Available online at www.sciencedirect.com
ScienceDirect
European Journal of Operational Research 196 (2009) 186–197
EUROPEAN JOURNAL OF OPERATIONAL RESEARCH
www.elsevier.com/locate/ejor

Stochastics and Statistics

Managing stochastic inventory systems with free shipping option

Bin Zhou [a,*], Michael N. Katehakis [b], Yao Zhao [b]

[a] *Department of Management, College of Business and Public Administration, Kean University, Union, NJ 07083, USA*
[b] *Department of Management Science and Information Systems, Rutgers University, Newark, NJ 07102, USA*

Received 27 September 2007; accepted 30 January 2008
Available online 6 February 2008

Abstract

In many industries, customers are offered free shipping whenever an order placed exceeds a minimum quantity specified by suppliers. This allows the suppliers to achieve economies of scale in terms of production and distribution by encouraging customers to place large orders. In this paper, we consider the optimal policy of a retailer who operates a single-product inventory system under periodic review. The ordering cost of the retailer is a linear function of the ordering quantity, and the shipping cost is a fixed constant K whenever the order size is less than a given quantity – the free shipping quantity (FSQ), and it is zero whenever the order size is at least as much as the FSQ. Demands in different time periods are i.i.d. random variables. We provide the optimal inventory control policy and characterize its structural properties for the single-period model. For multi-period inventory systems, we propose and analyze a heuristic policy that has a simple structure, the (s, t, S) policy. Optimal parameters of the proposed heuristic policy are then computed. Through an extensive numerical study, we demonstrate that the heuristic policy is sufficiently accurate and close to optimal.

Keywords: Free shipping; Inventory management; Multiple-time periods; Dynamic programming

1. Introduction

In many industries, customers or retailers are offered free shipping service whenever an order placed exceeds a minimum quantity specified by suppliers. This allows the suppliers to achieve economies of scale in terms of production and distribution by modifying their customers' purchasing behavior. The free shipping quantity (FSQ) practice is not uncommon even in international trade. Due to substantial production setup costs and longer lead times, international companies often use minimum free shipping quantity to avoid frequent small orders and thus achieve economies of scale by encouraging customers to place larger orders, less frequently.

Companies also provide more flexible options where their customers are allowed to order below the FSQ at an additional shipping or handling charge. Minimum order with free shipping option has significant impact and challenges the efficient control of retailers' inventory. In fact, our analysis demonstrates that the free shipping option helps to enhance retailers' flexibility in responding to demands, which are varied and highly uncertain, and consequently it decreases the retailers' total costs (ordering and inventory costs) when used properly.

In this paper, we consider a retailer who manages a single-product inventory system with free shipping option and stochastic demands. The ordering cost of the retailer is linear if its order size is no less than Q, i.e. the free shipping quantity (FSQ), and an additional fixed shipping cost K is imposed whenever the order quantity falls below the FSQ. We model this

* Corresponding author. Tel.: +1 908 737 4175.
E-mail address: bzhou@kean.edu (B. Zhou).

0377-2217/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2008.01.042

problem using stochastic dynamic programming. Under general assumptions on demand processes, we characterize structural properties of optimal cost functions and determine the optimal inventory policy for the retailer in single period. We show that the structure of optimal policies for the free shipping option problem is in general more complex than that of the models with fixed setup cost or minimum order quantity (or commitment) only. We discuss how our model can be extended to include constant lead times by following standard procedures (cf. Heyman and Sobel, 1984). In the case of stochastic lead times, additional assumptions are required on the sequence of order received, see Kaplan (1970) for the case of exclusion of order crossover, and Zalkind (1978) and Bradley and Robinson (2005) for the case of order crossover.

For multi-period inventory systems, optimal policies have quite complicated structures and hence they are difficult to identify and implement into practice. We propose and analyze a heuristic policy that has a simple structure, namely the (s,t,S) policy, where $S-s=Q$. The (s,t,S) policy works as follows: if the initial inventory position x is no more than s, order up to $s+Q$; if the initial inventory is above s but no more than t, order exactly the free shipping quantity Q; if the initial inventory exceeds t but no more than S, order amount $a(x)<Q$ (x is the initial inventory position); if the initial inventory is above S, do not order.

The paper is the first attempt to study and analyze the impacts of free shipping and free shipping quantity on inventory control policies. We focus on: (1) presenting structural analysis of optimal inventory policies, (2) proposing effective heuristic policies for multi-period inventory systems, and (3) quantifying the impacts of free shipping option and key parameters of the model. We conduct an extensive computational study and show that the heuristic policy is sufficiently accurate and it provides close-to-optimal solutions. We also discuss managerial insights regarding the free shipping option and the fixed cost K, the free shipping quantity Q, demand variability, and the service level on the inventory systems.

The remainder of this paper proceeds as follows. In Section 2, we offer a brief review of related literature. In Section 3, we formulate the problem mathematically and study the model in finite and infinite time horizons under the total discounted cost criterion. A discussion of incorporating lead times into the model is provided. In Section 4, we characterize structural properties of optimal cost functions and inventory policies for single-period inventory systems. In Section 5, we propose and analyze a heuristic policy for multi-period inventory systems. In Section 6, an extensive numerical study is conducted to gauge the effectiveness of the heuristic policy. Managerial insights with respect to systems parameters are also provided and analyzed. Finally, we give the conclusions of the paper in Section 7.

2. Literature review

One of the key concerns in production and inventory problems is how to effectively control production, inventory holding, and stockout costs. When replenishing inventories, customers or retailers generally face three different types of requirements from their suppliers: (a) fixed setup cost, (b) fixed batch size, or (c) minimum order quantity.

For fixed setup cost problems, the seminal work of Clark and Scarf (1960) unveils a critical inventory replenishment problem for a periodic-review, finite horizon, serial multi-echelon inventory system where demand occurs at the lowest stage. They show that the optimal inventory control policy of the overall system follows a base-stock policy for every stage, provided that there is a fixed cost only at the highest stage. Scarf (1960) proves that an (s,S) policy is optimal for periodic-review, single-product inventory systems with fixed ordering cost and general demand distribution. Federgruen and Zipkin (1984) later extend the results to a periodic-review, infinite horizon case under both discounted and average cost criteria.

In the second stream of research on fixed batch size, Veinott (1966) analyzes a dynamic inventory model with fixed batch size and proves that the (R,nQ) policy is optimal. This policy requires that if the inventory level is above the re-order point R, no order needs to be placed. Whereas, if the inventory level drops below point R, orders should be placed in batches to raise the inventory level to the smallest value above the re-order point. Chen and Zheng (1994) provide a simple way to show that the base-stock policy remains optimal when each stage orders in batches under some modified conditions.

