M&SOM MANUFACTURING & SERVICE OPERATIONS MANAGEMENT
Vol. 14, No. 3, Summer 2012, pp. 414–422
ISSN 1523-4614 (print) | ISSN 1526-5498 (online)
informs
http://dx.doi.org/10.1287/msom.1110.0373

Single-Stage Approximations for Optimal Policies in Serial Inventory Systems with Nonstationary Demand

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

Companies often face nonstationary demand due to product life cycles and seasonality, and nonstationary demand complicates supply chain managers' inventory decisions. This paper proposes a simple heuristic for determining stocking levels in a serial inventory system. Unlike the exact optimization algorithm, the heuristic generates a near-optimal solution by solving a series of independent single-stage systems. The heuristic is constructed based on three results we derive. First, we provide a new cost decomposition scheme based on echelon systems. Next, we show that the optimal base-stock level for each echelon system is bounded by those of two revised echelon systems. Last, we prove that the revised echelon systems are essentially equivalent to single-stage systems. We examine the myopic solution for these single-stage systems. In a numerical study, we find that the change of direction of the myopic solution is consistent with that of the optimal solution when system parameters vary. We then derive an analytical expression for the myopic solution and use it to gain insights into how to manage inventory. The analytical expression shows how future demand affects the current optimal local base-stock level; it also explains an observation that the safety stock at an upstream stage is often stable and may not increase when the demand variability increases over time. Finally, we discuss how the heuristic leads to a time-consistent coordination scheme that enables a decentralized supply chain to achieve the heuristic solution.

Key words: multiechelon; single-stage heuristic; nonstationary demand; myopic solution
History: Received: November 2, 2010; accepted: November 23, 2011. Published online in *Articles in Advance* April 13, 2012.

1. Introduction

Customer demand is often nonstationary in practice. Causes of nonstationary demand include product life cycles, seasonality, trends, and economic conditions. Nonstationary demand poses challenges for supply chain inventory managers. First, finding an optimal systemwide solution often requires solving interrelated, recursive cost functions between stages across time. An optimization algorithm, if it existed, is usually hard to understand and execute. Second, lacking a simple and intuitive solution approach makes the system less transparent. It is therefore difficult for the managers to foresee the impact of changes of the environment and react to them. Third, a supply chain is usually composed of self-interested firms. These firms may not be willing to implement the system optimal solution. In a nonstationary demand environment, the optimal solution is often time varying. It would seem difficult to design a simple incentive scheme that can induce each location to choose the optimal stocking level in each time period.

This paper proposes a simple heuristic that aims to resolve the above challenges. We consider a serial inventory system in a finite horizon. The system has N stages and materials flow from stage N to stage $N-1$, $N-1$ to $N-2$, etc. until stage 1, where a random, nonstationary demand occurs in each period. This model was first studied by Clark and Scarf (1960) who show that (time-varying) echelon base-stock policies are optimal. (Echelon j is a subsystem that includes stage j and all of its downstream stages.) Although the structure of the policy is simple, obtaining the optimal solution is quite complex because the optimal value function of an upstream stage depends on the optimal base-stock levels of its downstream stages. Clearly, the calculation becomes more cumbersome as N increases.

The heuristic we propose breaks down the dependance between stages. That is, it can generate an effective echelon base-stock level for stage j ($2 \le j \le N$) without knowing the base-stock level of stage i ($i < j$). More specifically, we show that the optimal echelon base-stock level for stage j is bounded by the optimal solutions of two single-stage systems. To establish this result, we first propose a cost decomposition scheme based on the echelon system. Then, we show that the optimal value function for echelon j is bounded above and below by that of a revised j-stage system.

We refer to these revised systems as the upper-bound system and the lower-bound system, respectively. The upper-bound system is constructed by requiring stage i ($< j$) to always order up to stage $i+1$'s echelon inventory level in each period. On the other hand, the lower-bound system is constructed by regulating stage i's ($i < j$) holding and order cost parameters. We further show that the optimal base-stock level for the upper-bound (lower-bound) system is a lower (upper) bound to that of the original echelon j system. Lastly, we show that solving these revised j-stage systems is equivalent to solving a single-stage system whose parameters are obtained from the original system. This result motivates us to propose a heuristic solution for each echelon by solving a single-stage system with a weighted average of the cost parameters obtained from the upper- and lower-bound systems. We provide a definitive guidance on choosing an effective weight based on a cost ratio, representing the service level of the system.

The above single-stage heuristic provides an approach to resolve the aforementioned challenges. The heuristic is very easy to understand and execute: it can generate an effective solution by solving N independent single-stage problems. This separation feature not only simplifies the computation but also shortens the computation time by allowing each stage to solve its own problem in parallel. (In other words, the heuristic can generate a solution at least N times faster than the exact algorithm, provided that parallel processing is possible.) To make the system more transparent, we investigate the myopic solution for the heuristic single-stage system. We find that the change of the myopic solution is consistent with that of the optimal echelon base-stock level when the system parameters vary. We then derive an analytical expression for the myopic solution to approximate the optimal local base-stock level. The expression shows how the system parameters affect the local base-stock level and safety stock. It also shows that the safety stock of an upstream stage is often stable and may not increase when the variance of the demand increases over time. Finally, our heuristic can lead to a remarkably simple, time-consistent contract that induces each stage to choose the heuristic solution in a decentralized supply chain. We refer the reader to §5.2 for details.

