, and charge the inventory holding and backlog costs of Stage 1 incurred at time $t + \sum_{i=1}^2 L_i$.

This cost accounting scheme only shifts costs across time points and hence, does not affect the long-run average inventory holding and backlog costs. It has an intuitive interpretation: an order placed by Stage 2 at time t does not have an effect on the inventory holding cost of Stage 2 until time $t + L_2$, and does not have an effect on the inventory holding and backlog cost of Stage 1 until time $t + \sum_{i=1}^2 L_i$ or later.

3.1.1. Outline of the Approach. Without the integer-ratio assumption, it is extremely hard to identify the *exact* regeneration points of the system operated under a given modified echelon (\mathbf{r}, \mathbf{Q}) policy. Instead, we track the *first* shipment time of an order placed by Stage 2 and being

sent to Stage 1. These time points are not necessarily the regeneration points of the system, because the system states (for example, the inventory positions of all stages) at these epochs are not necessarily the same. Nevertheless, these time epochs provide a structured way to analyze and bound the total operating costs. The time interval between two consecutive such time epochs can be *loosely* called a “cycle” for Stage 1.

Our approach is as follows. Because Stage 2 has an unlimited outside supplier, its expected cost rate can be accounted for by the standard approach (see, e.g., Zheng 1992). With Stage 1, in contrast, that is much more difficult. To tackle the difficulty, we first attempt to build an upper bound on Stage 1’s expected cost rate in terms of $IL_2(t)$. This involves categorizing “cycles” at Stage 1 into *regular* and *irregular* shipment periods. In a regular shipment period, Stage 1’s inventory position gradually decreases from $r_1 + Q_1$ to r_1 unit by unit. Because of such regularity, the expected cost rate can again be handled by the standard approach. We treat irregular shipment periods by considering their setup costs, and inventory holding and backlog costs separately. Bounding of the irregular shipment period’s setup costs is done by bounding the frequency of irregular shipment periods (see Observation 4). On the other hand, an upper bound on the irregular shipment period’s inventory holding and backlog costs can be expressed in terms of $IL_2(t)$ (see Lemma 6). Combining all costs at both stages, we can bootstrap an upper bound on $C(\mathbf{r}, \mathbf{Q})$ from $IL_2(t)$, whose distribution, by Equation (1), is linked to the uniform distribution of the steady-state echelon inventory position of Stage 2, $IP_2(t)$.

3.1.2. Cycles. Because Stage 2 is replenished from an incapacitated outside supplier, Stage 2 can always successfully raise its echelon inventory position from r_2 to $r_2 + Q_2$ after placing an order of size Q_2 , under a given modified echelon (\mathbf{r}, \mathbf{Q}) policy. For convenience of analysis, we view the time period between two consecutive orders of Stage 2, with Q_2 flow units in each order, as a “cycle.” To track how shipments are dispatched, flow units are ranked in chronological order of shipments, with an arbitrary tie-breaking rule within the same shipment. Without loss of generality, units are assumed to flow between stages in a *first-in-first-out* manner. For each stage, we define a “cycle” rigorously as follows.

DEFINITION 3 (CYCLE). For any Stage i , $i = 1, 2$, we call $[T_i^j, T_i^{j+1})$, for any $j \in \mathbb{N}$, the j th cycle, where T_i^j is the time epoch of the 1st unit, contained in the j th order of Stage 2, being sent to Stage i .

Because the modified echelon (\mathbf{r}, \mathbf{Q}) policy is not necessarily nested, there may not be an accompanying shipment to Stage 1 for each order arriving at Stage 2. As a result, the placing of an order from the outside supplier at time epoch T_2^j that triggers an immediate shipment of Q_2 units to Stage 2 may not lead to an immediate cross-dock shipment of Q_1 units to Stage 1 at time $T_2^j + L_2$. That is, T_1^j may be

Table 1. Illustration of a modified echelon (\mathbf{r}, \mathbf{Q}) policy.

Time t	Realized demand	Inventory levels at Stage 2 {[inventory in transit], [on-hand inventory]}	Inventory levels at Stage 1 {[inventory in transit], [on-hand inventory]}
$t = 0$		$\Rightarrow^2 \{(1, 1), \dots, (1, 7)\}, \text{zero}$	{zero, two}
$t = 1$	$D(0, 1] = 4$	{zero, [(1, 7)]}	$\Rightarrow^1 \{(1, 1), \dots, (1, 6)\}, \text{two backlogged}$
$t = 2$	$D(1, 2] = 2$	{zero, [(1, 7)]}	{zero, [(1, 5), (1, 6)]}
$t = 3$	$D(2, 3] = 1$	$\Rightarrow^2 \{(2, 1), \dots, (2, 7)\}, [(1, 7)]$	{zero, [(1, 6)]}
$t = 3.5$	$D(3, 3.5] = 1$	{[(2, 1), \dots, (2, 7)], zero}	$\Rightarrow^4 \{(1, 7)\}, \text{zero}$
$t = 4.5$	$D(3.5, 4.5] = 0$	{zero, [(2, 1), \dots, (2, 7)]}	{zero, [(1, 7)]}
$t = 5$	$D(4.5, 5] = 1$	{zero, [(2, 5), (2, 6), (2, 7)]}	$\Rightarrow^1 \{(2, 1), \dots, (2, 4)\}, \text{zero}$
$t = 6$	$D(5, 6] = 4$	{zero, zero}	$\Rightarrow^1 \{(2, 5), (2, 6), (2, 7)\}, \text{zero}$
$t = 7$	$D(6, 7] = 1$	$\Rightarrow^2 \{(3, 1), \dots, (3, 7)\}, \text{zero}$	{zero, [(2, 6), (2, 7)]}
$t = 8$	$D(6, 7] = 2$	{zero, [(3, 5), (3, 6), (3, 7)]}	$\Rightarrow^1 \{(3, 1), \dots, (3, 4)\}, \text{zero}$

Notes. $(r_1, r_2, Q_1, Q_2) = (0, 2, 4, 7)$ and $L_1 = L_2 = 1$. The sign \Rightarrow^i indicates that there is a shipment on the way to Stage i , and (j, k) is the k th unit in the j th order of Stage 2.

strictly greater than $T_2^j + L_2$. Moreover, for Stage 1, it is likely that a cycle $[T_1^j, T_1^{j+1})$ is an empty set. In that case, the entire j th order of Stage 2 is shipped to Stage 1 together with the 1st shipment of the $(j + 1)$ th order of Stage 2. Next we present an example to illustrate the definition of cycles. To facilitate discussion, we first index each flow unit by the time of its shipment from the outside supplier. Specifically, we denote, by Unit (j, k) , the k th unit in the j th order of Stage 2 to be released into the serial system by the outside supplier.

EXAMPLE 1 (CYCLES). Consider a modified echelon (\mathbf{r}, \mathbf{Q}) policy with $(r_1, r_2, Q_1, Q_2) = (0, 2, 4, 7)$ and $L_1 = L_2 = 1$. Table 1 illustrates the unit flow between stages, given a stream of realized demand. For simplicity, the last unit of demand in any listed interval of time is assumed to arrive exactly on the integral time epoch; e.g., a customer arrives exactly at $t = 0$, leaving the on-hand inventory of Stage 1 at two units, the 4th unit of demand in time interval $(0, 1]$ arrives exactly at $t = 1$, and the unit of demand in time interval $(3, 3.5]$ arrives exactly at $t = 3.5$. We first show how the modified echelon (\mathbf{r}, \mathbf{Q}) policy is implemented. For instance, when the second customer in time period $[0, 1]$ arrives, Stage 1’s inventory position drops to $r_1 = 0$. However, at that time, there is no on-hand inventory at Stage 2 and hence, no shipment will be triggered until time $t = 1$, when a shipment arrives at Stage 2. Because $IP_1^-(t = 1) = -2$, a shipment with a size of six units from Stage 2 to Stage 1 is triggered so as to increase the $IP_1(t = 1)$ to the desirable level $r_1 + Q_1 = 4$. Or, for another example, at time $t = 3.5$, a customer arrival causes $IP_1^-(t = 3.5)$ to hit $r_1 = 0$ and hence triggers a shipment from Stage 2 that contains all its on-hand inventory (only one unit in this case), although, unfortunately, Stage 2 does not have enough on-hand stock to raise $IP_1(t = 3.5)$ to the desired level of $r_1 + Q_1 = 4$. Note that the shipments to Stage 1 can be larger or smaller than Q_1 under modified echelon (r_1, Q_1) policy; see, e.g., $t = 1$ and $t = 3.5$. By definition, T_i^j is the time epoch at which Unit $(j, 1)$ is shipped to Stage i . Along the sample path presented in Table 1, for Stage 2, $T_2^1 = 0$, $T_2^2 = 3$, and


$T_2^3 = 7$, whereas for Stage 1, $T_1^1 = 1$, $T_1^2 = 5$ and $T_1^3 = 8$. Thus, $[0, 3)$ and $[3, 7)$ are the cycles of Stage 2, whereas $[1, 5)$ and $[5, 8)$ are the cycles of Stage 1. \square

3.1.3. Regular and Irregular Shipments. Because of the unlimited outside supply, cycles of Stage 2 are “standard” under any echelon (r_2, Q_2) policy, as in a single-stage system. However, Stage 1 is different. We zoom into each cycle of Stage 1 to investigate possible shipment behavior. Our analysis concentrates on non-empty cycles because Stage 1 incurs no costs for empty cycles. Over any non-empty cycle $[T_1^j, T_1^{j+1})$, the j th order of Stage 2 will be shipped to Stage 1 in one or more, say $M \in \mathbb{N}$, shipments in total. Let $T_1^{j,l}$ be the time of the l th shipment over $[T_1^j, T_1^{j+1})$, where $l = 1, 2, \dots, M$. By definition, $T_1^{j,1} \equiv T_1^j$ and we have $[T_1^j, T_1^{j+1}) = \bigcup_{l=1}^M [T_1^{j,l}, T_1^{j,l+1})$, where $T_1^{j,M+1} \equiv T_1^{j+1}$. We call $[T_1^{j,l}, T_1^{j,l+1})$ the l th shipment period over the cycle $[T_1^j, T_1^{j+1})$. Depending on the inventory positions of Stage 1 at the beginning and the end of a shipment period, we categorize Stage 1’s shipment periods into the following two types.