The third group of existing papers where order quantity is the major theme includes Fisher and Raman (1996), Robb and Silver (1998), Zhao and Katehakis (2006), and Zhou et al. (2007). Fisher and Raman (1996) consider an inventory system where the retailer faces both a minimum order quantity requirement and a capacity constraint. They formulate the problem into a stochastic program and provide optimal policies for a two time-period model. Robb and Silver (1998) study a periodic review inventory model with minimum order quantities. They conduct extensive simulation experiments that demonstrate situations where a simple 'rounding up' decision rule performs poorly. The papers that have close ties with this research are Zhao and Katehakis (2006) and Zhou et al. (2007).

Zhao and Katehakis (2006) explore a single-item stochastic inventory system where the retailer has to order either none or at least as much as the minimum order quantity. No fixed cost is considered in the system. They establish models in both finite and infinite horizons under discounted cost criterion using dynamic programming. By introducing a new "M-increasing" concept, they characterize the optimal policy in certain regions of the state space for any time period. Zhou et al. (2007) extend the work and develop a simple and effective heuristic policy for single-product stochastic inventory systems with minimum order quantity and a linear ordering cost. They develop an efficient algorithm to compute the optimal values of the heuristic parameters in the infinite time horizon under the average cost criterion. Through numerical studies, they show that the performance of the heuristic policy is very close to that of the optimal policy and it significantly outperforms the (s, S) policy under general conditions. In this paper, we consider a different and more complicated problem where the inventory system involves an order quantity requirement, a fixed cost as well as a free shipping option.

In all existing literature, either a fixed ordering cost or an order quantity requirement is considered. None of these papers analyze how the option of free shipping influences inventory control policies. In real life business practice, retailers sometimes need to order below the free shipping quantity and pay the extra shipping fee, due to demand variability, inventory capacity, and financial constraints. Hence, it is practically important to study and analyze ordering policies under the free shipping option and provide related managerial insights. This paper represents, to our knowledge, the first attempt towards this end.

3. The model

We model a stochastic inventory system where a retailer operates a single product and reviews its inventory periodically. The demand in time period n is denoted as D_n, where D_n are i.i.d. random variables with finite means $E(D_n)$. At the beginning of each time period, the retailer reviews its inventory and may place an order to its supplier; at the end of the time period, demand is realized and the retailer fills the demand as much as it can from the available stock. If the retailer is not able to satisfy all the demand, the excessive amount is backlogged. In each time period, the retailer incurs inventory holding cost h per unit for any unsold item and penalty cost π per unit for any extra demand. The supplier has ample supply with unlimited production capacity. Denote Q to be the free shipping quantity (FSQ) specified by the supplier. The ordering cost of a units is ca (i.e. linear) if the order size a of the retailer is at least as much as Q; whereas a fixed shipping cost K is imposed in addition to the linear ordering cost if the order size is below Q.

For the ease of exposition, we assume that products arrive at the retailer's site immediately after an order is placed. The model can be extended to include constant or stochastic lead times following standard procedures, see Heyman and Sobel (1984). In the case of stochastic lead times, additional assumptions are required, see Kaplan (1970) and Bradley and Robinson (2005).

Let x be the inventory position at the beginning of a time period which can take negative values, it is equal to the inventory on-hand minus backorders; and let y be the inventory position after the order decision is made. Then the action $a = a(x) = y - x$ is the amount ordered in that time period. Let $A_x^K = \{1, \dots, Q-1\}$ denote the set of all actions that order below the free shipping quantity, and let $A_x^0 = \{0\} \cup \{Q, Q+1, \dots\}$ denote the set of all feasible actions that order none or at least Q.

We next formulate the problem as a Markov Decision Process (MDP). We first establish the finite and infinite horizon models under the total discounted cost criterion, then we consider the model under the long-run average cost per unit time. In the latter case, it is known that the linear ordering cost can be ignored, see, e.g. Zheng and Federgruen (1991). For a given inventory position x and an action a, the single-period cost functions are

$$ca + E(L(x+a, D)), \quad a \in A_x^0, \tag{1}$$

$$K + ca + E(L(x+a, D)), \quad a \in A_x^K, \tag{2}$$

where $L(x+a, D) = h(x+a-D)^+ + \pi\,(D-x-a)^+$ is the holding or backorder penalty cost when the demand of the period is D, $E(\cdot)$ is its expectation with respect to the demand distribution, K is the optional shipping cost. Besides, $y = x + a$ (an integer) is the inventory position after ordering a units.

Given an initial inventory position x and a discount factor $\beta \in (0,1)$, let $U_N(x,\beta)$ (respectively $U(x,\beta)$) denote the expected minimum total discounted inventory and ordering costs for a finite horizon of length $N \geqslant 1$ (respectively infinite horizon) problem. Using standard dynamic programming arguments, we can show that the dynamic programming optimality equations for the finite and the infinite horizon problems are

$$U_n(x,\beta) = \min\left\{ \begin{array}{l} \inf_{a\in A_x^0}\{ca + E(L(x+a, D_n)) + \beta E(U_{n-1}(x+a-D_n, \beta))\}; \\ \min_{a\in A_x^K}\{K + ca + E(L(x+a, D_n)) + \beta E(U_{n-1}(x+a-D_n, \beta))\} \end{array} \right\}, \quad n = 1, \dots, N, \tag{3}$$

$$U(x,\beta) = \min\left\{ \begin{array}{l} \inf_{a\in A_x^0}\{ca + E(L(x+a, D)) + \beta E(U(x+a-D, \beta))\}; \\ \min_{a\in A_x^K}\{K + ca + E(L(x+a, D)) + \beta E(U(x+a-D, \beta))\} \end{array} \right\}. \tag{4}$$

Without loss of generality, we assume that the salvage value of any unsold item at the end of the planning horizon is zero, $U_0\,(x,\beta) = 0$. Define $V_n(x,\beta) = U_n(x,\beta) + cx$, $\forall\ n \geqslant 1$.

Based on action sets A_x^0 and A_x^K, we further define for the last time period $U_1^0(y,\beta) = cy + E(L(y, D_1, \beta))$ and $U_1^K(y,\beta) = K + cy + E(L(y, D_1, \beta))$. Similarly, for period $n \geqslant 2$, we define $U_n^0(y,\beta) = c(1-\beta)y + E(L(y, D_n)) + \beta E(V_{n-1}(y - D_n, \beta))$ and $U_n^K(y,\beta) = K + c(1-\beta)y + E(L(y, D_n)) + \beta E(V_{n-1}(y - D_n, \beta))$.