Several researchers have provided methods to simplify the computation for the Clark and Scarf (1960) model. Federgruen and Zipkin (1984) consider an infinite-horizon version of the model with independent and identically distributed demand and show that the optimal policy can be obtained by recursively solving two functional equations that have the form of a single-period problem. Chen and Zheng (1994) reinterpret Federgruen's and Zipkin's (1984) results, simplify the optimality proof, and present an optimization algorithm to facilitate the computation. Gallego and Özer (2003) consider an advanced demand information model and show the optimality of a myopic solution. Although the computation effort is much reduced, finding an upstream solution still is not easy because it depends on the downstream solutions. Thus, there is a stream of research that aims to further simplify the computation by solving independent single-stage problems. Noteworthy examples include Dong and Lee (2003), Shang and Song (2003), Gallego and Özer (2005), and Chao and Zhou (2007). Our paper can be viewed as a generalization of Shang and Song (2003) by considering the system with nonstationary demand.

Several papers have derived solutions for practical issues in supply chains under nonstationary demand. Erkip et al. (1990) consider a one-depot-multiwarehouse system in which the warehouses' demands are correlated. They derive an expression for the optimal safety stock as a function of the level of correlation through time. Ettl et al. (2000) consider a supply chain network that implements base-stock policies subject to service level requirements. They approximate the lead time demand for each location and suggest a rolling-horizon approach to find the base-stock levels for the nonstationary demand case. Abhyankar and Graves (2001) consider a two-stage serial system with a Markov-modulated Poisson demand process; they implement an inventory hedging policy to protect against cyclic demand variability. Graves and Willems (2008) consider a problem of allocating safety stocks in a supply chain network where the demand is bounded and there is a guaranteed service time between stages and customers; they propose an algorithm to determine safety stocks under a constant service time policy. Schoenmeyr and Graves (2009) examine the placement of safety stocks in a supply chain with an evolving demand forecast; they show that the algorithm developed in Graves and Willems (2000) can be applied to find the optimal safety stocks and that the system inventory level can be substantially reduced as the forecast improves over time. Similar to the above papers, the present work aims to provide a simple control policy.

Finally, our paper is also related to the coordination literature. Most coordination papers consider an infinite-horizon model with stationary demand. Because of the regenerative process, these infinite-horizon models are equivalent to single-cycle problems. These coordination papers often analyze a decentralized Nash equilibrium solution and provide contracts to induce the system to achieve the centralized (first best) solution, e.g., Lee and Whang (1999), Chen (1999), Cachon and Zipkin (1999), and Shang et al. (2009). However, when the system fails to form a regenerative process, studying the decentralized behaviors becomes more difficult. Donohue (2000) studies a two-period model with demand forecasts. She suggests using time-varying contract terms to coordinate the system. Parker and Kapuscinski (2011) consider a two-stage serial inventory system with capacity limits, where each stage aims to minimize its own costs. They show that there exists a Markov equilibrium policy for a dynamic game in the decentralized control system. In general, it is very difficult to derive a coordination contract in a finite-horizon model. Were such a contract to exist, it would be too difficult to implement because the contract terms are often time varying.

2. The Model and Echelon Cost Decomposition Scheme

We consider an N-stage serial inventory system, where stage 1 orders from stage 2, stage 2 from stage 3, etc., until stage N, which orders from an ample outside supplier. There is a lead time τ_j between stage j and stage $j+1$, and τ_j is a positive integer. Denote $\tau[i,j]=\sum_{k=i}^{j}\tau_k$ and $\tau[i,j]=0$ if $i>j$. Let h_j be the echelon holding cost rate at stage j and let b be the backorder cost rate at stage 1. Let $h[i,j]=\sum_{k=i}^{j}h_k$, and $h[i,j]=0$ if $i>j$. Define p_j as the unit order cost for stage j. We use t to index the time period and count the time backward. Let T be the planning horizon. Denote $D(t)$ the demand in period t. The demands are independent between periods, but the demand distributions may differ from period to period. Let $D[t,s]=\sum_{i=s}^{t}D(i)$, representing the total demand in period t, $t-1$, $t-2,\ldots,s$, where $t\geq s$.

The sequence of events in each period is as follows: At the beginning of a period, each stage j (1) receives a shipment sent τ_j periods ago from stage $j+1$; (2) receives an order from stage $j-1$; (3) places an order to stage $j+1$; and (4) sends a shipment to stage $j-1$. (Stage 1 skips events 2 and 4, whereas stage N orders from an outside ample source.) Stage 1 decides its order first, followed by stage 2, and so on, until stage N. The shipments are made in the opposite order, starting at stage N, then stage $N-1$, etc., until stage 1. After orders and shipments, demand occurs during the period. Inventory holding and backorder costs are assessed at the end of the period.

Clark and Scarf (1960) show that time-varying, echelon base-stock policies are optimal for the above model. Let the optimal echelon base-stock level be $s_j(t)$ for stage j in period t. To illustrate the policy, let us define the following inventory variables. For stage j at the beginning of period t, define

$x_j(t)$ = echelon inventory level after a shipment is received (after event (1)),
= on-hand inventory at stage j + inventory in transit to and held at stage i ($<j$) − backorders at stage 1,

$v_j(t)$ = echelon inventory position before an order is placed (before event (3)),
= inventory in transit to stage $j+x_j(t)$,

$y_j(t)$ = echelon inventory position for stage j after an order is placed (after event (3)).