DEFINITION 4 (REGULAR AND IRREGULAR SHIPMENT PERIOD). For any shipment period $[T_1^{j,l}, T_1^{j,l+1})$, $l = 1, \dots, M$, if $IP_1(T_1^{j,l}) = r_1 + Q_1$ and $IP_1^-(T_1^{j,l+1}) = r_1$, we call it a regular shipment period; otherwise, we call it an irregular shipment period.

EXAMPLE 2 (SHIPMENT PERIODS). We revisit Example 1 to illustrate the concept of shipment periods. Over the cycle $[1, 5)$ of Stage 1, two shipments are sent from Stage 2, at time $t = 1$ and $t = 3.5$. That is, $T_1^{1,1} = 1$ and $T_1^{1,2} = 3.5$. By definition, $T_1^{1,3} \equiv T_1^2 = 5$. Then this cycle is a union of two shipment periods, i.e., $[1, 5) = [1, 3.5) \cup [3.5, 5)$. Because $IP_1(t = 1) = 4 = r_1 + Q_1$ and $IP_1^-(t = 3.5) = 0 = r_1$, the first shipment period $[1, 3.5)$, is a regular shipment period. However, because $IP_1(t = 3.5) = 1 < 4 = r_1 + Q_1$, the second shipment period $[3.5, 5)$, is an irregular shipment period. Moreover, Because $IP_1^-(t = 1) = -2$, $t = 1$ is the ending point of an earlier irregular shipment period. Note that the situations where $t = 3.5$ and $t = 1$ are the two types of

Figure 1. Illustration of Stage 1’s inventory position over a cycle.



instances we imply in the definition of an irregular shipment period, namely, the starting inventory position of a shipment period $IP_1(T_1^{j,l}) \neq r_1 + Q_1$, and the ending inventory position of a shipment period $IP_1^-(T_1^{j,l+1}) \neq r_1$. \square

For some modified echelon (\mathbf{r}, \mathbf{Q}) policies, we may further characterize their cycles and shipment periods. For example, consider a special case where $Q_2 < Q_1$ and $r_2 + Q_2 \leq r_1$. In this case, Stage 1’s echelon inventory position will never reach $r_1 + Q_1$, and any shipment to Stage 2 will be immediately sent to Stage 1 at the moment it arrives at Stage 2. That is, $T_1^j = T_2^j + L_2$. Then, the cycles of Stage 1 are defined by $[T_2^j + L_2, T_2^{j+1} + L_2]$. Over any of these cycles, there is only one shipment that is an irregular shipment, and hence the cycle itself is an irregular shipment period.

In general, Figure 1 illustrates how Stage 1’s inventory position may change over a cycle. We show that the first $M - 1$ (possibly zero) shipment periods in any cycle have to be regular. Intuitively, that is because at the beginning of those shipment periods, there is enough on-hand inventory at Stage 2 to raise Stage 1’s echelon inventory position IP_1 to $r_1 + Q_1$. However, the last shipment period in a cycle may not be regular, with the following two scenarios that are not mutually exclusive. First, at the beginning of the last shipment period, Stage 2 does not have enough inventory to raise Stage 1’s inventory position to $r_1 + Q_1$; i.e., $IP_1(T_1^{j,M}) < r_1 + Q_1$. Second, at the end of the last shipment period, Stage 2 is out of stock and Stage 1’s inventory position is below r_1 ; i.e., $IP_1^-(T_1^{j,M+1}) < r_1$.

LEMMA 4 (IRREGULAR SHIPMENT FREQUENCY). *There exists at most one irregular shipment period in each cycle of Stage 1, whether the cycle is empty or not.*

3.1.4. Cost Assessment in Cycles. We evaluate and bound the expected cost rate for the two types of shipment periods. First we consider regular shipment periods. Over each regular shipment period $[T_1^{j,l}, T_1^{j,l+1})$, $l = 1, 2, \dots, M - 1$, Stage 1’s inventory position gradually

drops from $r_1 + Q_1$ to r_1 , at which another shipment is triggered and the regular shipment period ends. On the condition that Stage 1 is in a regular shipment period, the expected cost rate at Stage 1 should be the same as that in the single-stage problem when Stage 1 has access to an outside supplier with unlimited supply (Zheng 1992). See also the expected cost-rate expression of Stage 1 in Chen (1999), where Stage 1 is always in a regular shipment period under any integer-ratio (r, Q) policy.

OBSERVATION 2 (COST IN REGULAR AND IRREGULAR SHIPMENT PERIOD). (i) The expected cost rate at Stage 1 in a regular shipment period is $C_1(r_1, Q_1) = (1/Q_1)[\lambda K_1 + \int_{r_1}^{r_1+Q_1} G_1(y)dy]$, where $G_1(y)$ is defined in (6). (ii) For any time t in a non-empty irregular shipment period $[T_1^{j,M}, T_1^{j+1})$, Stage 1’s expected inventory holding and backlog cost accrues at a rate equal to $G_1(IP_1(t))$; furthermore, there is a fixed setup cost, K_1 , incurred for the irregular shipment.

3.1.5. Cost Upper Bound. Lemma 4 implies that the setup cost, K_1 , for a (possible) irregular shipment period is incurred at most once in any cycle $[T_1^j, T_1^{j+1})$. Consequently, we have the following result.

OBSERVATION 3. For any time $t \in [T_1^j, T_1^{j+1}) \neq \emptyset$, setup costs for irregular shipment periods accrue at a rate that is no more than $K_1/(T_1^{j+1} - T_1^j)$.

It is necessary to calculate the long-run average expected “cycle” length to assess the long-run average setup costs for irregular shipment periods. The following lemma provides an exact value for this quantity that is equal to the ratio between the batch size Q_2 and the demand rate λ , independent of inventory levels at any stage.

LEMMA 5 (CYCLE LENGTH). *Under any echelon (\mathbf{r}, \mathbf{Q}) policy, for each Stage $i = 1, 2$, the long-run average expected cycle length $\lim_{j \rightarrow \infty} E[(T_1^{j+1} - T_1^j)]/j = Q_2/\lambda$.*

Combining Observation 3 and Lemma 5, we immediately have an upper bound on the long-run average setup cost for irregular shipment periods, as follows.

OBSERVATION 4. The long-run average setup cost for irregular shipment periods has an upper bound $\lambda K_1/Q_2$.

Next we bound the expected cost rate at Stage 1, excluding setup costs for irregular shipment periods. Though the upper bound is provided in terms of Stage 2’s echelon inventory level $IL_2(t)$, we will eventually bound the total costs of both stages in terms of $IP_2(t)$, based on the relationship $IL_2(t + L_2) = IP_2(t) - D(t, t + L_2)$. This facilitates our analysis because $IP_2(t)$ is distributed uniformly.

LEMMA 6. *For any time $t \in [T_1^j, T_1^{j+1}) \neq \emptyset$, the expected cost rate at Stage 1 excluding an irregular shipment period’s setup costs, denoted by $\hat{\Gamma}_1(IL_2(t))$, can be bounded as follows:*

$$\hat{\Gamma}_1(IL_2(t)) \leq \begin{cases} G_1(IL_2(t)) & \text{if } IL_2(t) \leq r_1, \\ \max\{G_1(\omega_1), C_1(r_1, Q_1)\} & \text{otherwise,} \end{cases}$$

where $\omega_1 \equiv \arg \max_{r_1 < z \leq r_1 + Q_1} \{G_1(z)\}$.

Combining all cost terms of both stages, we are ready to present an upper bound on $C(\mathbf{r}, \mathbf{Q})$. For notation convenience, let $\hat{C}_1(r_1, Q_1) \equiv C_1(r_1, Q_1)$. Then, as with the induced-penalty cost allocation scheme (7), we define

$$\hat{G}_1(y) \equiv \begin{cases} G_1(y) - C_1(r_1, Q_1) & \text{if } y \leq r_1, \\ \max\{0, G_1(\omega_1) - C_1(r_1, Q_1)\} & \text{otherwise,} \end{cases}$$

$$\Lambda_2(y) \equiv \mathbb{E}[h_2(y - D_2) + \hat{G}_1(y - D_2)],$$

$$\hat{C}_2(r_2, Q_2) \equiv \frac{1}{Q_2} \left[\lambda K_2 + \int_{r_2}^{r_2+Q_2} \Lambda_2(y) dy \right].$$

Note that $\hat{G}_1(y)$ may not be convex for any given r_1 and Q_1 , but it is indeed so for $(r_1, Q_1) = (r_1^*, Q_1^*)$.

THEOREM 1 (AN UPPER BOUND). For any given modified echelon (\mathbf{r}, \mathbf{Q}) policy, the long-run average system-wide cost has an upper bound: $C(\mathbf{r}, \mathbf{Q}) \leq \sum_{i=1}^2 \hat{C}_i(r_i, Q_i) + \lambda K_1 / Q_2$.

3.2. Effectiveness and Asymptotic Optimality

Now we construct a modified echelon (\mathbf{r}, \mathbf{Q}) heuristic policy and investigate its performance in the two-stage serial system by identifying the gap between the induced-penalty lower bound (see §2.2.2) and its upper bound established in Theorem 1. We will also explore conditions under which this heuristic policy is asymptotically optimal.

