Available online at www.sciencedirect.com
ScienceDirect
Int. J. Production Economics 106 (2007) 523–531
international journal of production economics
www.elsevier.com/locate/ijpe

Effective control policies for stochastic inventory systems with a minimum order quantity and linear costs ☆

Bin Zhou, Yao Zhao*, Michael N. Katehakis

Department of Management Science and Information Systems, Rutgers University, the State University of New Jersey, 180 University Avenue, Newark, NJ 07102, USA

Received 11 November 2005; accepted 27 June 2006
Available online 1 November 2006

Abstract

We consider a model of single-item periodic-review inventory system with stochastic demand, linear ordering cost, where in each time period, the system must order either none or at least as much as a minimum order quantity (MOQ). Optimal inventory policies for such a system are typically too complicated to implement in practice. In fact, the (s, S) type of policies are often utilized in the real world. We study the performance of a simple heuristic policy that is easily implementable because it is specified by only two parameters (s, t). We develop an algorithm to compute the optimal values for these parameters in the infinite time horizon under the average cost criterion. Through an extensive numerical study, we demonstrate that the best (s, t) heuristic policy has performance close to that of the optimal policies when the coefficient of variation of the demand distribution is not very small. Furthermore, the best (s, t) policy always outperforms the best feasible (s, S) policies and on average the percentage differences are significant. Finally, we study the impact of MOQ on system performance.

Keywords: Minimum order quantity; Stochastic inventory systems; Heuristic policy; Multiple periods

1. Introduction

Economies of scale are important concerns in many industries. This is especially true for companies that manufacture and/or distribute pharmaceutical, apparel, consumer packaged goods or chemical products. Most companies use one of the following three ways to achieve the economies of scale in production and distribution: charging fixed ordering cost, or requiring batch orders, e.g., full truck load, or setting minimum order quantity (MOQ) for their customers.

MOQ is widely used in practice. A celebrated example of MOQ is the fashion sport ski-wear manufacturer and distributor: Sport Obermeyer (Hammond and Raman, 1996). While the production base of Sport Obermeyer in Hong Kong sets a MOQ of 600 garments per order, its production base in China requires 1200 garments. The MOQ constraints can be applied to a particular item or a group of items, such as all colors of a particular style. We refer the reader to Hammond and Raman (1996) for more detailed descriptions of the

☆Research supported in part by a Faculty Research Grant from Rutgers Business School–Newark and New Brunswick.
*Corresponding author. Tel.: +1 973 353 5003.
E-mail address: yaozhao@andromeda.rutgers.edu (Y. Zhao).

0925-5273/$ - see front matter
doi:10.1016/j.ijpe.2006.06.020

managerial situation. MOQ is not only required in fashion industries where customers, e.g., retailers, only place orders once or twice in a selling season, it also applies to many repeatedly ordered items carried by large retail chains such as Home-Depot and Wal-Mart stores. When a supplier sets a MOQ, the ordering cost from that supplier is typically a linear function of the ordering quantity because the economies of scale have been accounted for by the MOQ. Large MOQs present substantial challenges to the efficient management of supply chains, because they require customers (e.g., retailers) to order either none or many units, therefore they reduce customer flexibility in responding to demand, and eventually increase their inventory costs.

To provide tools and principles for companies that can guide their managerial actions in controlling inventory when they face MOQ from their suppliers, we consider a model of a single-item inventory system in which the demand is random and the ordering cost is linear. At each time period, the system can either order none or order at least as much as the MOQ. Our objective is two-fold: (1) designing effective inventory control policies that are easily implementable and computationally tractable; (2) developing insights with respect to the impact of the MOQ.

It was shown by Zhao and Katehakis (2006) that the optimal policies for such systems are typically too complicated to implement in practice. In fact, the (s, S) type of policies (the min–max policies) are commonly used in the real world for the inventory systems with MOQ. Herein, we propose and analyze a new class of inventory control policies with a simple structure, namely the (s, t) policies, with $s \leqslant t < s + M$, where M represents the MOQ. The (s, t) policy works as follows: when the initial inventory position is lower than or equal to s, order upto $s + M$; when the initial inventory exceeds s but is no more than t, order exactly M; otherwise, do not order. We provide bounds for the optimal t and develop an algorithm to compute the optimal values for s and t. Throughout the paper, optimality refers to the infinite time horizon average cost criterion.