There are two different notions of inventory positions in the literature, the echelon inventory *order* position (= inventory on order for stage $j+x_j(t)$) and the echelon inventory position $v_j(t)$. The difference between these two is the number of outstanding orders for stage j. The echelon inventory order position has been used to define the policy: Stage j reviews its echelon inventory order position at the beginning of each period. The stage orders up to $s_j(t)$ if the echelon inventory order position is less than $s_j(t)$, and does not order otherwise. This policy is equivalent to the following shipment scheme: If $v_j(t)$ is less than $s_j(t)$ and stage $j+1$ has positive on-hand inventory, stage $j+1$ sends a shipment to stage j to raise $v_j(t)$ to $s_j(t)$, if possible, or as close as possible to $s_j(t)$. Otherwise, no shipment is made. In other words, if stage $j+1$ sends a shipment after stage j places an order, the post-ordering echelon inventory position for stage j is $y_j(t)=\min\{s_j(t),x_{j+1}(t)\}$.

The optimal base-stock levels $s_j(t)$ can be found by solving N sets of functional equations sequentially. More specifically, finding $s_1(t)$ is equivalent to solving a single-stage system. With the known $s_1(t)$, one can compute an induced-penalty cost charged to stage 2. (The induced-penalty cost is a penalty charged to an upstream stage for not fulfilling the downstream order.) The optimal base-stock level $s_2(t)$ is the optimal solution obtained from the functional equation formulated from stage 2's cost. Continuing this procedure, with the known $s_i(t)$, $i<j$, one can compute the induced-penalty cost charged to stage j and find the corresponding optimal solution $s_j(t)$, $j=3,\ldots,N$. Although Clark and Scarf's (1960) algorithm significantly simplifies the computation by converting the original N-dimension problem into solving a series of N single-dimension problems, the computation is still quite involved because the upstream base-stock level depends on all of its downstream ones.

In this paper, we propose a heuristic for the optimal echelon base-stock levels by solving N independent single-stage systems. What we achieved is to show that the optimal cost and solution of an echelon system are bounded below and above by those of a single-stage system, respectively. To facilitate our analysis, we need to show the cost for each echelon by revising the original Clark and Scarf's (1960) cost decomposition scheme. We demonstrate the idea in a two-stage system with $\tau_1 = \tau_2 = 1$. A similar but more tedious analysis can be carried out for the general system. Notice that when $\tau_j = 1$, $v_j(t)$ and $x_j(t)$ are equal.

Define $L_j(x_j, t)$ the inventory cost incurred for stage j in period t when the echelon inventory level is x_j, namely,

$$L_1(x_1, t) = \mathsf{E}[h_1(x_1 - D(t)) + (b + h_1 + h_2)(x_1 - D(t))^-],$$
$$L_2(x_2, t) = \mathsf{E}[h_2(x_2 - D(t))],$$

where $(x)^- = \max\{0, -x\}$. The total inventory holding and backorder cost incurred in period t is then $L_1(x_1, t) + L_2(x_2, t)$. Let $f_t(x_1, x_2)$ be the optimal total discounted cost for the system with initial echelon inventory levels (x_1, x_2) when there are t periods to go. The dynamic program based on the echelon system for the Clark and Scarf (1960) model is as follows: Let $f_0(x_1, x_2) = 0$. For $t \geq 1$,

$$\begin{aligned} f_t(x_1, x_2) &= \min_{x_1 \le y_1 \le x_2 \le y_2} \{p_1(y_1 - x_1) + p_2(y_2 - x_2) + L_1(x_1, t) \\ &\qquad + L_2(x_2, t) + \alpha \mathsf{E}[f_{t-1}(y_1 - D(t), y_2 - D(t))]\} \\ &= C_t(x_1) + G_t(x_2), \end{aligned}$$

where

$$C_t(x_1) = L_1(x_1, t) - p_1 x_1 + [U_t(\max\{x_1, s_1(t)\}) - U_t(s_1(t))], \tag{1}$$
$$G_t(x_2) = U_t(\min\{x_2, s_1(t)\}) + L_2(x_2, t) - p_2 x_2 + V_t(\max\{x_2, s_2(t)\}), \tag{2}$$
$$U_t(y_1) = p_1 y_1 + \alpha \mathsf{E}[C_{t-1}(y_1 - D(t))],$$
$$V_t(y_2) = p_2 y_2 + \alpha \mathsf{E}[G_{t-1}(y_2 - D(t))],$$
$$s_1(t) = \arg\min_{y_1} \{U_t(y_1)\},$$
$$s_2(t) = \arg\min_{y_2} \{V_t(y_2)\}.$$

Here, α is the discount rate. Note that Equations (1) and (2) should be revised for $t = 1$ and $t = 2$. When $t = 1$, both stages will not order, so $C_1(x_1) = L_1(x_1, 1)$ and $G_1(x_2) = L_2(x_2, 1)$; when $t = 2$, stage 2 will not order, so $G_2(x_2)$ is the same as (2) except that the last term is changed to $V_2(x_2)$.

Under the above cost decomposition scheme, the total system cost $f_t(x_1, x_2)$ consists of two echelon cost functions. The $G_t(x_2)$ function is the cost for echelon 2, which includes the costs *directly and indirectly* determined by x_2, assuming that stage 1 will always order up to its optimal base-stock level in each period. The $C_t(x_1)$ function includes the remaining costs *directly* determined by x_1.