3.2.1. Heuristic. In the heuristic, for Stage 1, we select $(r_1, Q_1) = (r_1^*, Q_1^*)$, which minimizes $C_1(r_1, Q_1)$ (see its definition in §2.2.2). With this selection, we have $\hat{C}_1(r_1, Q_1) = C_1(r_1^*, Q_1^*) = C_1^*$. By Lemma 2(i) and the fact that $r_1^* < \omega_1 \leq r_1^* + Q_1^*$, we have $G_1(\omega_1) \leq C_1^*$. Then, by Definitions (7) and (8), $\hat{G}_1(y) = G_1(y)$, which in turn implies that $G_2(y) = \Lambda_2(y)$ and $C_2(r_2, Q_2) = \hat{C}_2(r_2, Q_2)$. Combining Observation 1 and Theorem 1, we immediately have

$$C_B^* \leq C(\mathbf{r}, \mathbf{Q})|_{(r_1, Q_1)=(r_1^*, Q_1^*)} \leq C_1^* + C_2(r_2, Q_2) + \frac{\lambda K_1}{Q_2}. \tag{9}$$

Then we can tighten the upper bound in (9) by optimizing r_2 and Q_2 . Specifically, we solve the following optimization problem:

$$\min_{r_2, Q_2} \tilde{C}_2(r_2, Q_2) \equiv \min_{r_2, Q_2} \left\{ C_2(r_2, Q_2) + \frac{\lambda K_1}{Q_2} \right\}$$

$$= \min_{r_2, Q_2} \frac{1}{Q_2} \left[\lambda(K_1 + K_2) + \int_{r_2}^{r_2+Q_2} G_2(y) dy \right]. \tag{10}$$

Problem (10) is a single-stage problem with a fixed setup cost equal to $K_1 + K_2$, and its objective function $\tilde{C}_2(r_2, Q_2)$

is jointly convex in r_2 and Q_2 (see §2.2.1). Therefore, the optimal solution can be efficiently computed (see Federgruen and Zheng 1992). Define $(\tilde{r}_2^*, \tilde{Q}_2^*) \equiv \arg \min_{r_2, Q_2} \tilde{C}_2(r_2, Q_2)$. We construct a heuristic modified echelon $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policy as follows:

$$(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) = (\hat{r}_1, \hat{Q}_1, \hat{r}_2, \hat{Q}_2) = (r_1^*, Q_1^*, \tilde{r}_2^*, \tilde{Q}_2^*). \tag{MERQ}$$

3.2.2. Effectiveness. The following theorems show the effectiveness of the heuristic policy.

THEOREM 2 ((K_1, K_2)-DEPENDENT PERFORMANCE BOUND). The modified echelon $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policy in (MERQ) is at least $(1 + K_1/K_2)$ -optimal.

Theorem 2 provides a primitive-dependent performance bound of the heuristic policy that depends only on fixed-cost primitives K_1 and K_2 . If $K_2 \geq K_1$ which tends to hold in practice, then the heuristic policy is guaranteed to be 2-optimal. Realizing that this performance bound, though very simple, may be loose, we propose an alternative performance bound that requires more computational efforts but can provide tighter bounds.

THEOREM 3 (ALTERNATIVE PERFORMANCE BOUND). (i) The absolute gap between C_B^* and $C(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ is bounded as follows: $0 \leq C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) - C_B^* \leq C_2^* (\sqrt{1 + C_1^*/(\beta_1^* C_2^*)} - 1)$, where $\beta_1^* \equiv Q_2^*/Q_1^*$.

(ii) The modified echelon $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policy in (MERQ) is at least $1 + 1/(2(\beta_1^* + \sqrt{\beta_1^*}))$ -optimal.

Theorem 3 provides a performance bound that depends on Q_1^* and Q_2^* , which are optimal solutions to two single-stage systems and can be computed efficiently. Both the absolute and relative performance gaps of the heuristic are decreasing in $\beta_1^* = Q_2^*/Q_1^*$. If the ratio β_1^* is not too small, the identified heuristic provides a good performance bound. As an immediate result of Theorem 3, we have the following corollary.

COROLLARY 1. For any $\rho > 0$, if the system primitives are such that $Q_2^*/Q_1^* \geq (1 + 1/\rho - \sqrt{1 + 2/\rho})/2$, then the modified echelon $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policy in (MERQ) is at least $(1 + \rho)$ -optimal. In particular, we have the performance bounds listed in Table 2.

Because we do not have closed-form expressions for Q_1^* and Q_2^* , we draw some insights from their deterministic counterparts to see when the ratio Q_2^*/Q_1^* tends to be large. By definition, Q_i^* , $i = 1, 2$, is the optimal solution to a single-stage system (see §2.2.1), where the fixed cost is K_i , the expected inventory holding and backlog cost function is $G_i(y)$ and demands arrive following a Poisson process with rate λ . When demand is assumed to be a deterministic constant

Table 2. Performance bound dependent on Q_2^*/Q_1^* .

$Q_2^*/Q_1^* \in$	[0.0429, 0.134)	[0.134, 0.382)	[0.382, 1)	[1, 3.21)	[3.21, ∞)
Performance bound	3-opt.	2-opt.	1.5-opt.	1.25-opt.	1.1-opt.

stream, ceteris paribus, the single-stage system boils down to the classic EOQ model with backorders allowed, which is considered as the first-order approximation of the stochastic model. It is easily verified that the deterministic counterparts of Q_1^* and Q_2^* are $Q_1^d = \sqrt{2\lambda K_1(1/h_1 + 1/(p + h_2))}$ and $Q_2^d = \sqrt{2\lambda K_2(1/h_2 + 1/p)}$, respectively. If $h_2 \leq h_1$ and $K_1 \leq K_2$, we always have $Q_1^d \leq Q_2^d$. This suggests that $Q_1^* \leq Q_2^*$ tends to hold if $h_2 \leq h_1$ and $K_1 \leq K_2$. We make the following remarks on the relationships between cost primitives in practical settings that indicates when the heuristic policy in (MERQ) is more likely to perform well.

REMARK 1 (FIXED COST). In a supply chain where the upstream stage procures products overseas and the downstream stage sells them domestically, the fixed cost incurred tends to be higher at the upstream stage than at the downstream stage.

In such a supply chain, there are usually larger economies of scale at the upstream stage. For instance, a standard container shipped by sea or a pallet of air cargo from an overseas manufacturer to a distribution center has a higher fixed cost than a truckload sent from the distribution center to a nearby retailer. We caution that our bounds can be loose if the upstream stage has a lower fixed cost than the downstream stage.

REMARK 2 (HOLDING COST). The echelon inventory holding cost tends to be lower at an upstream stage than that at a downstream stage.

The echelon inventory holding cost h_i at any Stage i usually includes financing costs and physical handling costs. Suppose c_i is the variable order cost per unit at Stage i ; then the financing cost is $r \cdot c_i$, where r is the interest rate. The variable order cost is usually lower at the upstream stage than that at the downstream stage; thus the financing cost is usually lower at the upstream stage. Physical handling also tends to be less expensive at the upstream stage. For instance, a downtown retail store would have a higher out-of-pocket inventory holding cost rate than a suburban warehouse.

3.2.3. Asymptotic Optimality. In the practical settings mentioned in Remarks 1 and 2, it is expected that Q_2^*/Q_1^* is more likely to be larger than 1, and hence a very good performance bound, e.g., 1.25-optimality, can be guaranteed for the heuristic regardless of the values of other system primitives. The following theorem further confirms asymptotic optimality of this heuristic if we take the dominant relationships of cost primitives to the extreme.

THEOREM 4 (ASYMPTOTIC OPTIMALITY). *The modified echelon (\hat{r}, \hat{Q}) policy in (MERQ) is asymptotically optimal if one of the following conditions holds: (i) $K_1/K_2 \rightarrow 0$. (ii) $h_2/h_1 \rightarrow 0$.*

We link Theorem 4 to existing results in the literature. First, when K_2 is fixed and K_1 is scaled down such that $K_1/K_2 \rightarrow 0$, the series of scaled systems converge to one

with setup costs incurred only at the upstream stage. Thus, our result is consistent with the existing finding that an echelon (\mathbf{r}, \mathbf{Q}) policy is optimal for such a system (see, e.g., Chen and Zheng 1994b). Second, when h_1 is fixed and h_2 is scaled down such that $h_2/h_1 \rightarrow 0$, the series of scaled systems converge to one in which Stage 2 will order an infinite amount just at the beginning of the horizon, and the Stage 1 problem is reduced to a single-stage problem with an (r, Q) policy known to be optimal (see, e.g., Zheng 1992).

4. Numerical Experiments

In this section, we present a set of comprehensive numerical experiments. The goal of this numerical study is threefold. First, we verify the effectiveness of the heuristic policy and test its sensitivity to system primitives with a large set of examples. Second, we verify the asymptotic optimality demonstrated in Theorem 4. Third, we compare our heuristic policy with integer-ratio (\mathbf{r}, \mathbf{Q}) policies.

To test the performance of the heuristic, we compare the upper bound of its system-wide cost (see Theorem 1 and (9)) with the induced-penalty lower bound of the optimal cost (see Lemma 3). Hence, the actual performance of the proposed heuristic should be better than the performance reported here. We denote by $\bar{C}(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ the cost *upper bound* of the heuristic, and by LB the induced-penalty *lower bound*. We define the following percentage difference:

$$\xi(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) \equiv \frac{\bar{C}(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) - \text{LB}}{\text{LB}} \times 100\%,$$

which is an *upper bound* on the effectiveness of the heuristic policy.

We examine two-stage systems with eight primitives: lead times L_1 and L_2 , setup costs K_1 and K_2 , holding costs h_1 and h_2 , the shortage penalty cost p and arrival rate λ . The complete test set of primitive values is given by: $L_1 \in \{0.2, 0.5, 1, 2, 5\}$, $\lambda \in \{2, 5, 15, 20\}$, $K_2 \in \{10, 30, 50, 100, 200\}$, $h_2 \in \{0.1, 0.2, 0.5, 1, 2\}$, $p \in \{0.5, 1, 3, 10\}$ with other primitives fixed as $L_2 = 1$, $K_1 = 10$, $h_1 = 2$. All combinations of these primitives provide $5 \times 4 \times 5 \times 5 \times 4 = 2,000$ test instances. To investigate the effect of the ratio $\beta_1^* = Q_2^*/Q_1^*$ on the performance of the heuristic, we present the numerical results in Table 3 by classifying them according to the value of β_1^* . As expected, we can see from Table 3 that the heuristic tends to perform better as the ratio β_1^* is larger, although not monotonically. When $\beta_1^* \in (2, 2.5]$, the average optimality gap achieves 0.61%. When $\beta_1^* > 4.5$ with a total of 662 instances, the average optimality gap is no more than 0.06% and the worst performance is only 0.51% that indicates that the heuristic is almost optimal for these instances. In the case of $\beta_1^* \in (1, 1.5]$, the average optimality gap is around 1.52% that is consistent with our theoretical result that the heuristic is 1.25-optimal. A natural question is whether the performance of the heuristic is consistent *within*

Table 3. Overall performance of modified echelon (\hat{r}, \hat{Q}) policies.

Range of $\beta_1^* = Q_2^*/Q_1^*$	(0, 1]	(1, 1.5]	(1.5, 2]	(2, 2.5]	(2.5, 3]	(3, 3.5]	(3.5, 4]	(4, 4.5]	(4.5, 5]	(5, ∞)
Number of instances	23	138	198	186	230	200	209	154	174	488
Average $\xi(\hat{r}, \hat{Q})$ (%)	1.28	1.52	1.33	0.61	0.50	0.23	0.16	0.15	0.06	0.05
st. dev. of $\xi(\hat{r}, \hat{Q})$ (%)	0.80	0.79	0.94	0.61	0.46	0.29	0.23	0.24	0.10	0.07
Minimum of $\xi(\hat{r}, \hat{Q})$ (%)	0.39	0.02	0.03	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Maximum of $\xi(\hat{r}, \hat{Q})$ (%)	3.42	3.38	3.63	2.31	1.96	1.40	1.28	1.13	0.51	0.46

Table 4. Primitives for base scenario.