Then, we can re-write Eq. (3) for the finite horizon as follows:

$$V_1(x,\beta) = \inf\{\{U_1^K(y,\beta) : x+1 \leqslant y \leqslant x+Q-1\};\ \{U_1^0(y,\beta) : x+Q \leqslant y\};\ U_1^0(x,\beta)\},$$
$$V_n(x,\beta) = \inf\{\{U_n^K(y,\beta) : x+1 \leqslant y \leqslant x+Q-1\};\ \{U_n^0(y,\beta) : x+Q \leqslant y\};\ U_n^0(x,\beta)\} + 2\beta c E(D_n),\quad n \geqslant 2. \tag{5}$$

Similarly, for the infinite horizon model with the demand of the period D, from Eq. (4) we have

$$V(x,\beta) = \inf\{\{U^K(y,\beta) : x+1 \leqslant y \leqslant x+Q-1\};\ \{U^0(y,\beta) : x+Q \leqslant y\};\ U^0(x,\beta)\} + 2\beta c E(D). \tag{6}$$

In the rest of the paper, we will consider policies under which the state changes in time following Markov Chain and satisfies the irreducibility conditions. It is easy to see that policies that violate the irreducibility conditions are either ordering too little or ordering too much, and this type of policies can not be optimal. Under the long-run average cost criterion, it is known that the limit $V(x)$ below exists and it is independent of x

$$V(x) = \lim_{\beta\to 1}(1-\beta)V(x,\beta). \tag{7}$$

The average cost optimal policies can be obtained either by solving linear programs (Derman, 1970), or by solving average cost optimality equations using value iteration (see Tijms, 1994, Chapter 3, pp. 206–209).

Lead times, deterministic and stochastic, can be included into the model under necessary conditions. This is typically done by defining the state variable to represent inventory on-hand plus inventory on-order and by propitious modifications to the cost functions and transition probabilities. For instance, consider the case in which there are τ time periods between order placement and order receipt of the goods for the retailer. Indeed, following Heyman and Sobel (1984), at time period n, re-define the retailer's inventory position x to be the inventory on-hand plus items on-order (when goods ordered τ periods ago arrive at retailer's site) before the ordering decision at the beginning of period n has been made. Demand is then met by the available inventory. Let a denote the amount ordered in period n. Then we can re-define the single-period cost functions corresponding to inventory position x and action a as follows:

$$ca + E(L(x+a, \widehat{D}_\tau)),\quad a \in A_x^0,$$
$$K + ca + E(L(x+a, \widehat{D}_\tau)),\quad a \in A_x^K,$$

where $\widehat{D}_\tau = \sum_{i=0}^{\tau} D_i$ and $L(y, \widehat{D}_\tau) = h(y - \widehat{D}_\tau)^+ + \pi(\widehat{D}_\tau - y)^+$ is the holding or backorder penalty cost when the demand of periods $n, n+1, \ldots, n+\tau$ is $\widehat{D}_\tau$. Then, the optimality Eqs. (3) and (4) can be revised by substituting the above expressions for the single-period cost functions.

4. Optimal policy for single-period systems

We next develop optimal ordering policies for single-period inventory systems with free shipping option in Theorems 4.1–4.3.

In view of the last time period in the planning horizon, we have (for simplicity, we omit discount factor β hereafter)

$$V_1(x) = \inf\{U_1^0(x); \{U_1^0(y) : x+Q \leqslant y\};\ \{U_1^K(y) : x+1 \leqslant y \leqslant x+Q-1\}\},$$

where x is the initial inventory position, and y is the inventory position after the order decision is made. Hence for $\pi > c$, clearly the single-period cost function $U_1^0(y)$ is convex and satisfies $U_1^0(y) \to +\infty$ as inventory position $|y| \to +\infty$. Let y_1^* be the value where $U_1^0(y)$ reaches its global minimum. Similarly, $U_1^K(y+q)$ ($q \in \{1, \ldots, Q-1\}$) is also convex and let $y_1^*(q)$ be the value where $U_1^K(y+q)$ has its global minimum.

For any convex function $U(y) \to +\infty$ as $|y| \to +\infty$, Zhao and Katehakis (2006, Lemma 3.2) shows that there is an unique junction y_0 such that $U(y_0) = U(y_0 + Q)$ and $U(x+Q) \leqslant U(x)\ \forall x < y_0$, $U(x) \leqslant U(x+Q')\ \forall x > y_0 \forall Q' \geqslant Q$. Clearly

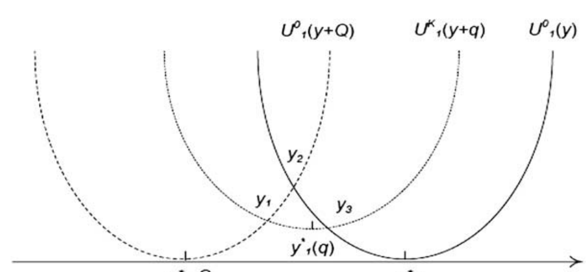

Fig. 1. Cost functions of a single-period model with free shipping option.

convex function $U_1^0(y)$ follows these properties. We let y_1 be the junction such that convex functions $U_1^0(y_1+Q)=U_1^K(y_1+q)$; y_2 be the junction where $U_1^0(y_2)=U_1^0(y_2+Q)$; and y_3 be the junction where $U_1^0(y_3)=U_1^K(y_3+q)$. In Fig. 1, we illustrate the structure of the ordering policies for a single-period inventory system.

We are now ready to develop the following theorems that characterize the optimal policies for single-period inventory systems with free shipping option.