3. Bounds for the Echelon Cost and Solution

We shall demonstrate that $s_j(t)$, $j = 1, \ldots, N$, is bounded by the solutions obtained from two single-stage systems. Let $\mathbb{S}_j^u[p_j^u, h_j^u, b_j, \mathfrak{T}_j]$ denote the upper-bound system for echelon j, where p_j^u, h_j^u, b_j, and $\mathfrak{T}_j$ denote the unit order cost, holding cost rate, backorder cost rate, and the lead time, respectively. Similarly, let $\mathbb{S}_j^l[p_j^l, h_j^l, b_j, \mathfrak{T}_j]$ denote the lower-bound system for echelon j. As stated in §2, finding the optimal base-stock level $s_1(t)$ is the same as solving a single-stage system. That is, for $j = 1$, the upper-bound system and lower-bound system have the same parameters: $p_1^u = p_1^l = p_1$, $h_1^u = h_1^l = h_1$, $b_1 = b + h[2, N]$, and $\mathfrak{T}_1 = \tau[1, 1]$. Below we derive the parameters of the upper-bound and lower-bound systems for echelon $j \geq 2$.

3.1. Upper-Bound System

Consider echelon j with a more restrictive policy: stage i *always* orders up to x_{i+1} in each period except $t \leq \tau[1, i]$ for $i < j$. (When $t \leq \tau[1, i]$, stage i would not order because of the end of the horizon.) Let $s_j^l(t)$ be the resulting optimal echelon base-stock level for echelon j. Clearly, such a policy is suboptimal and the resulting echelon cost is an upper bound to that of the original system. For this reason, we call this restrictive system the upper-bound system. In Online Appendix A (available at http://msom.journal.informs.org/), we show the resulting dynamic program formulation under this suboptimal policy for the above two-stage system example.

Under this policy, the resulting $y_j(t)$ is $x_{j+1}(t)$ instead of $\min\{s_j(t), x_{j+1}(t)\}$ in the original system. It is intuitive that the solution $s_j^l(t)$ is a lower bound to $s_j(t)$ because more stocks will be shipped to the downstream stages in a period. Thus, the optimal echelon base-stock level in the upper-bound system should be lower. Theorem 1 confirms this intuition. (All proofs are available in Online Appendix B.)

Theorem 1. *$s_j(t) \geq s_j^l(t)$ for $t > \tau[1, j]$, provided that $v_i(t) < s_i(t)$ for $i < j$.*

The condition in Theorem 1 is to ensure that the downstream stage i in the original echelon system will place an order in each period.

Let us take a closer look at the upper-bound system. To construct the upper-bound system for echelon j, we regulate stage i ($< j$) to always order up to x_{i+1} in each period. By doing so, any unit ordered by stage j will eventually arrive at stage 1 in $\tau[1, j]$ periods. Thus, we can view that each stage $i, i = 2, \ldots, j$ as a transit point and that the upper-bound system is effectively a single-stage system with a lead time of $\mathfrak{T}_j = \tau[1, j]$ periods. The holding cost and the backorder cost at stage 1 in echelon j are $h_j^u = h[1, j]$ and

$b_j = b + h[j+1, N]$, respectively. (See Proposition 2 of Shang and Song 2003 for an explanation of how these cost parameters are derived.)

It is more complicated to compute the unit order cost for the upper-bound system. In a single-stage system, we do not consider the inventory holding cost before a unit arrives at the stage. (That is, the pipeline inventory holding cost is "external" to the single-stage system.) However, this is not the case for the upper-bound system. Any unit ordered by stage j will incur an order cost when it arrives at each of its downstream stages and a holding cost in each period before it arrives at stage 1. Thus, the unit order cost for stage j in the upper-bound system is the total cost incurred by a unit when it is ordered until it arrives at stage 1.

More specifically, for each unit ordered by stage j in the current period t, p_j is incurred. This unit will arrive at stage j at the beginning of period $t - \tau_j$. Then, stage $j-1$ has to order it with a cost of p_{j-1} to ensure that it continues to move to stage $j-1$. In addition, this unit will incur a holding cost h_j in each period before it arrives at stage $j-1$. Thus, the cost (by considering the discount effect) incurred by this unit when it travels from stage j to $j-1$ is $p_{j-1}(\alpha^{\tau_j}) + h_j(\sum_{i=1}^{\tau_{j-1}} \alpha^{\tau_j+i-1})$. Continuing this logic, when this unit arrives at stage 2 at the beginning of period $t-\tau[2, j]$, stage 1 has to order it, incurring an order cost of p_1; this unit will incur a holding cost of $h[2, j]$ in each period before it arrives at stage 1. The cost incurred by this unit is $p_1(\alpha^{\tau[2,j]}) + h[2, j](\sum_{i=1}^{\tau_1} \alpha^{\tau[2,j]+i-1})$. The sum of these order costs and holding costs incurred in transit will be the unit order cost for the upper-bound system, denoted by p_j^u, where

$$p_j^u = p_j + \sum_{k=2}^{j}\left(p_{k-1}(\alpha^{\tau[k,j]}) + h[k,j]\left(\sum_{i=1}^{\tau_{k-1}} \alpha^{\tau[k,j]+i-1}\right)\right).$$

Proposition 1. *The lower-bound solution $s_j^l(t)$ can be obtained by solving $\mathbb{S}_j^u[p_j^u, h_j^u, b_j, \mathfrak{T}_j]$, for $j = 2, \ldots, N$.*

3.2. Lower-Bound System

The approach of constructing the lower-bound system for echelon j is different: We set $h_i = 0$ and $p_i = 0$ for $i < j$. Clearly, the resulting cost is a lower bound to the echelon j's cost. Under this construction, stage i would order up to x_{i+1} in each period as there is no benefit of carrying inventory at stage $i+1$. Consequently, the optimal order policy is the same as that for the upper-bound system. In Online Appendix A, we continue the two-stage example and show the corresponding dynamic program formulation for the lower-bound system.