L_1	L_2	K_1	K_2	h_1	h_2	p	λ
2	1	10	100	2	1	3	5

each range of β_1^* . To answer that question, we define the standard deviation (denoted by st. dev.) for each range as

$$\text{st. dev.} \equiv \sqrt{\frac{(\text{Individual Error} - \text{Average } \xi_1)^2}{\text{Number of Instances}}}$$

which measures the variation of $\xi(\hat{r}, \hat{Q})$ within each range. We can see from Table 3 that the st. dev. in each range is relatively small and decreases as β_1^* increases. This set of numerical results shows that the effectiveness of the heuristic is consistent within each range.

4.1. Sensitivity of Model Primitives

We investigate the effect of system primitives on the performance of the heuristic. To that end, we set a base scenario and

vary the primitive values once at a time. Table 4 summarizes the values of all primitives in the base scenario.

The numerical results on sensitivity are presented in Tables 5–7. Table 5 demonstrates the effectiveness of the heuristic under various values of K_2 and h_2 . It is observed that the heuristic performs better as K_2 increases and is almost optimal when $K_2 = 200$ with the ratio $K_1/K_2 = 1/20$. That is consistent with Theorem 4(i), which states that the heuristic is asymptotically optimal as $K_1/K_2 \rightarrow 0$. Table 5 also shows that the performance of the heuristic tends to be better as h_2 becomes smaller, though not monotonically. Note that the alternative performance bound links the effectiveness to the ratio of Q_2^*/Q_1^* . As h_2 decreases, Q_2^* becomes larger, but Q_1^* does not change, and thus, the ratio of Q_2^*/Q_1^* becomes larger and the performance tends to be better. Table 6 illustrates the effectiveness of the heuristic under various values of λ and p . It is observed that the effectiveness does not seem monotone in p or λ . Table 7 shows how the lead time of Stage 1, L_1 , affects the effectiveness of the heuristic. It is observed that the effectiveness seems to increase in L_1 , though Q_2^*/Q_1^* seems to decrease in L_1 . This

Table 5. Comparisons between lower bounds and upper bounds under various K_2 and h_2 .

h_2	K_2	$r_1^* = \hat{r}_1$	$Q_1^* = \hat{Q}_1$	r_2^*	Q_2^*	\hat{r}_2	\hat{Q}_2	β_1^*	LB	UB	$\xi(\hat{r}, \hat{Q})$ (%)
2	10	7	10	7	11	6	14	1.1	41.6361	42.1926	1.34
	30	7	10	5	16	4	19	1.6	49.0639	49.2433	0.37
	50	7	10	3	21	2	23	2.1	54.3771	54.4876	0.20
	100	7	10	0	29	-1	31	2.9	64.3240	64.3771	0.08
	200	7	10	-5	41	-5	42	4.1	78.5235	78.5419	0.02
1	10	6	11	8	12	7	17	1.09	30.3858	30.8014	1.37
	30	6	11	6	20	5	24	1.82	36.3789	36.5382	0.44
	50	6	11	4	27	4	28	2.45	40.6223	40.6515	0.072
	100	6	11	2	37	1	39	3.36	48.5221	48.5579	0.074
	200	6	11	-2	52	-2	53	4.73	59.7721	59.7847	0.021
0.5	10	6	10	8	17	7	23	1.7	24.0472	24.4436	1.65
	30	6	10	7	27	6	31	2.7	28.6264	28.7185	0.32
	50	6	10	6	34	5	38	3.4	31.8546	31.9042	0.16
	100	6	10	4	48	3	51	4.8	37.8378	37.8543	0.044
	200	6	10	1	68	0	71	6.8	46.3527	46.3710	0.039
0.2	10	5	11	9	24	8	34	2.18	19.3626	19.6224	1.34
	30	5	11	8	41	7	47	3.73	22.4304	22.5076	0.34
	50	5	11	7	52	7	57	4.73	24.5725	24.6085	0.15
	100	5	11	6	73	5	78	6.64	28.5414	28.5650	0.08
	200	5	11	4	103	4	106	9.36	34.1794	34.1852	0.02
0.1	10	5	11	10	33	9	47	3	17.3108	17.4988	1.09
	30	5	11	9	56	8	66	5.09	19.5334	19.5949	0.31
	50	5	11	8	73	8	79	6.64	21.0753	21.0998	0.12
	100	5	11	7	102	7	107	9.27	23.9291	23.9393	0.04
	200	5	11	6	144	5	148	13.09	27.9842	27.9878	0

Table 6. Comparisons between lower bounds and upper bounds under various λ and p .

λ	p	$r_1^* = \hat{r}_1$	$Q_1^* = \hat{Q}_1$	r_2^*	Q_2^*	\hat{r}_2	\hat{Q}_2	β_1^*	LB	UB	$\xi(\hat{r}, \hat{Q})$ (%)
2	0.5	-1	8	-22	35	-23	36	4.38	17.8350	17.8396	0.026
	1	0	7	-12	28	-13	30	4	21.9137	21.9328	0.087
	3	1	7	-3	24	-3	25	3.43	28.3382	28.3691	0.11
	10	3	6	3	21	2	23	3.5	34.3645	34.4061	0.12
	20	4	6	4	22	4	23	3.67	37.4369	37.4689	0.085
5	0.5	2	13	-29	55	-31	58	4.23	31.8924	31.9187	0.082
	1	4	11	-14	45	-15	47	4.09	38.3572	38.3797	0.059
	3	6	11	2	37	1	39	3.36	48.5221	48.5579	0.074
	10	9	10	10	34	10	36	3.4	57.9641	58.0123	0.083
	20	11	9	13	34	13	35	3.78	62.7535	62.7819	0.045
15	0.5	17	22	-31	95	-34	99	4.32	67.9317	67.9553	0.035
	1	19	21	-5	78	-7	82	3.71	79.1425	79.1949	0.066
	3	24	18	22	64	21	68	3.56	96.7356	96.8254	0.093
	10	29	16	37	59	36	62	3.69	113.0254	113.0854	0.053
	20	31	16	42	58	41	62	3.63	121.1066	121.2110	0.086
20	0.5	25	25	-28	110	-32	116	4.4	83.7988	83.8512	0.06
	1	28	23	3	90	1	94	3.91	96.7415	96.7862	0.046
	3	33	21	34	74	33	77	3.52	117.0624	117.1090	0.040
	10	39	19	50	69	50	72	3.63	135.8600	135.9287	0.051
	20	42	18	56	68	56	70	3.78	145.1236	145.1645	0.028

Table 7. Comparisons between lower bounds and upper bounds under various L_1 .

L_1	$r_1^* = \hat{r}_1$	$Q_1^* = \hat{Q}_1$	r_2^*	Q_2^*	\hat{r}_2	\hat{Q}_2	β_1^*	LB	UB	$\xi(\hat{r}, \hat{Q})$ (%)
0.2	-2	8	-7	37	-7	39	4.625	37.5574	37.5993	0.11
0.5	-1	9	-5	37	-6	39	4.11	39.3931	39.4321	0.099
1	1	10	-3	37	-3	38	3.7	42.4770	42.4932	0.038
2	6	11	2	37	1	39	3.36	48.5221	48.5579	0.074
5	21	12	16	37	15	39	3.08	66.0201	66.0594	0.059

Table 8. Performance of modified echelon (\hat{r}, \hat{Q}) policies when $Q_2^* < Q_1^*$.

K_2	K_1	$r_1^* = \hat{r}_1$	$Q_1^* = \hat{Q}_1$	r_2^*	Q_2^*	\hat{r}_2	\hat{Q}_2	β_1^*	LB	UB	$\xi(\hat{r}, \hat{Q})$ (%)
10	20	5	14	7	12	5	20	0.86	33.4201	34.4253	3.01
	50	3	20	5	12	0	29	0.6	39.9605	43.4241	8.67
	80	1	25	3	13	-3	35	0.52	44.8682	50.2240	11.94
	100	0	28	2	13	-5	39	0.46	47.6625	54.3480	14.03
	200	-4	40	-2	13	-12	53	0.33	58.7656	70.3018	19.63
	300	-6	47	-5	13	-18	65	0.28	67.3711	83.1980	23.49
	500	-11	62	-9	12	-27	83	0.19	81.0860	103.4457	27.58

observation cautions that though the effectiveness of the heuristic can be linked to the value of Q_2^*/Q_1^* , the sensitivity of Q_2^*/Q_1^* may not move in the same direction as that of the heuristic performance.

4.2. Performance When $Q_2^* < Q_1^*$

By Corollary 1, the heuristic is at least 1.25-optimal when $Q_2^* \geq Q_1^*$. We are particularly interested in the performance of the heuristic for the case of $Q_2^* < Q_1^*$. We vary $K_1 \in \{20, 50, 80, 100, 200, 300, 500\}$ and let $K_2 = 10$ with other primitives fixed as in Table 4. Most of the cases here, in which K_1 is more than a double-digit multiple of K_2 , would be extremely rare in practical settings. There are

two observations from Table 8. First, despite the strong dominance of K_1 over K_2 , which is rare, the value of β_1^* is hardly ever very low. For example, β_1^* is around 0.19 even when $K_1 = 500$, which is 50 times of K_2 . In most of the cases, β_1^* is above 0.134 and therefore the heuristic is at least 2-optimal by Corollary 1. Second, we observe consistently that the actual effectiveness is much better than the theoretical bound. For example, the effectiveness of the heuristic can be as good as 1.2758-optimal, even though β_1^* is as small as 0.19. For this particular case when $\beta_1^* = 0.19$, the theoretical performance bound is 1.7989-optimal by Corollary 1.

We caution that the performance of our heuristic policy decreases when $\beta_1^* = Q_2^*/Q_1^*$ becomes smaller. When β_1^* is

Table 9. Comparisons between modified echelon ($\hat{\mathbf{r}}, \hat{\mathbf{Q}}$) policies and the optimal integer-ratio ($\mathbf{r}^n, \mathbf{Q}^n$) policies.