Theorem 4.1. *In single-period inventory systems with free shipping option, if $\pi > c$ and $y_2 < y_1^*(q)$, the optimal policy for the retailer, given initial inventory position x, is to order*

$$a^*(x)=\begin{cases} y_1^*-x, & if\ x\in(-\infty,y_1^*-Q],\\ Q, & if\ x\in(y_1^*-Q,y_1],\\ y_1^*-x, & if\ x\in(y_1,y_2],\\ y_1^*-x, & if\ x\in(y_2,y_1^*(q)],\\ q(x), & if\ x\in(y_1^*(q),y_3],\\ 0, & otherwise.\end{cases}$$

Proof. For $x\leqslant y_1^*-Q$, obviously, it is optimal for the retailer to raise its inventory up to global minimum y_1^*. If $y_1^*-Q<x\leqslant y_1$, we first notice that $U_1^0(x+Q)\leqslant U_1^0(x)$, which implies that the system should order at least Q; next, because $U_1^0(y)$ is convex and $x+Q>y_1^*$, then $U_1^0(x+Q')\geqslant U_1^0(x+Q)$, for all $Q'>Q$. Hence, ordering exactly Q is optimal. For $y_1<x\leqslant y_2$ and $y_2<x\leqslant y_1^*(q)$, it is optimal to order up to y_1^* as is the case of $x\leqslant y_1^*-Q$. If $y_1^*(q)<x\leqslant y_3$, it is optimal to order $q(x)$, $1+x\leqslant q(x)\leqslant Q+x-1$ and it is dependent upon initial inventory position x, since $U_1^K(x+q)<U_1^0(x)$ for any $q\in\{1,\ldots,Q-1\}$. And finally, if $x>y_3$, since $U_1^K(x+q')\geqslant U_1^0(x)$, for all $q'\geqslant q$, thus the optimal policy is to order zero. □

Theorem 4.2. *In single-period inventory systems with free shipping option, if $\pi > c$ and $y_2 > y_1^*(q)$, the optimal policy for the retailer, given initial inventory position x, is to order*

$$a^*(x)=\begin{cases} y_1^*-x, & if\ x\in(-\infty,y_1^*-Q],\\ Q, & if\ x\in(y_1^*-Q,y_1],\\ y_1^*-x, & if\ x\in(y_1,y_1^*(q)],\\ q(x), & if\ x\in(y_1^*(q),y_2],\\ q(x), & if\ x\in(y_2,y_3],\\ 0, & otherwise.\end{cases}$$

Proof. If $x\leqslant y_1^*-Q$, we proved in Theorem 4.1, it is optimal to order up to y_1^*. If $y_1^*-Q<x\leqslant y_1$, ordering Q is the optimal policy. If $y_1<x\leqslant y_1^*(q)$, it is again optimal to order up to y_1^*. For $y_1^*(q)<x\leqslant y_2$ and $y_2<x\leqslant y_3$, the optimal policy is to order $q(x)$. Proof is shown in Theorem 4.1. And finally, for $x>y_3$, order nothing is optimal. □

Theorem 4.3. *In single-period inventory systems with free shipping option, if $\pi > c$ and $y_2 = y_1^*(q)$, the optimal policy for the retailer, given initial inventory position x, is to order*

$$a^*(x)=\begin{cases} y_1^*-x, & if\ x\in(-\infty,y_1^*-Q],\\ Q, & if\ x\in(y_1^*-Q,y_1],\\ y_1^*-x, & if\ x\in(y_1,y_1^*(q)=y_2],\\ q(x), & if\ x\in(y_1^*(q)=y_2,y_3],\\ 0, & otherwise.\end{cases}$$

Proof. The proof is shown in Theorems 4.1 and 4.2.

Theorems 4.1–4.3 demonstrate that the optimal policies for single-period inventory systems with free shipping option have different and more complex structures than the (s,S) policy that is optimal for the models with fixed cost only. For a detailed discussion of the single-period model with fixed cost only, we refer to Ross (1992, Chapter 8, pp. 169–171). We point out that if the fixed cost K is sufficiently large such that $U_1^0(y)\leqslant U_1^K(y)$, then the cost function of the single time period problem takes the form below:

$$V_1(x)=\inf\{U_1^0(x);\ U_1^0(y):y\geqslant x+Q\}.$$

Hence, in this case the ordering decision only involves the order quantity Q without the free shipping option.

From Fig. 1, for instance, we can observe that if the retailer's inventory position is less than or equal to $y_1^* - Q$, it is optimal for the retailer to raise its inventory to y_1^*; if the inventory position is greater than $y_1^* - Q$ but less than or equal to y_1, then it is optimal to order the free shipping quantity Q. If the inventory position is greater than y_1 but less than or equal to $y_1^*(q)$, it is again optimal to order up to y_1^*; if the inventory position is greater than $y_1^*(q)$ but less than or equal to y_3, the optimal policy is to order $q(x)$ and pay the extra shipping cost K; finally, if the inventory position is above y_3, no order needs to be placed.

The optimal inventory policy is inherently complicated even for the single-period model due to the option of ordering below Q at a fixed cost K. The multiple intersections of the cost functions separate the inventory position into several regions, which make the replenishment policies hard to identify and difficult to implement practically. We next propose and analyze a simple and effective heuristic policy for multi-period inventory systems with free shipping option. □

5. Heuristic policy for multiple-period systems

In light of the optimal control policies of single-period inventory systems, we construct the following class of stationary policies, namely the (s,t,S) policies. Given integers s,t, and S, where $s \leqslant t \leqslant S < s+Q$, and initial inventory position x, the policy is to order

$$a(x) = \begin{cases} s+Q-x, & \text{if } x \in (-\infty, s], \\ Q, & \text{if } x \in (s,t], \\ \phi(x), & \text{if } x \in (t,S], \\ 0, & \text{otherwise.} \end{cases}$$

An (s,t,S) policy works as follows. When the inventory position x is smaller than or equal to s, order up to $s+Q$; when x is greater than s but smaller than or equal to t, order exactly the free shipping quantity Q; when x is above t but smaller than or equal to S, order $\phi(x) \in \{1,\dots,Q-1\}$; and when x is above S, order nothing. One needs to calculate the minimum cost numerically to determine the corresponding order quantity below Q for $t < x \leqslant S$. Notice that when the fixed cost K is very large, ordering below the free shipping quantity is not optimal, and the heuristic policy reduces to the (s,t) policy discussed in Zhou et al. (2007).