Unlike the upper-bound system, it is not clear whether an order relationship exists between the resulting solution, $s_j^u(t)$, and the optimal solution $s_j(t)$. This is because setting downstream cost parameters equal to zero makes the echelon to stock more. On the other hand, more inventory pushed to the downstream stages makes the echelon stock less. Thus, it is not clear about the joint effect. Theorem 2 shows that an order relationship does hold.

Theorem 2. *$s_j^u(t) \geq s_j(t)$, for $t > \tau[1, j]$.*

Because the downstream stages use the same optimal policy, the lower-bound system is equivalent to a single-stage system. We can simply set $h_i = 0$ and $p_i = 0$ for $i < j$ in the upper-bound system to obtain the cost parameters for the lower-bound system. That is, $p_j^l = p_j + h_j(\sum_{i=\tau_j}^{\tau[1,j]-1} \alpha^i)$ and the holding cost rate $h_j^l = h_j$.

Proposition 2. *The upper-bound solution $s_j^u(t)$ can be obtained by solving $\mathbb{S}_j^l[p_j^l, h_j^l, b_j, \mathfrak{T}_j]$, for $j = 2, \ldots, N$.*

3.3. Single-Stage Heuristic

We suggest a heuristic that solves a single-stage system with a weighted average of the cost parameters obtained from the upper- and lower-bound systems. Specifically, for $j \geq 1$, let $p_j^a = wp_j^u + (1-w)p_j^l$ and $h_j^a = wh_j^u + (1-w)h_j^l$, where $0 \leq w \leq 1$. We call the resulting single-stage system *heuristic system* j, denoted by $\mathbb{S}_j^a(p_j^a, h_j^a, b_j, \mathfrak{T}_j)$ and define the resulting optimal solution as $s_j^a(t)$. In §4, we shall provide a guidance for choosing an effective w based on the cost parameters.

4. Numerical Study

The goal of the numerical study is to test the effectiveness of the heuristic and provide a definitive guidance to choose the weight w for the heuristic. We consider two-stage systems with a horizon of $T = 10$ periods. Assume that the period demand follows a Poisson distribution with rate $\lambda(t)$ in period t. We test the following demand patterns: constant (C), linear increasing (I), linear decreasing (D), concave (V), and convex (X) forms. For the constant demand, we set $\lambda(t) = 5.5$, for $1 \leq t \leq 10$; for the increasing demand, $\lambda(t) = 11 - t$, for $1 \leq t \leq 10$; for the decreasing demand, $\lambda(t) = t$, for $1 \leq t \leq 10$; for the convex demand, $\lambda(t) = 11 - 2t$ for $1 \leq t \leq 5$ and $\lambda(t) = 2t - 10$ for $6 \leq t \leq 10$; finally, for the concave demand, $\lambda(t) = 2t - 1$ for $1 \leq t \leq 5$, and $\lambda(t) = 22 - 2t$ for $6 \leq t \leq 10$.

We first fix $h_1 = 1$, $p_1 = 4$ and change the cost parameters at stage 2: $h_2 \in \{0.5, 1, 1.5\}$, $p_2 \in \{2, 6\}$. We then swap the stage index in the above order and holding cost parameters to generate another set of instances. The other parameters are $\tau_1 = \tau_2 = 1$ and $b \in \{15, 50\}$. The total number of instances is 120 and the total number of optimal base-stock levels for stage 2 is 960. The discount rate is $\alpha = 0.95$ for all instances.

Table 1 Number of the Heuristic Solutions Equal to the Optimal Solutions Under Different Backorder Cost Rates and Weights

w	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0	Max
$b = 50$												
#	88	150	217	296	360	402	411	393	346	303	250	480
$\bar{\epsilon}$ (%)	4.14	3.27	2.53	1.83	1.21	0.81	0.73	0.86	1.35	1.74	1.97	—
$b = 15$												
#	32	63	96	150	233	302	357	396	397	351	280	480
$\bar{\epsilon}$ (%)	8.05	6.78	5.64	4.63	3.59	2.52	1.68	1.13	1.04	1.54	2.53	—
Total												
#	120	213	313	446	593	704	768	789	743	654	530	960
$\bar{\epsilon}$ (%)	6.10	5.03	4.09	3.23	2.40	1.67	1.21	1.00	1.20	1.64	2.25	—

To test the effectiveness of the heuristic under different weights, we test $w \in \{0.0, 0.1, \ldots, 0.9, 1.0\}$. Note that $w = 0.0$ and $w = 1.0$ correspond to the $s_2^u(t)$ and $s_2^l(t)$, respectively. We define

$$\epsilon = \frac{|s_2(t) - s_2^a(t)| \times 100\%}{s_2(t)}$$

to represent the percentage error between the optimal solution and the heuristic solution. Table 1 shows the number of optimal solutions (denoted by #) generated by the heuristic and the average ϵ (denoted by $\bar{\epsilon}$) under different ws.