λ	K_2	r_1^n	Q_1^n	r_2^n	Q_2^n	$C(\mathbf{r}^n, \mathbf{Q}^n)$	\hat{r}_1	\hat{Q}_1	\hat{r}_2	\hat{Q}_2	$C(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ or UB	LB	$\xi(\mathbf{r}^n, \mathbf{Q}^n)$ (%)	$\xi(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ (%)
$L_1 = 1, L_2 = 2, K_1 = 10, h_1 = 0.5, h_2 = 1, p = 5$														
1	5	0	6	1	6	8.3828	0	7	1	7	8.3948	8.0216	4.50	4.65
	100	0	8	-1	16	17.2446	0	7	-1	17	17.1451	17.1317	0.66	0.08
	200	0	8	-2	24	22.5396	0	7	-2	23	22.3690	22.3692	0.77	0.009
	400	0	8	-3	32	30.0028	0	7	-3	31	29.8456	29.8456	0.53	0
5	5	4	14	12	14	21.4394	4	16	11	15	21.5758	20.6433	3.86	4.51
	100	4	18	8	36	41.3823	4	16	7	37	41.1538	41.1448	0.58	0.02
	200	4	17	5	51	53.1838	4	16	5	51	52.8492	52.8355	0.66	0.026
	400	5	14	1	70	69.9265	4	16	1	71	69.5690	69.5613	0.52	0.01
10	5	9	19	26	19	33.2192	9	22	25	21	33.3534	32.1126	3.45	3.86
	100	8	25	20	50	61.4478	9	22	19	53	61.1632	61.1264	0.53	0.06
	200	8	24	16	72	78.1643	9	22	16	72	77.6698	77.6614	0.65	0.01
	400	10	20	10	100	101.8232	9	22	11	100	101.3087	101.3039	0.51	0.005
15	5	14	23	41	23	43.4355	13	28	39	26	43.6150	42.0567	3.28	3.70
	100	13	31	33	62	78.0018	13	28	32	64	77.6335	77.6118	0.50	0.03
	200	13	29	29	87	98.4661	13	28	28	88	97.8904	97.8676	0.61	0.02
	400	15	25	21	125	127.4852	13	28	22	122	126.8234	126.8215	0.52	0.005
$L_1 = 1, L_2 = 2, K_1 = 100, h_1 = 0.5, h_2 = 1, p = 5$														
1	10	-1	15	-1	15	17.5833	-1	21	-2	17	17.8676	14.1869	23.94	25.94
	20	-1	15	-1	15	18.2500	-1	21	-2	17	18.4559	15.6931	16.29	17.60
	50	-1	16	-1	16	20.1719	-1	21	-2	19	20.3289	18.7718	7.46	8.29
	100	-2	19	-2	19	22.9605	-1	21	-3	22	23.1310	22.4025	2.49	3.25
5	10	2	32	7	32	42.1400	1	47	5	37	42.5875	34.4419	22.35	23.65
	20	2	33	7	33	43.6654	1	47	4	39	44.3012	37.7927	15.54	17.22
	50	2	37	6	37	47.9091	1	47	4	43	48.3893	44.7731	7.00	8.07
	100	1	42	5	42	54.1384	1	47	3	49	54.4414	52.8987	2.43	2.92
10	10	6	46	19	46	62.5360	4	67	16	52	63.1548	51.5881	21.22	22.42
	20	6	47	19	47	64.6939	4	67	15	55	65.5001	56.3312	14.85	16.27
	50	6	53	17	53	70.7028	4	67	14	61	71.4751	66.2391	6.74	7.90
	100	5	60	16	60	79.5638	4	67	13	70	79.9998	77.7434	2.34	2.90
15	10	11	55	32	55	79.3476	8	82	27	65	80.4344	65.9332	20.34	21.99
	20	10	58	31	58	81.9860	8	82	27	67	82.8912	71.7467	14.27	15.53
	50	10	64	30	64	89.3438	8	82	26	75	90.2690	83.8715	6.52	7.63
	100	9	74	27	74	100.2038	8	82	24	86	100.7649	97.9737	2.28	2.85

very small, our heuristic can perform poorly. The reason is as follows. In our heuristic, when determining \hat{Q}_2 , we balance the cost terms of $C_2(r_2, Q_2)$ and $\lambda K_1/Q_2$ (see Problem (10)). To obtain the alternative performance bound, we compare the sum of these two cost terms evaluated at $Q_2 = \hat{Q}_2$ with $C_2(r_2^*, Q_2^*)$ that is part of the lower bound. Considering the term $C_2(r_2, Q_2)$ alone, intuitively, it is more desirable to force \hat{Q}_2 to be as close to Q_2^* as possible. However, when Q_1^* is large, e.g., induced by an extremely large K_1 , our heuristic requires \hat{Q}_2 to be large as well, hence deviating from Q_2^* , to dampen the effect of high setup cost K_1 that might otherwise be incurred more often for irregular shipment periods.

4.3. Comparison with Integer-Ratio (\mathbf{r}, \mathbf{Q}) Policies

Now we compare numerically the proposed modified echelon ($\hat{\mathbf{r}}, \hat{\mathbf{Q}}$) heuristic policy with the integer-ratio echelon (\mathbf{r}, \mathbf{Q}) policy (i.e., echelon (r, nQ) policy) that also charges shipment-based fixed costs; see, e.g., Chen and Zheng (1994a). For some special cases of policy parameters, we show analytically that the two types of policies are equivalent in the sense of generating the same sample path of inventory flows.

LEMMA 7. For a two-stage system under any modified echelon (\mathbf{r}, \mathbf{Q}) = (r_1, r_2, Q_1, Q_2) policy, if $r_2 - Q_2 \leq r_1$ and $Q_2 \leq Q_1$, then the long-run average cost is the same as that under an integer-ratio policy with the policy parameters (\mathbf{r}, \mathbf{Q}') = (r_1, r_2, Q_2, Q_2).

We denote by r_i^n and Q_i^n the policy parameters for Stage i of the optimal integer-ratio (\mathbf{r}, \mathbf{Q}) policy, and by $C(\mathbf{r}^n, \mathbf{Q}^n)$ its system-wide cost. We list numerical comparisons between our heuristic and the optimal integer-ratio (\mathbf{r}, \mathbf{Q}) policy in Table 9, where the first half scenarios are adopted directly from Table 1 of Chen and Zheng (1994a). By Lemma 7, we can calculate the real cost of those modified echelon ($\hat{\mathbf{r}}, \hat{\mathbf{Q}}$) policies that satisfy $\hat{r}_2 - \hat{Q}_2 \leq \hat{r}_1$ and $\hat{Q}_2 \leq \hat{Q}_1$, because they are equivalent to integer-ratio (\mathbf{r}, \mathbf{Q}) policies whose explicit cost expression has been given by Chen and Zheng (1994a). Therefore, we replace the cost upper bound $\bar{C}(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ by the real cost in evaluating the heuristic’s effectiveness for these instances. It can be seen from Table 9 that for these cases, our proposed heuristic policy, equivalent to a specific integer-ratio (\mathbf{r}, \mathbf{Q}) policy, performs no better than the optimal integer-ratio (\mathbf{r}, \mathbf{Q}) policy. These cases provide strong evidence to show that the well-studied integer-ratio

policies are also very efficient in many scenarios. On the other hand, for these cases, again by the equivalency established in Lemma 7, if the optimal integer-ratio (\mathbf{r}, \mathbf{Q}) policy happens to satisfy $r_2^n - Q_2^n \leq r_1^n$ and $Q_1^n = Q_2^n$ (which is indeed the case for all our test instances such that $\hat{r}_2 - \hat{Q}_2 \leq \hat{r}_1$ and $\hat{Q}_2 \leq \hat{Q}_1$), the modified echelon (\mathbf{r}, \mathbf{Q}) policy with policy parameters $(\mathbf{r}, \mathbf{Q}) = (r_1^n, r_2^n, Q_2^n, Q_2^n)$ is equivalent to the optimal integer-ratio (\mathbf{r}, \mathbf{Q}) policy. Hence, for all our test instances, there exists a modified echelon (\mathbf{r}, \mathbf{Q}) policy that is at least *as good as* the optimal integer-ratio (\mathbf{r}, \mathbf{Q}) policy.

We also observe that our heuristic is more likely to outperform the optimal integer-ratio (\mathbf{r}, \mathbf{Q}) policy when $\hat{Q}_2 > \hat{Q}_1$ or $K_2 > K_1$. By design, the modified echelon (\mathbf{r}, \mathbf{Q}) policies are more flexible than the integer-ratio (\mathbf{r}, \mathbf{Q}) policies. First, modified echelon (\mathbf{r}, \mathbf{Q}) policies do not require Q_2 to be an integral multiple of Q_1 . This flexibility can be desirable when the system primitives intrinsically demand an (r, Q) policy at each stage that does not satisfy the integer-ratio constraint across stages. Second, under a modified echelon (\mathbf{r}, \mathbf{Q}) policy, the actual shipment quantity from Stage 2 to Stage 1 may not be an integral multiple of Q_1 and can be of any size no greater than the on-hand inventory of Stage 2. This flexibility at the operational level can be valuable when the downstream stage has been starved of inventory. The disadvantage of modified echelon (\mathbf{r}, \mathbf{Q}) policies is that the frequency of incurring Stage 1's fixed setup costs can be high because of shipments that may be less than full batches. Our numerical comparisons show that the benefits of more flexible modified echelon (\mathbf{r}, \mathbf{Q}) policies tend to outweigh the costs when $\hat{Q}_2 > \hat{Q}_1$ or $K_2 > K_1$, either of which is more likely to result in less frequent partial shipments to Stage 1.

Moreover, it can be observed that our heuristic policy is very sensitive to K_2 and seems to converge quickly and become near-optimal as K_2 increases, whereas the effectiveness of the optimal integer-ratio (\mathbf{r}, \mathbf{Q}) policies does not seem to converge quickly when K_2 becomes larger. Both policies tend to perform better when λ becomes larger without demonstrating any strictly monotone relationship in λ .

5. Conclusion

In this paper, we have studied the classic serial inventory system. In the presence of setup costs per shipment at downstream stages, the optimal policy is unknown. We have studied a class of modified echelon (\mathbf{r}, \mathbf{Q}) policies that do not have Q_{i+1}/Q_i to be integer-ratioed. We have evaluated the performance of a constructed modified echelon (\mathbf{r}, \mathbf{Q}) heuristic policy by identifying its performance gap with the optimal cost, and explored conditions under which the heuristic policy is asymptotically optimal.