Let $L_{x,\alpha}$ be the retailer's one period inventory holding or backorder penalty cost when its inventory position is x and the action (order decision) is α, where $\alpha \in A = \{0,1,\dots,+\infty\}$. Let $\Lambda_{x,\alpha}$ be the long-term fraction of time (the state – action frequency) the system is in inventory position x taking action α. And let $A^K = \{1,2,\dots,Q-1\}$ be the action set of order below the free shipping quantity Q. To identify the inventory control policy, we focus on minimizing the long-run average cost of the retailer, which is denoted by $C(s,t,S)$:

$$\text{minimize } C(s,t,S) = \sum_{x=l}^{s} L_{x,s+Q-x} \cdot \Lambda_{x,s+Q-x} + \sum_{x=s+1}^{t} L_{x,Q} \cdot \Lambda_{x,Q} + \sum_{x=t+1}^{S} L_{x,q} \cdot \Lambda_{x,q} + \sum_{x=S+1}^{u} L_{x,0} \cdot \Lambda_{x,0} + K \cdot \sum_{x=t+1}^{S} \Lambda_{x,q},$$

where l and u are, respectively, the lower and the upper state space of x, and $x \in X = \{l, l+1,\dots,u\}$, and $q \in \{1,\dots,Q-1\}$ is the order quantity below Q. Following Derman (1970), the optimal values of s^*, t^*, and S^* can be obtained by solving the following linear program of the average cost:

$$\begin{aligned} &\text{minimize } C(s,t,S) \\ &\text{subject to } \sum_{x\in X}\sum_{\alpha\in A} \Lambda_{x,\alpha} = 1 \\ &\qquad \sum_{x\in X}\sum_{\alpha\in A} \Lambda_{x,\alpha} \cdot p_{x,y,\alpha} = \sum_{\alpha\in A} \Lambda_{y,\alpha}, \quad y \in X \\ &\qquad \Lambda_{x,\alpha} \geqslant 0, \end{aligned} \tag{8}$$

where $p_{x,y,\alpha}$ is the one-step transition probability from state (inventory position) x to state (inventory position) y when taking action α.

6. Computational results

In this section, we conduct an extensive computational study on the (s,t,S) policy and present our findings. Our goal is to determine its effectiveness and performance, and understand the impacts and insights with respect to the free shipping option, different levels of free shipping quantities, the fixed cost K, and other key parameters of the model.

We conduct computations for three commonly used demand distributions in inventory control literature: the Normal, the Poisson, and the Uniform distributions. To achieve discrete demand for the Normal distribution, we use the standard approach to sum the decimals to the nearest integer (see, e.g. Aviv and Federgruen, 2001). The Normal distribution is a widely used demand distribution for inventory problems, we refer our readers to Silver et al. (1998) for a detailed justification of its significance in business practice. We design the numerical analysis on the following non-dimensional parameters: $Q/E(D), \pi/(\pi+h)$, and demand variability.

We test Normal demand distributions with a variety of expected mean demands and standard deviations (Normal(μ, σ) has mean μ and standard deviation σ). Specifically, for mean demand $\mu = 10$, we use Normal(10, 2), Normal(10, 5), and Normal(10, 10); for $\mu = 20$, we use Normal(20, 2), Normal(20, 5), and Normal(20, 10); and for $\mu = 30$, we use Normal(30, 2), Normal(30, 5), and Normal(30, 10). We study Poisson demand distributions with expected mean demand $\lambda = 10$, 20, and 30, and Uniform distributions [0, 20] and [0, 40].

The unit holding cost h was equal to 1. The ratio of the free shipping quantity and the mean demand $Q/E(D)$ was varied from 0 to 10 in increment of 1. We let the ratio $\pi(\pi+h)$ equal to 0.80, 0.85, 0.90 and 0.95. The value of the fixed shipping cost K was varied from 1 to 5 with an increment of 1. Altogether, there were a total of 3080 instances in our computational setup.

6.1. Effectiveness of the (s, t, S) policy

We evaluate the heuristic policy by comparing the average costs of the optimal (s, t, S) policies with those of the true optimal policies. We measure the effectiveness of the (s, t, S) policy with respect to the optimal policy by observing the percentage difference between the average costs of these two policies, denoted by Deviation 1, using the formula

$$\text{Dev. } 1 = 100 \times \frac{(s, t, S) \text{ average costs} - \text{Optimal average costs}}{\text{Optimal average costs}}.$$

Second, we compare the average costs of the (s, t, S) policy with those of the policy which includes minimum order quantity but does not consider the free shipping option, i.e. the (s, t) policy provided by Zhou et al. (2007). The (s, t) policy works as follows: if the inventory position is no more than s, order up to $s+Q$; if the inventory position is above s but no more than t, order exactly Q; if the inventory position is above t, order nothing. To achieve the average costs of the (s, t) policy, we follow the method discussed in Zhou et al. (2007). The percentage difference between the average costs of these two policies is denoted by Deviation 2 and computed using the formula as follows:

$$\text{Dev. } 2 = 100 \times \frac{(s, t) \text{ average costs} - (s, t, S) \text{ average costs}}{(s, t, S) \text{ average costs}}.$$

We first report our computational results for Normal demand distributions in Table 1. The last two columns of the table contrast Dev. 1 and Dev. 2. Under Normal demand distribution with $\mu = 10$, for instance, when the demand standard deviation is quite small $\sigma = 2$, the average percentage deviation between the (s, t, S) policies and the optimal policies varies from 1.67% to 2.19%. However, the average percentage deviation between the (s, t) and the (s, t, S) policies is significantly larger, which varies from 15.72% to 16.05%. When $\sigma = 5$, the average percentage deviation between the (s, t, S) policies and the optimal policies declines to the range of 0.00% to 0.03%. In fact, as the standard deviation increases, the percentage deviations tend to decrease. At $\sigma = 10$, the average percentage deviation reduces to 0.00% under the (s, t, S) policy.

Table 1 also presents the results for the cases where the values of expected mean demands become larger. When μ increased to 20, σ was equal to 2, 5, and 10, the average percentage deviation between the (s, t, S) policies and the optimal policies varies from 0.00% to 3.94%. When $\mu = 30$, the percentage deviation varies from 0.00% to 5.09%. On the other hand, the average percentage deviations between the (s, t) and the (s, t, S) policies are all quite substantial. They vary from 10.98% to 18.26% when $\mu = 20$, and from 15.18% to 21.28% when $\mu = 30$. Results for both cases show that the percentage deviations between the (s, t, S) heuristic policies and the optimal policies are still very small and the (s, t, S) policies present quite robust and close-to-optimal performance.

From Table 1, we observe that for a fixed value of mean demand ($\mu = 10$, 20 and 30), the percentage deviations decrease as σ increases, and they decline to 0.00% when their standard deviations reach the biggest values. The intuitive interpretation for these results is that when there is large variability (higher standard deviation) in the demand, the retailer can benefit more from the flexibility of either ordering below the free shipping quantity or ordering big lot size to receive free shipping, than in the case where demand is much less varied (smaller standard deviation). Hence the (s, t, S) heuristic policies perform better and are closer to the true optimal policies, and they outperform the (s, t) policies in which no free shipping option is offered.

In addition, for smaller standard deviations, bigger percentage deviations indicate that the optimal policies may not have same structures as the (s, t, S) policies. To explore the true structure and understand the complexity of the optimal policies, we next compute the optimal policy for the case in which the standard deviation is very small, i.e. $\sigma = 1$.