To demonstrate the effectiveness of the (s, t) policy, we conduct an extensive numerical study to compare the performance of the best (s, t) policy to that of the optimal policy as well as that of the best feasible (s, S) policy with $S - s \geqslant M$. In practice, many companies have a build-in periodic review structure and utilize the min–max policies to control inventories, i.e., if the inventory drops below the minimum level s, an order is placed that will refill the stock up to the maximum level S. Under the constraint of MOQ, the difference between S and s is no less than the MOQ as demanded by the supplier. Indeed, the difference between S and s is often set exactly to the MOQ due to the lack of fixed ordering cost. Our numerical study demonstrates that the (s, t) policy has close to optimal performance under reasonable conditions, and it significantly outperforms the best feasible (s, S) in all instances. We also study the trade-off between inventory costs and the MOQ, and the interplay of the MOQ and various parameters, e.g., demand variability, penalty and holding costs.

2. Literature review

Extensive research has been conducted on stochastic production-inventory systems in which the economies of scale in production and transportation are main concerns. Most of the literature focuses on models with either fixed ordering costs or fixed batch sizes, see Veinott (1966), Heyman and Sobel (1984), Chen (1998) and Zipkin (2000) for excellent reviews, while little work has been done on systems with MOQ.

Chan and Muckstadt (1999) analyze a production-inventory system in which the production quantity is constrained by a minimum and a maximum level in each period, i.e., the production smoothing problem. They characterized the optimal policy in finite and infinite time horizons under the discounted cost criterion. The problem of MOQ is different from the production smoothing problem because the order quantity in the former is either zero or at least the MOQ, and therefore the action sets are disjoint and unbounded, while in the later, the action sets are connected and compact. It was shown in Zhao and Katehakis (2006) that these disjoint action sets significantly complicate the structure of the optimal policy.

Fisher and Raman (1996) is perhaps the first paper considering MOQ in a stochastic inventory system. Its focus is on fashion items with short product life-cycles and therefore the inventory system under study has only two review periods. There are multiple items subject to the MOQ requirements as well as a production capacity constraint. The paper formulates the problem into a stochastic program, and quantifies the impact of MOQ on the inventory costs for fashion items.

Zhao and Katehakis (2006) consider a single-item stochastic inventory system with MOQ in finite and infinite time horizons under the discounted cost criterion. The model is suitable for repeatedly ordered items. The authors characterized the optimal ordering policy everywhere in the state space outside of an interval for each time period, and develops simple upper bound and asymptotic lower bound for these intervals. In a simple example with two review periods, Zhao and Katehakis (2006) demonstrate the complexity of the optimal policy by showing that the cost functions may have multiple local minimums in these intervals.

The main contribution of this paper is the performance analysis and optimization of the easily implementable (s,t) policies. We demonstrate by numerical studies that they outperform the (s,S) policies and have performance close to the optimal policies under a reasonable condition of demand distribution. The rest of this paper is organized as follows: in Section 3, we formally define the model. In Section 4, we present and analyze the (s,t) policy and develop an algorithm to determine its optimal parameters. An extensive computational study is conducted in Section 5 to demonstrate the effectiveness of the (s,t) heuristic policy and to develop managerial insights. Lastly, conclusions are summarized in Section 6.

3. The model

We consider a periodic review inventory system managing a single item. The demand D for this item are i.i.d. random variables with finite mean $E(D)$. At each time period, the inventory system can either order none or order any amount as long as it equals to or exceeds the MOQ, M. There is no fixed ordering cost, but there is an inventory holding cost h per unit per period, and a penalty cost π per unit per period. We set the purchasing cost equal to zero because linear ordering costs can be ignored under the average cost criterion, see, e.g., Veinott and Wagner (1965) and Zheng and Federgruen (1991). For the ease of exposition, we assume that the retailer faces zero replenishment lead-time. The model can be easily extended to systems with positive replenishment lead-times following the standard procedure as in Heyman and Sobel (1984, p. 75).