Our results can be extended in several directions. First, we extend most of our results for the two-stage system to a multiple-stage system (see Online Appendix A). Second, because an assembly system can be transformed into a serial

system under some mild conditions (see Rosling 1989), similar results can be obtained for assembly systems. Third, our results *may* also hold in the setting with stochastic sequential lead times such that orders do not cross in time, for which we leave the verification to future research.

Our approach has limitations. First, in our model, we consider 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 worth noting that even Chen 1999's 94%-effective result is not readily extendable to these settings.) The reason for that is as follows. For the periodic-review system, shipment decisions are made only at the end of each period, and thus, unlike the continuous-review system with a Poisson arrival process, a replenishment is not necessarily triggered when the echelon inventory position falls exactly to a preset level even with ample supply. Therefore, the analysis on the steady-state distribution of echelon inventory positions becomes very complicated. The echelon inventory position at the end of each period has a jump with the size of the realized demand, instead of a gradual decrease in steps of one unit. As a result, even for the single-stage periodic-review model, the steady-state inventory position under an (s, S) policy no longer follows a uniform distribution over $\{s + 1, \dots, S\}$, but with a complex steady-state distribution (see Morse 1959). Thus, for a two-stage system, even when Stage 1 is always in a regular shipment period, the expected cost of Stage 1 does not have a simple form like that in the continuous-review system with Poisson demand (see Equation (5)). Because the performance bound analysis relies crucially on the properties of Equation (5), our analysis cannot be readily extended to the periodic-review system. For similar reasons, our results are not readily extendable to the continuous-review system with a compound Poisson arrival process.

Supplemental Material

Supplemental material to this paper is available at <http://dx.doi.org/10.1287/opre.2014.1291>.

Acknowledgments

The authors thank the Department Editor Chung Piaw Teo, the Associate Editor, and two anonymous reviewers for their guidance and constructive comments. The authors are grateful to Lifeng Chen for providing the inspiration to study the proposed heuristic, Awi Federgruen for proposing the terminology of modified echelon (r, Q) policy, and Frank Chen, Retsef Levi, Özalp Özer, Jeannette Song and Sean Zhou for their helpful suggestions. The first author is grateful for financial support by the 2008–2010 Innovation Research Program Award of Hewlett-Packard Labs and the 2011 Connaught New Researcher Award of University of Toronto. The second author is grateful for financial support by the National Natural Science Foundation of China [Grant 71201142].

Appendix. Proofs

PROOF OF LEMMA 2. (i) For any $y \in [r^*, r^* + Q^*]$, there exists $\mu \in [0, 1]$ such that $y = \mu r^* + (1 - \mu)(r^* + Q^*)$. By the convexity

of $G(y)$, $G(y) = G(\mu r^* + (1 - \mu)(r^* + Q^*)) \leq \mu G(r^*) + (1 - \mu)G(r^* + Q^*) = C^*$, where the last equality is because of Lemma 1(ii).

(ii) From Lemma 1(v), $A(Q^*) = Q^*H(Q^*) - \int_0^{Q^*} H(y) dy = \lambda K$. Then, Lemma 1(iv) implies that $\lambda K \leq Q^*H(Q^*) - \frac{1}{2}Q^*H(Q^*) = \frac{1}{2}Q^*H(Q^*) = \frac{1}{2}Q^*C^*$. \square

PROOF OF LEMMA 4. We prove the lemma by contradiction. Suppose over the cycle $[T_1^j, T_1^{j+1})$ of Stage 1, there exist two irregular shipment periods. First, it must be the case that at the end of the first irregular shipment period, the inventory position of Stage 1 exactly hits r_1 ; otherwise if the inventory position of Stage 1 drops below r_1 at the end of the first irregular shipment period, there is a contradiction: By the stipulation that there exist two irregular shipment periods in the cycle $[T_1^j, T_1^{j+1})$, at the end of the first irregular shipment period, there must be a second shipment to Stage 1 containing units from the j th order of Stage 2. This implies that at the end of the first shipment period, Stage 2 should have some on-hand inventory to raise the inventory position of Stage 1 above r_1 ; otherwise, the second shipment containing units of the j th order of Stage 2 would not exist. Second, by the conclusion of the first part and the definition of an irregular shipment period, the first irregular shipment period must start with an inventory position of Stage 1 less than $r_1 + Q_1$. However, again, that contradicts the assumption that there is a second shipment containing units from the j th order of Stage 2. That is because under our modified echelon (r, Q) policy, at the beginning of the first irregular shipment period, Stage 2 must use its j th order to raise Stage 1's echelon inventory position as close as possible to $r_1 + Q_1$; hence if at the beginning of the first irregular shipment period, Stage 1's inventory position is still below $r_1 + Q_1$ after replenishment, units in the j th order of Stage 2 must have all been shipped to Stage 1 and there can be no second shipment containing units of the j th order of Stage 2. It is worth noting that this proof always holds without assuming $Q_1 \leq Q_2$. \square

PROOF OF LEMMA 5. STAGE 2. Because Stage 2 orders from an amply stocked supplier, for any j , there are Q_2 units of demand in total over the $(j + 1)$ th cycle $[T_2^j, T_2^{j+1})$ at Stage 2 by Definitions 1 and 3. Because demand arrives following a homogeneous Poisson process with rate λ , the expectation of the length of any cycle j at Stage 2 is $\mathbb{E}[T_2^{j+1} - T_2^j] = Q_2/\lambda$. Hence, the result holds for Stage 2.

STAGE 1. The order in the j th cycle will arrive at Stage 2 at time $T_2^j + L_2$. Denote by Δt_1^j the time lag, for the first time, in receiving units contained in this order at Stage 1, i.e., $\Delta t_1^j \equiv T_1^j - (T_2^j + L_2)$. Clearly, $\Delta t_1^j \geq 0$. By the definition of T_2^j , we know that $IP_2^-(T_2^j) = r_2$, $IP_2(T_2^j) = r_2 + Q_2$, and $IL_2(T_2^j + L_2) = IP_2(T_2^j) - D(T_2^j, T_2^j + L_2) = r_2 + Q_2 - D(T_2^j, T_2^j + L_2)$. At time $T_2^j + L_2$ before units contained in the Stage 2's j th order are sent to Stage 1, there are at most $(r_2 - D(T_2^j, T_2^j + L_2) - r_1)^+$ units in addition to r_1 in Stage 1's inventory position. Let $t(x) \equiv \inf\{t \geq 0 \mid D(0, t] \geq x\}$. Thus, for all j , we have

$$0 \leq \Delta t_1^j \leq t(r_2 - D(T_2^j, T_2^j + L_2) - r_1)^+ \leq t(r_2 - r_1)^+. \quad (11)$$

Note that the lower and upper bound on Δt_1^j depend only on the policy parameters. Therefore,

$$\begin{aligned} & \lim_{j \rightarrow \infty} \mathbb{E}[T_1^{j+1} - T_1^j]/j \\ &= \lim_{j \rightarrow \infty} \mathbb{E}[(T_2^{j+1} + L_2 + \Delta t_1^{j+1}) - (T_2^j + L_2 + \Delta t_1^j)]/j \end{aligned}$$

$$\begin{aligned} &= \lim_{j \rightarrow \infty} \{\mathbb{E}[T_2^{j+1} - T_2^j]/j + \mathbb{E}[\Delta t_1^{j+1} - \Delta t_1^j]/j\} \\ &= \lim_{j \rightarrow \infty} \mathbb{E}[T_2^{j+1} - T_2^j]/j = jQ_2/(j\lambda) = Q_2/\lambda, \end{aligned}$$

where the third equality follows from (11) and the second-to-last equality is because of $\mathbb{E}[T_2^{j+1} - T_2^j] = Q_2/\lambda$ for all j . Therefore, the result also holds for Stage 1. \square

PROOF OF LEMMA 6. If $IL_2(t) \leq r_1$, by the definition of modified (r_1, Q_1) policy at Stage 1, then $IP_1(t) = IL_2(t) \leq r_1$, which implies that Stage 1 must be in an irregular shipment period at time t . Therefore, in this case, $\hat{\Gamma}_1(IL_2(t)) = G_1(IP_1(t)) = G_1(IL_2(t))$.

If $IL_2(t) > r_1$, the inventory position of Stage 1 must be in the range $(r_1, r_1 + Q_1]$, i.e., $r_1 < IP_1(t) \leq r_1 + Q_1$. In this case, it is possible that Stage 1 is either in a regular or in an irregular shipment period. To obtain an upper bound on $\hat{\Gamma}_1(IL_2(t))$, we charge the larger one between expected cost rates of the regular and irregular shipment period. That is, $\hat{\Gamma}_1(IL_2(t)) \leq \max\{G_1(IP_1(t)), C_1(r_1, Q_1)\} \leq \max\{G_1(\omega_1), C_1(r_1, Q_1)\}$, where the last inequality follows from the definition of ω_1 and the fact that $r_1 < IP_1(t) \leq r_1 + Q_1$. \square

PROOF OF THEOREM 1. Note that because the entire horizon is a union of cycles $[T_1^j, T_1^{j+1})$ for all $j \in \mathbb{N}$, the upper bound stated in Lemma 6 holds for any time. By the definition of $\hat{G}_1(y)$ in (8), we obtain

$$\hat{\Gamma}_1(IL_2(t)) \leq \hat{G}_1(IL_2(t)) + C_1(r_1, Q_1). \quad (12)$$

We denote by $\Gamma_2(IP_2(t))$ the total expected cost rate of both stages at time t when the inventory position of Stage 2 is $IP_2(t)$, excluding the setup costs incurred at Stage 2 and the setup costs incurred at Stage 1 for irregular shipment periods. By this definition, $\Gamma_2(IP_2(t))$ constitutes two parts: (i) the inventory holding cost at Stage 2, and (ii) the total costs at Stage 1 excluding irregular shipment period's setup costs. That is, according to the cost accounting scheme (see Definition 2),

$$\begin{aligned} \Gamma_2(IP_2(t)) &= \mathbb{E}[h_2(IP_2(t) - D_2)] + \mathbb{E}[\hat{\Gamma}_1(IP_2(t) - D_2)] \\ &\leq \mathbb{E}[h_2(IP_2(t) - D_2)] + \mathbb{E}[\hat{G}_1(IP_2(t) - D_2)] \\ &\quad + C_1(r_1, Q_1) \\ &= \Lambda_2(IP_2(t)) + C_1(r_1, Q_1), \end{aligned} \quad (13)$$

where the first equality follows from (1), the inequality from (12), and the last equality from (8).