We present the policies in Fig. 2. Notice that the optimal ordering policy does not possess the same structure as the (s, t, S) heuristic policy given small demand variability (small standard deviation). In fact, the optimal policy can be very

Table 1
Effectiveness of the (s, t, S) policy

Demand distribution	$\pi/(\pi+h)$	$Q/E(D)$	K	Avg. Dev. 1	Avg. Dev. 2
Normal $\mu=10$, $\sigma=2$	0.80	[0,10]	[1,5]	2.03	15.97
	0.85	[0,10]	[1,5]	2.19	15.87
	0.90	[0,10]	[1,5]	1.89	15.72
	0.95	[0,10]	[1,5]	1.67	16.05
Normal $\mu=10$, $\sigma=5$	0.80	[0,10]	[1,5]	0.01	9.77
	0.85	[0,10]	[1,5]	0.03	9.70
	0.90	[0,10]	[1,5]	0.00	9.66
	0.95	[0,10]	[1,5]	0.00	9.83
Normal $\mu=10$, $\sigma=10$	0.80	[0,10]	[1,5]	0.00	8.19
	0.85	[0,10]	[1,5]	0.00	8.07
	0.90	[0,10]	[1,5]	0.00	7.89
	0.95	[0,10]	[1,5]	0.00	8.27
Normal $\mu=20$, $\sigma=2$	0.80	[0,10]	[1,5]	3.82	17.23
	0.85	[0,10]	[1,5]	3.94	17.03
	0.90	[0, 10]	[1,5]	3.41	16.91
	0.95	[0,10]	[1,5]	3.04	18.26
Normal $\mu=20$, $\sigma=5$	0.80	[0,10]	[1,5]	0.01	15.86
	0.85	[0,10]	[1,5]	0.04	15.62
	0.90	[0,10]	[1,5]	0.01	15.48
	0.95	[0,10]	[1,5]	0.00	15.93
Normal $\mu=20$, $\sigma=10$	0.80	[0,10]	[1,5]	0.00	11.40
	0.85	[0, 10]	[1,5]	0.00	11.27
	0.90	[0, 10]	[1,5]	0.00	10.98
	0.95	[0, 10]	[1,5]	0.00	11.51
Normal $\mu=30$, $\sigma=2$	0.80	[0,10]	[1,5]	4.83	20.05
	0.85	[0, 10]	[1,5]	5.09	19.13
	0.90	[0, 10]	[1,5]	4.37	18.82
	0.95	[0, 10]	[1,5]	3.85	21.28
Normal $\mu=30$, $\sigma=5$	0.80	[0,10]	[1,5]	0.02	19.16
	0.85	[0, 10]	[1,5]	0.05	18.49
	0.90	[0, 10]	[1,5]	0.01	18.23
	0.95	[0, 10]	[1,5]	0.00	19.74
Normal $\mu=30$, $\sigma=10$	0.80	[0,10]	[1,5]	0.00	15.69
	0.85	[0, 10]	[1,5]	0.00	15.41
	0.90	[0, 10]	[1,5]	0.00	15.18
	0.95	[0, 10]	[1,5]	0.00	15.83

Normal distribution.

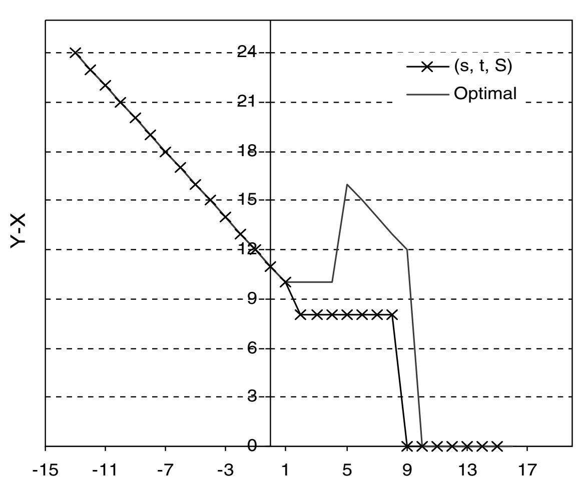

Fig. 2. Optimal order quantity, optimal policy vs. (s, t, S) policy.

complicated and hard to identify. Observe that the optimal policy has a 'peak' for the range $x \in \{5, \ldots, 10\}$, which corresponds to the order quantity above Q, while the heuristic policy is much simpler and smoothes out the 'peak'.

The (s,t,S) heuristic policies also perform very well when demand follows Poisson and Uniform distributions. Table 2 provides our computational results.

For Poisson demand distributions, we test three cases where λ was assigned the values 10, 20, and 30. For Uniform distributions, we use Uniform $[0,20]$ and Uniform $[0,40]$. The results obtained herein are consistent with those of Normal distributions: average percentage deviations between the (s,t,S) policies and the optimal policies are all very small, varying from 0.00% to 0.92% for Poisson and 0.00% for Uniform distributions (Table 2, column 5); while the percentage deviations between the (s,t) policies and the (s,t,S) policies vary from 13.25% to 19.57% for Poisson and from 9.11% to 10.35% for Uniform distributions (column 6), under different levels of penalty cost ratio $\pi/(\pi+h)$.

Comparing Tables 1 and 2, we observe that the (s,t,S) policies tend to perform slightly better under Uniform than under Poisson and Normal distributions. Intuitively, this is also due to the reason that retailer benefits more from the free shipping option given higher demand variability when demand follows Uniform distribution than Normal and Poisson distributions, at same mean demand values. The (s,t,S) policies are therefore more robust.

All results show that the (s,t,S) policies are very accurate and close to the optimal policies under a variety of combinations of the parameters under all three different demand distributions with varied mean demands and standard deviations. Moreover, for the problem of free shipping option, (s,t,S) policies perform much better than the (s,t) policies that do not consider the choice of free shipping.

To summarize the performance of the (s,t,S) policy, our numerical studies reveal that:

- In general, optimal policies of the inventory systems with free shipping option have quite complicated structures and they are difficult to identify and implement into practice.
- The (s,t,S) policies are sufficiently accurate and have close-to-optimal performance under different demand distributions with various combinations of parameters; and they outperform the (s,t) policies that are suitable for problems with minimum order quantity only.
- In all instances, as values of demand standard deviations increase, average percentage deviations between the (s,t,S) policies and the optimal policies decrease, implying that the heuristic tend to perform better when demand variability is larger.