The sequence of events is as follows. At the beginning of a time period, the retailer reviews the inventory and places 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 stock. If the retailer cannot fulfill all the demand, the excessive amount is backlogged. Let x be the initial inventory position at the beginning of a time period, and y be the inventory position after order is placed, the single-period cost function can be written as

$$L(y) = E_D[h(y-D)^+ + \pi(D-y)^+],$$

where y either equals x or y is not smaller than $x+M$. x and y are integers. Clearly, $L(y)$ is a convex function and $L(y) \to +\infty$ as $|y| \to \infty$. Let y^* be the smallest global minimum of $L(y)$.

4. The heuristic policy

As Zhao and Katehakis (2006) point out, the optimal policies for multi-period stochastic inventory systems with MOQs do not have a simple structure. Indeed, the optimal policies are too complicated to implement in practice. Therefore, it is of practical importance to identify and analyze easily implementable heuristic policies which also have close to optimal performance under reasonable conditions. Based on the analysis of the single-period problem as well as multiple period Markov decision process associated with the stochastic inventory systems with MOQs (see Zhao and Katehakis, 2006), we propose the (s,t) policy: given integers s,t where $s \leqslant t < s+M$, and an initial inventory position x, the policy is to order $y(x)-x$, where

$$y(x)-x = \begin{cases} s+M-x & \text{if } x \leqslant s, \\ M & \text{if } s < x \leqslant t, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

That is, when x is smaller than or equal to s, order upto $s+M$; when x is greater than s but smaller than or equal to t, order exactly the MOQ; and when x is above t, order nothing. Note that when $t=s$, the (s,t) policy reduces to a (s,S) policy with $S-s=M$.

To identify the optimal s and t that minimize the long-run average cost within this class of policies, we focus on the inventory position y at the beginning of a time period after the order decision is made. y can choose one of these values, $t+1, t+2, \ldots, t+M$, and it evolves according to a discrete time Markov chain (DTMC) with the transition matrix P, where

$$P_{i,j} = \begin{cases} p_{(i-j)^+} & \text{for } j = t+1, t+2, \ldots, s+M-1 \\ & \quad \forall i = t+1, \ldots, t+M, \\ \sum_{k=i-s}^{+\infty} p_k + p_{(i-j)^+} & \text{for } j = s+M \\ & \quad \forall i = t+1, \ldots, t+M, \\ p_{i-j+M} + p_{(i-j)^+} & \text{for } j = s+M+1, s+M+2, \ldots, t+M \\ & \quad \forall i = t+1, \ldots, t+M, \end{cases} \tag{2}$$

where $p_i = P\{D = i\}$, $p_{(i-j)^+}$ equals to p_{i-j} if $i \geqslant j$, and zero otherwise. Let q_i be the long-run fraction of time such that $y = t + i$, and $\overline{q} = \{q_1, q_2, \ldots, q_M\}$. If $\{t+1, t+2, \ldots, t+M\}$ forms a close communicating class, then,

$$\overline{q} = \overline{q}P,$$

$$\sum_{i=1}^{M} q_i = 1. \tag{3}$$

The steady-state distribution of the inventory position is not uniform in general. Using renewal reward process, Veinott and Wagner (1965) derive the exact expression of the long-run average cost of the (s, S) policies, and Zheng and Federgruen (1991) develop an efficient method to calculate the optimal (s, S) policy. The key idea is to identify the expected total cost in the cycle between two successive orders where a cycle always starts with inventory position S. However, in the (s, t) policies, the inventory position just after ordering can take any value from $s + M$, $s + M + 1, \ldots,$ to $t + M$, instead of a unique value S in the case of (s, S) policies. Thus, there are multiple types of cycles where each starts with a different inventory position. To calculate the long-run average cost using renewal reward theory, one needs to identify the steady-state probability of a cycle starting with every possible inventory position. However, it is not clear how to calculate these steady-state probabilities more efficiently than directly solving the linear system in Eq. (3).

Observe that given any M and demand distribution, the transition matrix P and therefore $\overline{q}$ only depends on $\Delta = t - s$, we define the long-run average cost $C(\Delta, t) = \sum_{i=1}^{M} q_i L(t+i)$. We next establish an important property for the heuristic policy.