Because Stage 2 has an unlimited supply from the external supplier, under the modified echelon (r, Q) policy, the inventory position of Stage 2, $IP_2(t)$, is uniformly distributed on $\{r_2 + 1, \dots, r_2 + Q_2\}$. Therefore, by the definition of $\Gamma_2(IP_2(t))$, the long-run average system-wide cost, excluding the setup costs incurred at Stage 1 for irregular shipment periods, can be expressed under the continuous approximation and bounded as:

$$\begin{aligned} & \frac{1}{Q_2} \left[\lambda K_2 + \int_{r_2}^{r_2+Q_2} \Gamma_2(y) dy \right] \\ & \leq \frac{1}{Q_2} \left[\lambda K_2 + \int_{r_2}^{r_2+Q_2} [\Lambda_2(y) + C_1(r_1, Q_1)] dy \right] \\ & = \sum_{i=1}^2 \hat{C}_i(r_1, Q_i), \end{aligned} \quad (14)$$

where the inequality is because of (13).

Finally, combining (14) and Observation 4, we find that the long-run average system-wide cost can be bounded as: $C(\mathbf{r}, \mathbf{Q}) \leq \sum_{i=1}^2 \hat{C}_i(r_i, Q_i) + \lambda K_1/Q_2$. \square

PROOF OF THEOREM 2. We have $C_1^* + C_2^* \leq C_B^* \leq C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) \leq C_1^* + C_2(r_2^*, Q_2^*) + (\lambda K_1)/Q_2^* = C_1^* + C_2^* + (K_1/K_2)((\lambda K_2)/Q_2^*) \leq C_1^* + (1 + K_1/K_2)C_2^* \leq (1 + K_1/K_2)(C_1^* + C_2^*)$, where the first inequality is due to the induced-penalty lower bound (see Lemma 3) and the second and third inequalities are because of (9) and the optimal selection of $(\hat{r}_2^*, \hat{Q}_2^*)$. \square

PROOF OF THEOREM 3. We have

$$\begin{aligned} C_B^* &\leq C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) \leq C_1^* + C_2(r_2, Q_2) + \frac{\lambda K_1}{Q_2} \\ &\leq C_1^* + \epsilon \left(\frac{Q_2}{Q_2^*} \right) C_2^* + \frac{C_1^* Q_1^*}{2Q_2}, \end{aligned} \quad (15)$$

where the first and second inequalities are because of (9) and the optimal selection of (r_2^*, Q_2^*) , and the third inequality is due to Lemma 1(iii) and Lemma 2(ii). Then, we can further tighten up the upper bound in (15) by selecting Q_2 as

$$\begin{aligned} \check{Q}_2 &\equiv \arg \min_{Q_2} \left\{ \epsilon \left(\frac{Q_2}{Q_2^*} \right) C_2^* + \frac{C_1^* Q_1^*}{2Q_2} \right\} \\ &= \sqrt{\frac{(Q_2^*)^2 C_2^* + Q_1^* Q_2^* C_1^*}{C_2^*}} = Q_2^* \sqrt{1 + \frac{C_1^*}{\beta_1^* C_2^*}}. \end{aligned}$$

(i) By (15), we have $0 \leq C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) - C_B^* \leq C_1^* + \epsilon(\check{Q}_2/Q_2^*)C_2^* + C_1^* Q_1^*/(2\check{Q}_2) - C_B^* \leq C_1^* + \epsilon(\check{Q}_2/Q_2^*)C_2^* + C_1^* Q_1^*/(2\check{Q}_2) - (C_1^* + C_2^*) = C_2^*(\sqrt{1 + C_1^*/(\beta_1^* C_2^*)} - 1)$, where the third inequality is because of the induced-penalty lower bound (see Lemma 3) and the last equality follows from the selection of $Q_2 = \check{Q}_2$ and $r_2 = r_2(\check{Q}_2)$.

(ii) $(C(\hat{\mathbf{r}}, \hat{\mathbf{Q}}) - C_B^*)/C_B^* \leq C_2^*(\sqrt{1 + C_1^*/(\beta_1^* C_2^*)} - 1)/C_B^* \leq C_2^*(\sqrt{1 + C_1^*/(\beta_1^* C_2^*)} - 1)/(C_1^* + C_2^*)$, where the first inequality is by part (i) and the second inequality holds because of the induced-penalty lower bound (see Lemma 3). To obtain the desired result, it is sufficient to prove the following stronger statement: for any $x_1, x_2 > 0$, $[x_2(\sqrt{1 + x_1/(\beta_1^* x_2)} - 1)]/(x_1 + x_2) \leq \alpha \equiv 1/(2(\beta_1^* + \sqrt{\beta_1^*}))$, which is equivalent to $\alpha^2(x_1/x_2)^2 + [2\alpha(1 + \alpha) - 1/\beta_1^*](x_1/x_2) + \alpha(2 + \alpha) \geq 0$. The result follows immediately by verifying that the quadratic function $f(x) = \alpha^2 x^2 + [2\alpha(1 + \alpha) - 1/\beta_1^*]x + \alpha(2 + \alpha)$ has a zero discriminant, i.e., $\Delta = 4\alpha^2 - 4\alpha(1 + \alpha)/\beta_1^* + 1/(\beta_1^*)^2 = 0$. \square

PROOF OF COROLLARY 1. It is easily verified that if $\beta_1^* = Q_2^*/Q_1^* \geq (1 + 1/\rho - \sqrt{1 + 2/\rho})/2 \equiv F(\rho)$, then $1 + 1/(2(\beta_1^* + \sqrt{\beta_1^*})) \leq 1 + \rho$, and by Theorem 3, $(1 + \rho)$ -optimality can be guaranteed. In particular, $F(\rho = 0.1) = (11 - \sqrt{21})/2 \simeq 3.21$, $F(\rho = 0.25) = 1.25$, $F(\rho = 0.5) = (3 - \sqrt{5})/2 \simeq 0.382$, $F(\rho = 1) = (2 - \sqrt{3})/2 \simeq 0.134$, and $F(\rho = 2) = (1.5 - \sqrt{2})/2 \simeq 0.0429$. \square

PROOF OF THEOREM 4. It follows from Theorem 3(ii) that to prove asymptotic optimality of the modified echelon $(\hat{\mathbf{r}}, \hat{\mathbf{Q}})$ policy, it is sufficient to show the following statements: (i) $\lim_{K_1/K_2 \rightarrow 0} \beta_1^* = \infty$; (ii) $\lim_{h_2/h_1 \rightarrow 0} \beta_1^* = \infty$.

(i) Let $r_i(Q_i) \equiv \arg \min_{r_i} C_i(r_i, Q_i)$. For $i = 1, 2$, define $A_i(Q_i) \equiv Q_i G_i(r_i(Q_i)) - \int_0^{Q_i} G_i(r_i(y)) dy$. By Lemma 1(v) and Zheng (1992), $A_i(Q_i)$ is a continuous and strictly increasing function such that $A_i(Q_i^*) = \lambda K_i$ and $A_i(0) = 0$. Let $A_i^{-1}(x)$ be the inverse function

of $A_i(Q_i)$. Then, $Q_i^* = A_i^{-1}(\lambda K_i)$ and $A_i^{-1}(0) = 0$. Let $\gamma \equiv K_1/K_2$ and we have

$$\begin{aligned} \lim_{K_1/K_2 \rightarrow 0} \beta_1^* &= \lim_{K_1/K_2 \rightarrow 0} \frac{A_2^{-1}(\lambda K_2)}{A_1^{-1}(\lambda K_1)} = \lim_{\gamma \rightarrow 0} \frac{A_2^{-1}(\lambda K_2)}{A_1^{-1}(\lambda \gamma K_2)} \\ &= \frac{A_2^{-1}(\lambda K_2)}{A_1^{-1}(\lambda K_2)} \lim_{\gamma \rightarrow 0} \frac{A_1^{-1}(\lambda K_2)}{A_1^{-1}(\lambda \gamma K_2)} = \infty, \end{aligned}$$

where the last equality follows from the fact that $A_i^{-1}(x)$ is continuous and $A_1^{-1}(0) = 0$.

(ii) Let $\zeta \equiv h_2/h_1$. Then, $G_1(y)$ can be written as $G_1(y) = \mathbb{E}[h_1(y - D_1)] + \mathbb{E}[(h_1 + \zeta h_1 + p)(y - D_1)^-]$. It follows that for any given $p > 0$ and $h_1 > 0$, Q_1^* converges to a finite constant, when $\zeta \rightarrow 0$. Next, it is sufficient to show that $\lim_{\zeta \rightarrow 0} \check{Q}_2^* = \infty$. Let $\tilde{H}_2(Q) \equiv ((h_2 p)/(h_2 + p))Q$, $A_2(Q) \equiv ((h_2 p)/(2(h_2 + p)))Q^2$, $\tilde{Q}_2 \equiv \sqrt{(2\lambda K_2(h_2 + p))/(h_2 p)}$. Then, it is easy to check that $\tilde{A}_2'(Q) = Q\tilde{H}_2'(Q)$ and $\tilde{A}_2(Q) = \lambda K_2$. Define $H_2(Q) = G_2(r_2(Q))$. Then, $\tilde{A}_2'(Q) = Q\tilde{H}_2'(Q)$. By Lemma 1(iv) and Lemma 3(i), $H_2'(Q) \leq (h_2 p)/(h_2 + p) = \tilde{H}_2'(Q)$, which implies that $\tilde{A}_2'(Q) \leq \tilde{A}_2'(Q)$. Note that $\tilde{A}_2(0) = \tilde{A}_2(0) = 0$. Therefore, we must have $\tilde{A}_2(Q) \leq \tilde{A}_2(Q)$. It then follows that $\tilde{A}_2(\tilde{Q}_2) \leq \tilde{A}_2(\tilde{Q}_2) = \lambda K_2$. Because $\tilde{A}_2(Q_2^*) = \lambda K_2$ and $\tilde{A}_2(Q)$ is an increasing function, we have $Q_2^* \geq \tilde{Q}_2$. On the other hand, because $K_2 > 0$, $\lim_{\zeta \rightarrow 0} \tilde{Q}_2 = \lim_{\zeta \rightarrow 0} \sqrt{(2\lambda K_2(\zeta h_1 + p))/(\zeta h_1 p)} = \infty$. Hence $\lim_{\zeta \rightarrow 0} \check{Q}_2^* = \infty$ and therefore, $\lim_{\zeta \rightarrow 0} \beta_1^* = \infty$. \square

PROOF OF LEMMA 7. It suffices to show that the inventory flows under the two policies are the same. Note that the initial inventory levels do not affect the long-run average cost. Thus, without loss of generality, we assume that the initial inventory levels under both policies satisfy $IP_2 = IP_1 = r_1$. This assumption implies that the sum of initial inventories in transit to and on hand at Stage 2 equals zero. Notice that given any demand sample path, if both policies have the same policy parameters r_2 and Q_2 , the shipment processes from the outside supplier to Stage 2 under both policies are the same, with the same amount Q_2 in each shipment.