We have some observations. First, the optimal solution $s_2(t)$ tends to be closer to $s_j^l(t)$: there are a total of 530 (120) $s_2^l(t)$ ($s_2^u(t)$) solutions that coincide the optimal solution $s_2(t)$. The performance of the heuristic under different weights is more robust when b is large. This is because when b is large, the gap between $s_2^u(t)$ and $s_j^l(t)$ is smaller. When $b = 50$ ($b = 15$), the best w is 0.6 (0.8), and the heuristic generates 411 (397) optimal solutions. This observation indicates that when b increases, one should choose a smaller weight, making the heuristic solution leaning toward $s_2^u(t)$. The choice of w does not seem critical as long as w falls in some range. For example, when $b = 50$ ($b = 15$), using $w \in [0.4, 0.8]$ ($w \in [0.6, 0.9]$) generates at least 72% of the optimal solutions. We do not observe a significant difference for $\bar{\epsilon}$ between the demand forms for a given w. Together, these observations suggest that the choice of w is closely related to the backorder cost rate, or equivalently, the system's service level. We therefore suggest inferring w through the critical fractile $b/(b + h[1, N])$.[1] In our test bed, the critical ratio among all instances ranges from 0.85 to 0.975. For each instance with a particular critical fractile, we then find the best w. Table 2 is a summary of showing the relationship between the best w and the ratio. This table can be used as a guidance to choose an effective w.

To test the effectiveness of the heuristic solution under the guidance set in Table 2, we consider a four-stage system with $T = 20$. We fix $h[1, 4] = 1$ and consider four different holding cost forms: $(h_1, h_2, h_3, h_4) \in \{(0.25, 0.25, 0.25, 0.25), (0.4, 0.2, 0.2, 0.2), (0.1, 0.1, 0.4, 0.4), (0.2, 0.2, 0.4, 0.2)\}$, representing linear, affine, kink, and jump forms, respectively. Similarly, we fix $\tau[1, 4] = 6$ and let $(\tau_1, \tau_2, \tau_3, \tau_4) \in \{(2, 2, 1, 1), (1, 2, 2, 1), (1, 1, 2, 2)\}$, representing long lead times at the downstream, middle, and upstream stages, respectively. For the order costs, we consider two scenarios: $(p_1, p_2, p_3, p_4) \in \{(1, 1, 2, 2), (2, 2, 1, 1)\}$, representing high order costs at the upstream and downstream stages, respectively. For the backorder cost, let $b \in \{15, 50\}$. Finally, for the demand form, we consider both convex demand $(D(1), D(2), \ldots, D(20)) = (10, 9, \ldots, 1, 1, 2, \ldots, 10)$ and concave demand $(D(1), D(2), \ldots, D(20)) = (1, 2, \ldots, 10, 10, 9, \ldots, 1)$. There are a total of 96 instances, with 1,632 optimal solutions at stage 2, 1,472 at stage 3, and 1,344 at stage 4. Based on Table 2, we set $w = 0.5$ for $b = 50$ (with $b/(b + h[1, 4]) = 0.9804$) and $w = 0.7$ for $b = 15$ (with $b/(b + h[1, 4]) = 0.9375$). Table 3 summarizes the average percentage error $\bar{\epsilon}$ at each stage. The result indicates that the heuristic is effective when the number of stages increases.

5. Applications of the Heuristic

This section demonstrates the usefulness of the single-stage heuristic. For managers, it is crucial to learn how the system parameters affect the stocking decision in a supply chain. Section 5.1 provides a simple analytical expression to approximate the optimal local base-stock level and the safety stock at each stage. Section 5.2 considers a decentralized supply chain. As we shall demonstrate, the heuristic leads to a simple, time-consistent contract that enables the supply chain to achieve the heuristic solution.

[1] We may use a more refined critical fractile derived from the myopic solution of the upper-bound system, for example, $(\alpha^{\tilde{\tau}_j} b_j - p_j^u(1 - \alpha))/(\alpha^{\tilde{\tau}_j}(b_j + h_j^u))$, to infer the weight for each stage j. However, we find such a refined approach does not significantly improve the effectiveness of the heuristic as it is quite robust for w in some range. Therefore, for easiness of use, we suggest to infer w by the ratio $b/(b + h[1, N])$.

Table 2 Selecting an Effective Weight Based on the Cost Ratio

$b/(b+h[1,N])$	(0, 0.85]	(0.85, 0.925]	(0.925, 0.95]	(0.95, 0.975]	(0.975, 0.99]	(0.99, 1)
w	0.9	0.8	0.7	0.6	0.5	0.4

Table 3 Average Percentage Error $\bar{\epsilon}$ for the Heuristic Solution $s_j^a(t)$ in the Four-Stage System

Backorder cost rate	Stage 2 (%)	Stage 3 (%)	Stage 4 (%)
$b=15$	1.00	1.60	1.71
$b=50$	0.68	1.14	1.31

5.1. Approximating the Optimal Base-Stock Level

We propose using the myopic solution to approximate the optimal base-stock level. It is well known that the myopic solution is an upper bound to the optimal solution for a single-stage system (Zipkin 2000, pp. 378–379). In our numerical study, we find that in all cases the myopic solution $s_j^m(t)$ moves in the same direction as the optimal solution $s_j(t)$ when the system parameters change. (Table 4 is an example when demand has a concave form.) This result motivates us to use $s_j^m(t)$ as an approximation for $s_j^a(t)$. In addition, for any echelon base-stock policy, there exists an equivalent local base-stock policy, and vice versa (Zipkin 2000, p. 306); see Equation (3) below. With these two results, we can use the myopic solution to derive an approximation for the optimal local base-stock level. Below we lay out the detailed steps.