6.2. Managerial insights

We next focus on providing managerial insights with regard to the impacts of the free shipping option, different levels of free shipping quantity Q, and the fixed cost K on the average costs of the retailer. We use Normal demand distribution with $\mu = 10$ unless otherwise mentioned.

Table 2
Effectiveness of the (s,t,S) policy

Demand distribution	$\pi/(\pi+h)$	$Q/E(D)$	K	Avg. Dev. 1	Avg. Dev. 2
Poisson $\lambda = 10$	0.80	[0,10]	[1,5]	0.81	13.42
	0.85	[0,10]	[1,5]	0.92	13.31
	0.90	[0,10]	[1,5]	0.74	13.25
	0.95	[0,10]	[1,5]	0.61	13.47
Poisson $\lambda = 20$	0.80	[0,10]	[1,5]	0.02	16.10
	0.85	[0,10]	[1,5]	0.05	15.79
	0.90	[0,10]	[1,5]	0.01	15.54
	0.95	[0,10]	[1,5]	0.00	16.22
Poisson $\lambda = 30$	0.80	[0,10]	[1,5]	0.00	18.88
	0.85	[0,10]	[1,5]	0.02	18.22
	0.90	[0,10]	[1,5]	0.00	17.97
	0.95	[0,10]	[1,5]	0.00	19.57
Uniform $[0,\ldots,20]$	0.80	[0,10]	[1,5]	0.00	9.17
	0.85	[0,10]	[1,5]	0.00	9.13
	0.90	[0,10]	[1,5]	0.00	9.11
	0.95	[0,10]	[1,5]	0.00	9.21
Uniform $[0,\ldots,40]$	0.80	[0,10]	[1,5]	0.00	10.27
	0.85	[0,10]	[1,5]	0.00	10.13
	0.90	[0,10]	[1,5]	0.00	9.98
	0.95	[0,10]	[1,5]	0.00	10.35

Poisson and Uniform distributions.

We first analyze the impact of free shipping option and the fixed cost K on the effect of the (s,t,S) policies. The values of K were allowed to vary from 1 to 5, incremented by 1. We compare three different types of inventory policies: (a) the optimal policy with free shipping option; (b) the optimal policy without free shipping option; and (c) the optimal (s,t,S) policy with free shipping option.

The percentage differences between the average costs of these policies are denoted by Gap 1 and Gap 2, respectively, as follows (for simplicity, optimal policy with (without) free shipping is denoted as Optimal I (II)):

$$\text{Gap 1} = 100 \times \frac{(s,t,S)\text{ average costs} - \text{Optimal I average costs}}{\text{Optimal I average costs}},$$
$$\text{Gap 2} = 100 \times \frac{\text{Optimal II average costs} - \text{Optimal I average costs}}{\text{Optimal I average costs}}.$$

We present our findings in Table 3.

The results in Table 3 illustrate that the retailer will indeed benefit from the free shipping option provided by the supplier. First, the retailer's cost savings as a result of taking the advantage of the free shipping option can reach around 13% in the event that the shipping cost K is not very large (see columns 6 and 7). Second, as K increases, average costs of both the optimal policy with free shipping and the (s,t,S) heuristic policy increase. In fact, when K becomes larger the (s,t,S) policy is closer to optimal, since both policies tend to converge to the optimal policy without free shipping which does not change with K. Third, when the shipping cost K becomes very large, the retailer will be better off ordering as least as much as Q to avoid the significant shipping cost. Finally, the results again confirm our previous findings that the heuristic policy's performance is close to that of the optimal policy.

Next, to further illustrate the impact of the shipping cost on the performance of the heuristic policies, we did a plot of the average costs as a function of different levels of shipping cost K. Fig. 3 compares the average costs of the (s,t,S) policy and those of the optimal policies with and without free shipping option at different values of K. The penalty cost ratio $\pi(\pi+h)$ was set to be 0.90, and $\sigma = 3$.

Fig. 3 confirms our results in Table 3. Indeed, notice that average costs of the (s,t,S) policy and the optimal policy with free shipping (Optimal I) increase as K increases and both policies converge to the optimal policy without free shipping (Optimal II). We conclude with the following managerial insights from our analysis:

- The free shipping option does allow the retailer to reduce its inventory costs. For a given free shipping quantity Q, the average costs under the (s,t,S) policy are significantly lower than those under the policies without free shipping option.

Table 3
The impact of free shipping option

$\pi/\pi(\pi+h)$	K	Optimal I	(s,t,S)	Optimal II	Gap 1	Gap 2
0.80	1	4.51	4.60	5.05	2.00	11.97
	2	4.72	4.78	5.05	1.27	6.99
	3	4.89	4.94	5.05	1.02	3.27
	4	4.98	5.05	5.05	1.41	1.41
	5	5.04	5.05	5.05	0.20	0.20
0.85	1	4.94	5.11	5.56	3.44	12.55
	2	5.17	5.28	5.56	2.13	7.54
	3	5.36	5.43	5.56	1.31	3.73
	4	5.48	5.54	5.56	1.09	1.46
	5	5.56	5.56	5.56	0.00	0.00
0.90	1	5.60	5.75	6.29	2.68	12.32
	2	5.86	5.93	6.29	1.19	7.34
	3	6.03	6.10	6.29	1.16	4.31
	4	6.14	6.23	6.29	1.47	2.44
	5	6.24	6.29	6.29	0.80	0.80
0.95	1	6.50	6.72	7.35	3.38	13.08
	2	6.78	6.91	7.35	1.92	8.41
	3	6.99	7.09	7.35	1.43	5.15
	4	7.15	7.23	7.35	1.12	2.80
	5	7.27	7.35	7.35	1.10	1.10

$Q = 10$, $\sigma = 3$.

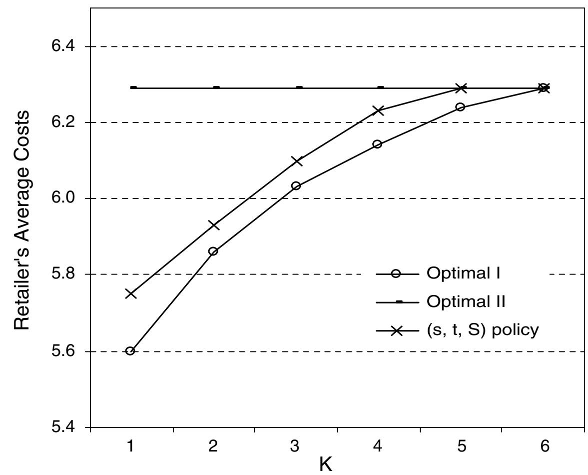

Fig. 3. The Impact of shipping cost K on retailer's average costs.