Next we prove the following results by induction.

R1. Under both policies, the on-hand inventory level of Stage 2 satisfies $OI_2(t) = 0$ or Q_2 .

R2. Under both policies, the amount in each shipment from Stage 2 to Stage 1 is Q_2 .

R3. Under both policies, the inventory levels at both stages are the same.

The induction is conducted with respect to the arrival time of a shipment from the outside supplier to Stage 2. Let t_m be the arrival time of the m th shipment from the outside supplier to Stage 2. With the initial inventory levels, $OI_2(t) = 0$, $t \in [0, t_1)$, and $OI_2(t) = Q_2$, $t = t_1$, i.e., R1 holds just before time t_1 . The first shipment from the outside supplier to Stage 2 will reach Stage 1 just at time t_1 under both policies, i.e., R2 holds at time t_1 . For both stages, given any demand sample path, the flow-in and flow-out processes under both policies are the same before time t_1 , and so, therefore, are the inventory levels. That is, R3 holds before time t_1 . Now suppose all the results hold just before t_m .

We first prove $OI_2(t_m^-) = 0$ by contradiction. If $OI_2(t_m^-) \neq 0$, then by R1, $OI_2(t_m^-) = Q_2$. Denote by $ION_2(t)$ the sum of inventories in transit to and on hand at Stage 2 at time t . Clearly, $ION_2(t_m^-) \geq 2Q_2$. Note that $IP_2(t_m^-) \leq r_2 + Q_2$. Thus, $IP_1(t_m^-) = IP_2(t_m^-) - ION_2(t_m^-) \leq r_2 + Q_2 - 2Q_2 = r_2 - Q_2 \leq r_1$, where the last inequality follows from the policy stipulation in Lemma 7. This result implies that the $(m - 1)$ th shipment must be shipped before time t_m^- . Consequently, at time t_m , there is only the m th shipment on hand at Stage 2.

By R3, we know that all the inventory levels under both policies are the same before time t_m ; that implies that under both policies, IP_1 will fall to r_1 at the same time, denoted by \hat{t}_m . Note that \hat{t}_m could be t_m , which means that IP_1 has already been at or below r_1 , when the m th shipment arrives at Stage 2. In this case, it is clear that the m th shipment will be shipped to Stage 1 immediately under both policies. Now, consider the case $\hat{t}_m > t_m$. The shipment quantity should be $\min\{OI_2(\hat{t}_m), Q_1\} = \min\{Q_2, Q_1\} = Q_2$. Overall, under both policies, the m th shipment with Q_2 units will be shipped to Stage 1 at the same time before the $(m+1)$ th shipment arrives at Stage 2. As a result, given any demand sample path, the inventory levels at both stages are the same under both policies before time t_{m+1} . That is, R2 and R3 also hold before time t_{m+1} . Clearly, $OI_2(t) = Q_2$ when $t \in [t, \hat{t}_m)$, and $OI_2(t) = 0$ when $t \in [\hat{t}_m, t_{m+1})$; i.e., R1 holds before time t_{m+1} . Therefore, the induction is completed. \square

References

- Atkins D, Sun D (1995) 98%-effective lot sizing for series inventory systems with backlogging. *Oper. Res.* 43(2):335–345.
- Axsäter S (1990) Simple solution procedures for a class of two-echelon inventory problems. *Oper. Res.* 38(1):64–69.
- Axsäter S (1993) Exact and approximate evaluation of batch ordering policies of two-level inventory systems. *Oper. Res.* 41(4):777–785.
- Chan LA, Muriel A, Shen M, Simchi-Levi D, Teo C-P (2002) Effective zero-inventory-ordering policies for the single-warehouse multiretailer problem with piecewise linear cost structures. *Management Sci.* 48(11):1446–1460.
- Chao X, Zhou SX (2009) Optimal policy for a multi-echelon inventory system with batch ordering and fixed replenishment intervals. *Oper. Res.* 57(2):377–390.
- Chen F (1999) 94%-effective policies for a two-stage serial inventory system with stochastic demand. *Management Sci.* 45(12):1679–1696.
- Chen F (2000a) Optimal policies for multi-echelon inventory problems with batch ordering. *Oper. Res.* 48(3):376–389.
- Chen F (2000b) Effectiveness of (R, Q) policies in one-warehouse multiretailer systems with deterministic demand and backlogging. *Naval Res. Logist.* 47(5):422–439.
- Chen F, Zheng YS (1994a) Evaluating echelon stock (R, nQ) policies in serial production/inventory systems with stochastic demand. *Management Sci.* 40(10):1262–1275.
- Chen F, Zheng YS (1994b) Lower bounds for multi-echelon stochastic inventory systems. *Management Sci.* 40(11):1426–1443.
- Chen F, Zheng YS (1998) Near-optimal echelon-stock (R, nQ) policies in multistage serial systems. *Oper. Res.* 46(4):592–602.
- Chu Y, Shen M (2010) A power-of-two ordering policy for one-warehouse multiretailer systems with stochastic demand. *Oper. Res.* 58(2):492–502.
- Clark A, Scarf H (1960) Optimal policies for a multi-echelon inventory problem. *Management Sci.* 6(4):475–490.
- Clark A, Scarf H (1962) Approximation solutions to a simple multi-echelon inventory problem. Arrow KJ, Karlin S, Scarf H, eds. *Studies in Applied Probability and Management Science* (Stanford University Press, Stanford, CA), 88–110.
- De Bodt M, Graves S (1985) Continuous review policies for a multi-echelon inventory problem with stochastic demand. *Management Sci.* 31(10):1286–1295.
- Dong LX, Lee HL (2003) Optimal policies and approximations for a serial multi-echelon inventory system with time-correlated demand. *Oper. Res.* 51(6):969–980.
- Federgruen A, Zheng YS (1992) An efficient algorithm for computing an optimal (r, Q) policy in continuous review stochastic inventory systems. *Oper. Res.* 40(4):808–813.
- Federgruen A, Zipkin P (1984) Computational issues in an infinite horizon multi-echelon inventory model. *Oper. Res.* 32(4):818–836.
- Gallego G, Özer Ö (2003) Optimal replenishment policies for multiechelon inventory problems under advance demand information. *Manufacturing Service Oper. Management* 5(2):157–175.
- Gallego G, Özer Ö (2005) A new algorithm and a new heuristic for serial systems. *Oper. Res. Lett.* 33(4):349–362.
- Graves S (1985) A multi-echelon inventory model for a repairable item with one-for-one replenishment. *Management Sci.* 31(10):1247–1256.
- Jackson P (1988) Stock allocation in a two-echelon distribution system or “what to do until your ship comes in”. *Management Sci.* 34(7):880–895.
- Levi R, Roundy R, Truong VA (2006) Provably near-optimal balancing policies for multi-echelon stochastic inventory control models. Working paper, Cornell University, Ithaca, NY.
- Levi R, Roundy R, Shmoys DB, Sviridenko M (2008a) A constant approximation algorithm for the one-warehouse multi-retailer problem. *Management Sci.* 54(4):763–776.
- Levi R, Roundy R, Shmoys DB, Truong VA (2008b) Approximation algorithms for capacitated stochastic inventory control models. *Oper. Res.* 56(5):1184–1199.
- Moinedeh K, Lee HL (1986) Batch size and stocking levels in multi-echelon repairable systems. *Management Sci.* 32(12):1567–1581.
- Morse M (1959) Solution of a class of discrete-time inventory problems. *Oper. Res.* 7(1):67–78.
- Rosling K (1989) Optimal inventory policies for assembly systems under random demands. *Oper. Res.* 37(4):565–579.
- Roundy R (1985) 98%-effective integer-ratio lot-sizing for one-warehouse multi-retailer systems. *Management Sci.* 31(11):1416–1430.
- Shang K (2008) A simple heuristic for serial inventory systems with setup costs. *Oper. Res.* 56(4):1039–1043.
- Shang K, Song JS (2003) Newsvendor bounds and heuristic for optimal policies in serial supply chains. *Management Sci.* 49(5):618–638.
- Shang K, Zhou X (2010) Optimal and heuristic echelon (r, nQ, T) policies in serial inventory systems with fixed costs. *Oper. Res.* 58(2):414–427.
- Shang K, Song JS, Zipkin P (2009) Coordination mechanisms in decentralized serial inventory systems with batch ordering. *Management Sci.* 55(4):685–695.
- Svoronos A, Zipkin P (1988) Estimating the performance of multi-level inventory systems. *Oper. Res.* 36(1):57–72.
- Svoronos A, Zipkin P (1991) Evaluation of one-for-one replenishment policies for multi-echelon inventory systems. *Management Sci.* 37(1): 68–83.
- van Houtum G-J, Scheller-Wolf A, Yi J (2007) Optimal control of serial inventory systems with fixed replenishment intervals. *Oper. Res.* 55(4):674–687.
- Yang L, Yang J, Yu G, Zhang H (2011) Near-optimal (r, Q) policies for a two-stage serial inventory system with Poisson demand. *Internat. J. Production Econom.* 133:728–735.
- Zheng YS (1992) On properties of stochastic inventory systems. *Management Sci.* 38(1):87–103.
- Zipkin P (2000) *Foundations of Inventory Management* (McGraw-Hill, New York).

Ming Hu is an assistant professor of operations management at Rotman School of Management, University of Toronto. His research interests include revenue management, supply chain management and service operations management. Most recently, he studies operations management in the context of social buying, crowdfunding, crowdsourcing and two-sided markets.

Yi Yang is an associate professor in the Department of Management Science and Engineering Management at Zhejiang University. His main research interests include supply chain management, in particular inventory management, and revenue management.