Proposition 1. *For a given Δ, the smallest optimal t^* satisfies $t^* < y^* \leqslant t^* + M$ and $C(\Delta, t)$ is convex in t.*

Proof. To prove these bounds for t^*, we focus on $L(y)$ since the q_i remain as constants for a fixed Δ. Assume that $y^* \leqslant t^*$. Since $L(y)$ is nondecreasing for $y \geqslant y^*$, $L(t^* + i - 1) \leqslant L(t^* + i)$ for $i = 1, 2, \ldots, M$, hence $C(\Delta, t^* - 1) = \sum_{i=1}^{M} q_i L(t^* - 1 + i) \leqslant \sum_{i=1}^{M} q_i \; L(t^* + i) = C(\Delta, t^*)$. This contradicts with the definition of t^*.

Similarly, assume that $t^* + M < y^*$. Since $L(y)$ is nonincreasing for $y \leqslant y^*$ and y^* is the smallest global minimum, $L(t^* + i + 1) < L(t^* + i)$ for $i = 1, 2, \ldots, M$, which again contradicts with the definition of t^*. Finally, the convexity of $C(\Delta, t)$ follows directly from the convexity of $L(y)$.

For a given Δ, we only need to calculate $\overline{q}$ once for all t. We design the following algorithm to compute the t^* and s^* that minimize the long-run average cost within the class of the (s, t) heuristic policies:

1. For each $\Delta \in \{0, 1, 2, \ldots, M-1\}$, calculate P and solve Eq. (3) for $\overline{q}$.
2. For each Δ and $\overline{q}$, identify $t = t(\Delta)$ from the set $\{y^* - M, y^* - M + 1, \ldots, y^* - 1\}$ that minimizes $C(\Delta, t)$.
3. Find Δ^* that minimizes $C(\Delta, t(\Delta))$, and let $t^* = t(\Delta^*)$ and $s^* = t^* + \Delta^*$.

The first step requires a computational time proportional to $O(M^4)$ if a Gaussian elimination with partial pivots is used to solve the linear system in Eq. (3) with rank M. The second step requires a computational time proportional to $O(\frac{1}{2}M^2)$ because it can take advantage of the convexity of $C(\Delta, t)$. For real-world systems in which M is hundreds and even thousands, we can improve the efficiency of the algorithm by discretizing demand distribution into bins of appropriate size. □

5. Computational study

The objective of this section is two-fold: first, to demonstrate the effectiveness of the (s, t) policy, and second, to develop insights on the impact of MOQ.

We conduct numerical studies with respect to the following nondimensional parameters, $M/E(D)$, $\pi/(\pi+h)$ (the penalty cost ratio) and demand coefficient of variation (c.v.). We assume that the demand in one period follows normal distribution with $E(D)=10$ unless otherwise mentioned. To ensure nonnegative demand, we truncate the normal distribution and let $P\{D=0\}$ be the probability that the normal random variable is equal to or small than zero. Normal distribution is one of the commonly used distributions for inventory applications; we refer the reader to Nahmias (2001, p. 246–247) and Silver et al. (1998, p. 272–273) for detailed justification of its significance.

5.1. Effectiveness of the (s,t) policy

We study the performance of the (s,t) policy by comparing the long-run average costs of the best (s,t) policies to those of the optimal policies. Due to the zero fixed ordering cost, many companies utilize simple (s,S) policies with $S-s=M$ to control inventories in practice. To demonstrate the effectiveness of the best (s,t) policy, we also compare its performance to that of the best feasible (s,S) policy with $S-s\geqslant M$.

To calculate the optimal policy under the average cost criterion, we solve numerically the average cost optimality equations of the inventory system with MOQ using value iteration, see e.g., Veatch and Wein (1996) or Bertsekas (1995). The state space is truncated, and its size is determined by testing larger and larger state space until the results are insensitive to the increments.