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

Proposition 3. *For $t > \mathfrak{T}_j$,*

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

where

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

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

To obtain a simple analytical expression for $s_j^m(t)$, we apply normal approximation on $D[t, t-\mathfrak{T}_j]$. Let the mean of $D[t, t-\mathfrak{T}_j]$ be $\lambda[t, t-\mathfrak{T}_j]$ and the standard deviation be $\sigma[t, t-\mathfrak{T}_j] = \sqrt{\mathsf{Var}[D[t, t-\mathfrak{T}_j]]}$. Then,

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

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

$$\begin{aligned} s_j^{m'}(t) &= s_j^m(t) - s_{j-1}^m(t) \\ &= \lambda[t-\mathfrak{T}_{j-1}-1, t-\mathfrak{T}_j] + \sigma[t, t-\mathfrak{T}_j]\Phi^{-1}(\beta_j) \\ &\quad - \sigma[t, t-\mathfrak{T}_{j-1}]\Phi^{-1}(\beta_{j-1}). \end{aligned} \tag{3}$$

The first term in Equation (3) is the average pipeline inventory in period t, which depends on the average τ_j periods of *future* demand in $[t-\mathfrak{T}_{j-1}-1, t-\mathfrak{T}_j]$. The second term is the safety stock for stage j in period t, denoted by $ss_j^m(t)$, which depends on the cost ratios β_j and β_{j-1}, and the variability of the demand in $[t, t-\mathfrak{T}_j]$.

Equation (3) allows us to analytically investigate how the system parameters affect the optimal base-stock level and the safety stock at each stage. For example, if we are interested in the change to the amount of safety stock of the upstream stage in a two-stage system, we can define the change to stage 2's safety stock in period t as

$$\begin{aligned} &\Delta ss_2^m(t) \\ &= ss_2^m(t-1) - ss_2^m(t) \\ &= (\sigma[t-1, t-1-\mathfrak{T}_2] - \sigma[t, t-\mathfrak{T}_2])\Phi^{-1}(\beta_2) \\ &\quad - (\sigma[t-1, t-1-\mathfrak{T}_1] - \sigma[t, t-\mathfrak{T}_1])\Phi^{-1}(\beta_1). \end{aligned} \tag{4}$$

From the above equation, we can see that $\Delta ss_2^m(t)$ will be fairly small unless there is a significant difference between $\mathsf{Var}[D(t-1-\mathfrak{T}_2)]$ and $\mathsf{Var}[D(t-1-\mathfrak{T}_1)]$. This implies that the safety stock at the upstream stage should be fairly stable. (A similar conclusion is numerically observed in Graves and Willems 2008.) In addition, $\Delta ss_2^m(t)$ may not be positive even if $\mathsf{Var}[D(t)] < \mathsf{Var}[D(t-1)]$, $\forall\, t$. More specifically, when either p_2 is large, or h_2 is large, or b is small, $\Phi^{-1}(\beta_2)$ tends to be smaller than $\Phi^{-1}(\beta_1)$, causing the difference in (4) to become negative even when the demand variance increases over time. This suggests that the safety stock at an upstream stage may not increase with the demand variability.

Example 1. We consider a two-stage system with $\tau_1 = \tau_2 = 1$, $p_2 = 6$, $h_2 = 1$, $p_1 = 4$, $h_1 = 1$, $b = 15$, and $\alpha = 0.95$. The demand follows a Poisson distribution with mean rate shown in Table 4. We report the optimal echelon, heuristic, and myopic base-stock levels in each period. We also report the optimal local base-stock level as well as the corresponding safety stock for stage 2. As shown, the change of the

Table 4 Two-Stage Example with the Optimal, Heuristic, and Myopic Solutions, and the Local Base-Stock Levels and the Resulting Safety Stocks

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

myopic solution is consistent with that of the optimal solution; the safety stock may decrease although the demand rate increases (e.g., from $t = 10$ to $t = 9$). Online Appendix C summarizes the performance of the myopic solution based on the same 120 two-stage instances in §4.

5.2. Coordination Mechanism

In reality, a supply chain is often composed of independent firms, each pursuing its own best interest. Because the centralized solution is often time varying, a coordination contract, if it existed, would have nonstationary parameters. The nonstationary contract terms would make the implementation difficult. Following a similar idea of Shang et al. (2009), we show that our heuristic can lead to a simple, time-consistent contract that induces the supply chain partners to choose the heuristic solution $s_j^a(t)$ in each period.

A key enabler for implementing this contract is a supply chain *integrator* who knows the heuristic solution and is responsible for payment transfers between the firms. The integrator designs a contract for each stage j with three cost parameters $(\theta_j p_j^a, \theta_j h_j^a, \theta_j b_j)$, aiming to induce the stage manager to choose $s_j^a(t)$ in each period, where $\theta_j \geq 0$ is a stage-specific constant. The calculation for obtaining θ_j is explained in Online Appendix D.

The players in this game are the integrator and each of the echelon managers. The contract specifies the payment transfers between these players at the end of each period after the demand realizes. More specifically, the following payment scheme is announced to all players before the game: At the end of each period after the cost is evaluated, echelon j manager first pays the integrator based on the *accounting* echelon inventory level $\bar{x}_j$ according to the contract cost terms. (The accounting echelon inventory level is defined as the echelon inventory level by assuming that there is an ample supply from upstream.) Then, the integrator compensates the actual cost incurred for echelon j after echelon j implements the contract.

The goal of each player is to minimize the total cost in T periods. In Online Appendix D, we demonstrate that such a contract is implementable and can achieve the heuristic cost.