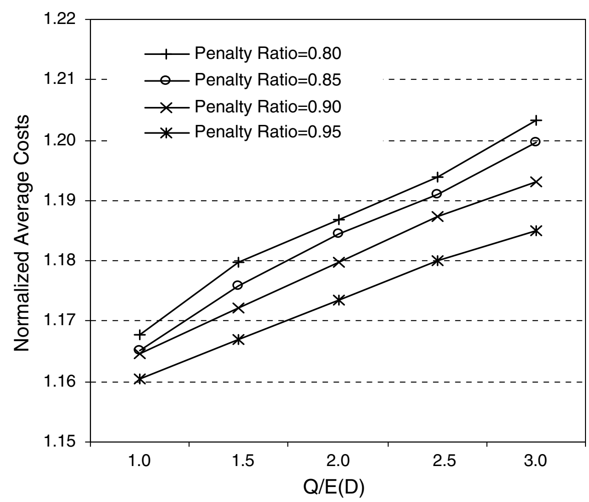

Fig. 4. Impact of the free shipping quantity.

- As the shipping cost K increases, the retailer's average costs increase, and the cost savings diminish when K becomes larger.
- When K is very large, the retailer will be better off ordering at least the free shipping quantity Q to receive free shipping and therefore reduce the total inventory costs.

Lastly, we plot the long-run average costs as a function of $Q/E(D)$ under different penalty cost ratios $\pi/(\pi + h)$. The average costs are normalized by the global minimum value of $L(y)$ of the optimal policy. Fig. 4 depicts our results for the case of $\sigma = 3$ and $K = 3$. The insight obtained herein is consistent with our intuition:

- For a given fixed cost K, the average costs of the retailer increase as the free shipping quantity Q increases, and higher free shipping quantity requirement results in higher average costs for the retailer.

7. Summary and conclusion

In this paper, we considered the problem of a stochastic inventory system with free shipping option and a fixed cost K. The problem arises in situations where a supplier offers free shipping with linear ordering cost if the order size of the retailer is at least as much as the free shipping quantity; an additional fixed shipping cost K is applied otherwise. We formulated the problem using stochastic dynamic programming and provided the optimal policy for the single-period inventory system. For the multi-period case, we introduced a new class of simple but effective heuristic policy, the (s, t, S) policy.

By conducting extensive numerical studies, we illustrated that: (a) the new (s, t, S) heuristic policy has a robust performance and it provides close-to-optimal solutions for a wide range of system parameters and different demand distributions;

(b) for inventory systems with free shipping, the heuristic policy outperforms the (s,t) policy that only considers minimum order quantity; (c) the heuristic policy tends to have better performance when the demand variability is larger. Furthermore, we discussed and provided important managerial insights. Our analysis demonstrated that the free shipping option can indeed improve retailer's inventory management and generate considerable cost savings. As the shipping cost K increases, the retailer's inventory costs increase (cost savings diminish). And finally, higher values of free shipping quantity Q lead to higher average costs for the retailer. Throughout the paper, we assumed that $\pi > c$ to exclude the cases in which it is never optimal for the retailer to carry inventory before demand is realized.

There are several research questions that can stem from this current paper. We have not found the optimal policies for the single-item, multi-period inventory system. One interesting extension worth pursuing is to investigate and obtain some structural properties of the true optimal inventory policies for multi-period systems with free shipping option. Another future research effort may consider the free shipping problem in multi-item, multi-period stochastic inventory systems where the retailer replenishes multiple products from the same supplier, who offers free shipping if the total order size of all products is no less than the free shipping quantity (FSQ) and requires the retailer to pay an additional fixed cost if the total order size is below the FSQ. Questions may include: What are the structures of the optimal policy for the multi-item inventory systems? Are there any effective heuristics that can provide good approximation for the problem?

References

Aviv, Y., Federgruen, A., 2001. Capacitated multi-item inventory systems with random and seasonally fluctuating demands: Implications for postponement strategies. Management Science 47, 512–531.

Bradley, J., Robinson, L., 2005. Improved base-stock approximations for independent stochastic lead times with order crossover. Manufacturing and Service Operations Management 7, 319–329.

Chen, F., Zheng, Y., 1994. Lower bounds for multi-echelon stochastic inventory problems. Management Science 40, 1426–1443.

Clark, A., Scarf, S., 1960. Optimal policies for a multi-echelon inventory problem. Management Science 6, 475–490.

Derman, C., 1970. Finite State Markovian Decision Processes. Academic Press, New York.

Federgruen, A., Zipkin, P., 1984. Computational issues in an infinite horizon, multi-echelon inventory model. Operations Research 32, 818–836.

Fisher, M., Raman, A., 1996. Reducing the cost of demand uncertainty through accurate response to early sales. Operations Research 44, 87–99.

Heyman, D., Sobel, M., 1984. In: Stochastic Models in Operations Research, vol. 2. McGraw Hill, New York.

Kaplan, R., 1970. A dynamic inventory model with stochastic lead times. Management Science 17, 491–507.

Robb, D., Silver, E., 1998. Inventory management with periodic ordering and minimum order quantities. The Journal of the Operational Research Society 49, 1085–1094.

Ross, S., 1992. Applied Probability Models with Optimization Applications. Dover, New York.

Scarf, H., 1960. The optimality of (s,S) policies in the dynamic inventory problem. In: Arrow, K., Karlin, S., Suppes, P. (Eds.), Mathematical Methods in the Social Sciences. Stanford University Press, Stanford.

Silver, E., Pyke, D., Peterson, R., 1998. Inventory Management and Production Planning and Scheduling. Wiley, New York.

Tijms, H.C., 1994. Stochastic Models: An Algorithmic Approach. Wiley, New York.

Veinott, A.F., 1966. The status of mathematical inventory theory. Management Science 12, 745–777.

Zalkind, D., 1978. Order-level inventory systems with independent stochastic leadtimes. Management Science 24, 1384–1392.

Zhao, Y., Katehakis, M.N., 2006. On the structure of optimal ordering policy of stochastic inventory systems with minimum order quantity. Probability in the Engineering and Information Sciences 20, 257–270.

Zheng, Y., Federgruen, A., 1991. Finding the optimal (s, S) policies is about as simple as evaluating a single policy. Operations Research 39, 654–665.

Zhou, B., Zhao, Y., Katehakis, M.N., 2007. Effective control policies for stochastic inventory systems with a minimum order quantity and linear costs. International Journal of Production Economics 106, 523–531.