6. Concluding Remarks

The solution bounds and the heuristic can be extended to more general supply chain structures. Rosling (1989) shows that an assembly system can be transformed to a serial system with modified lead times. Because Rosling's result does not require any demand assumptions, our heuristic can be applied to the assembly system. The proposed heuristic can be applied to a one-warehouse-multiretailer distribution system with identical retailers. Under the so-called *balance* assumption (i.e., the inventory levels of the retailers can be freely and instantaneously redistributed as needed), the distribution system is equivalent to a two-stage serial system, where the downstream stage can be viewed as a composite stage that includes all retailers' demands. As shown in Federgruen and Zipkin (1984), one can apply the Clark and Scarf (1960) serial algorithm to obtain the echelon base-stock level for the warehouse and the composite stage. Then, one could apply the myopic allocation rule in each period to the retailers to determine the retailers' base-stock levels. Federgruen and Zipkin (1984) reports that this approach can generate a very effective solution. Clearly, we can apply the same technique to generate single-stage approximations for the resulting two-stage system and then apply the myopic allocation for the system. It will be interesting to examine whether our results can be applied to the nonidentical retailer system. We leave this for future research.

Our solution bounds as well as the heuristic can be extended to a system with Markov-modulated demand. More specifically, assume that the demand process is driven by a homogeneous, discrete-time Markov chain $\mathbf{W}$ with K states. It is known that a state-dependent echelon base-stock policy is optimal (e.g., Chen and Song 2000, Muharremoglu and Tsitsiklis 2008). Let $s_j(k, t)$ be the optimal echelon base-stock level for stage j when the demand state is k in period t, $k = 1, \ldots, K$. Following a similar analysis, we can derive single-stage bounds $s_j^l(k, t)$ and $s_j^u(k, t)$ for each demand state k in each period t such that $s_j^l(k, t) \leq s_j(k, t) \leq s_j^u(k, t)$. A detailed analysis is available from the author.

Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at http://msom.journal.informs.org/.

Acknowledgments

The author thanks Stephen C. Graves and the anonymous reviewers for useful suggestions. He is grateful to Li Chen and Jing-Sheng Song for comments on a previous version of this paper.

References

Abhyankar, S., S. Graves. 2001. Creating an inventory hedge for Markov-modulated Poisson demand: An application and a model. *Manufacturing Service Oper. Management* **3**(4) 306–320.

Cachon, G., P. Zipkin. 1999. Competitive and cooperative inventory policies in a two-stage supply chain. *Management Sci.* **45**(3) 936–953.

Chao, X., S. Zhou. 2007. Probabilistic solution and bounds for serial inventory system with discounted and average cost. *Naval Res. Logist.* **54**(6) 623–631.

Chen, F. 1999. Decentralized supply chains subject to information delays. *Management Sci.* **45**(8) 1076–1090.

Chen, F., J.-S. Song. 2000. Optimal policies for multiechelon problems with Markov-modulated demand. *Oper. Res.* **49**(2) 226–234.

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

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

Dong, L., H. Lee. 2003. Optimal policies and approximations for a serial multi-echelon inventory system with time-correlated demand. *Oper. Res.* **51**(6) 969–980.

Donohue, K. 2000. Efficient supply contracts for fashion goods with forecast updating and two production modes. *Management Sci.* **46**(11) 1397–1411.

Erkip, N., W. Hausman, S. Nahmias. 1990. Optimal centralized ordering policies in multi-echelon inventory systems with correlated demands. *Management Sci.* **36**(3) 381–392.

Ettl, M., G. Feigin, G. Lin, D. Yao. 2000. A supply network model with base-stock control and service requirements. *Oper. Res.* **48**(2) 216–232.

Federgruen, A., P. Zipkin. 1984. Computational issues in an infinite horizon multi-echelon inventory problem with stochastic demand. *Oper. Res.* **32**(4) 818–836.

Gallego, G., Ö. Özer. 2003. Optimal replenishment policies for multiechelon inventory problems under advance demand information. *Manufacturing Service Oper. Management* **5**(2) 157–175.

Gallego, G., Ö. Özer. 2005. A new algorithm and a new heuristic for serial supply systems. *Oper. Res. Lett.* **33**(4) 349–362.

Graves, S., S. Willems. 2000. Optimizing strategic safety stock placement in supply chains. *Manufacturing Service Oper. Management* **2**(1) 68–83.

Graves, S., S. Willems. 2008. Strategic inventory placement in supply chains: Nonstationary demand. *Manufacturing Service Oper. Management* **10**(2) 278–287.

Lee, H., S. Whang. 1999. Decentralized multi-echelon supply chains: Incentives and information. *Management Sci.* **45**(5) 633–640.

Muharremoglu, A., J. Tsitsiklis. 2008. A single-unit decomposition approach to multiechelon inventory systems. *Oper. Res.* **56**(5) 1089–1103.

Parker, R. P., R. Kapuscinski. 2011. Managing a noncooperative supply chain with limited capacity. *Oper. Res.* **59**(4) 866–881.

Rosling, K. 1989. Optimal inventory policies for assembly systems under random demands. *Oper. Res.* **37**(4) 565–579.

Schoenmeyr, T., S. Graves. 2009. Strategic safety stocks in supply chains with evolving forecasts. *Manufacturing Service Oper. Management* **11**(4) 657–673.

Shang, K., J.-S. Song. 2003. Newsvendor bounds and heuristic for optimal policies in serial supply chains. *Management Sci.* **49**(5) 618–638.

Shang, K., J.-S. Song, P. Zipkin. 2009. Coordination mechanisms in decentralized serial inventory systems with batch ordering. *Management Sci.* **55**(4) 685–695.

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