The numerical examples are chosen as follows. The holding cost h is set to 1. M varies from 0 to 50 in increment of 1, $\pi/(\pi+h)$ (penalty ratio for convenience) takes values of 0.80, 0.85, 0.90, or 0.95, and the demand c.v. varies from $0.1, 0.2, 0.3$ to 0.4. Altogether, the combination of these choices results in 816 instances in the numerical study. For each instance, we measure the effectiveness of the best (s,t) policy with respect to the optimal policy by the % gap, G_1, between the costs of these policies, as follows:

$$G_1 = 100 \times [\text{best } (s,t) \text{ average cost} - \text{optimal average cost}]/\text{optimal average cost}.$$

For each instance, we also measure the effectiveness of the best (s,t) policy with respect to the best feasible (s,S) policy by the % gap, G_2, between the costs of these policies, as follows:

$$G_2 = 100 \times [\text{best feasible } (s,S) \text{ average cost} - \text{best } (s,t) \text{ average cost}]/\text{best } (s,t) \text{ average cost}.$$

Table 1 summarizes the maximum and average % gaps, G_1, over all instances between the optimal policies and the best (s,t) (in columns 4 and 5), and maximum and average % gaps, G_2, between the best (s,t) policy and the best feasible (s,S) policy (in columns 6 and 7).

Table 1 demonstrates that the performances of the best (s,t) policies are very close to those of the optimal policies for most combinations of the parameters except for a few cases with very small values of the demand coefficient of variation, e.g., c.v. $=0.1$. For instance, at c.v. $=0.2$, the maximum and average G_1 vary between 0.72% and 1.87% and between 0.03% and 0.13%, respectively. Similar results are obtained at c.v. $=0.3$ and 0.4. Indeed, as c.v. increases, G_1 tends to decrease. When c.v. $=0.4$, both the maximum and average G_1 reduce to nearly 0%. These results indicate that the best (s,t) policy has a close to optimal performance for a wide range of system parameters, and it tends to perform better as demand variability increases. Fig. 1 shows that the long-run average costs (normalized by the global minimum value of $L(y)$) of the optimal policy and the best (s,t) policy almost completely overlap for c.v. $=0.2, 0.3$, and 0.4.

The only cases in which the (s,t) policies may not perform as well are those at very small demand coefficient of variation, e.g., c.v. $=0.1$. The maximum % gap between the optimal policy and the

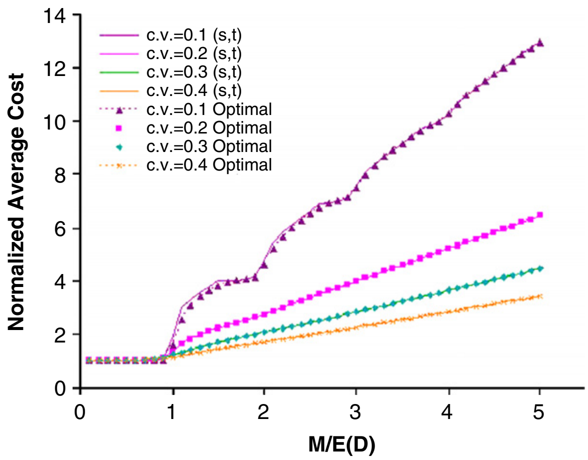

Fig. 1. The best (s,t) policy vs. the optimal policy. $\pi/(\pi+h)=0.90$.

best (s, t) policy can be substantial: it is 16.65% when $\pi/(\pi+h) = 0.80$, and 24.51% when $\pi/(\pi+h) = 0.95$. However, the average % gaps are much smaller, varying from 1.16% to 2.15% (see also Fig. 1).

Large % gaps in these cases clearly indicate that the optimal policies do not have the same structure as the (s, t) policy. To develop insights into the structure and the complexity of the optimal policies, we compute the optimal policy for the case corresponding to the maximum G_1 in column 4 at $\pi/(\pi+h) = 0.90$ (where $M/(E(D) = 1)$). The optimal policy as well as the best (s, t) policy are presented in Fig. 2 in which the X-axis is the initial inventory position x and Y-axis is the optimal ordering quantities $y - x$.

Fig. 2 demonstrates that the optimal ordering policy under the average cost criterion may not have the same structure as the (s, t) policy for normal demand distribution with small coefficient of variation. In the above example, the optimal policy is identical to the best (s, t) policy outside the interval [5, 10] where it has a "peak", i.e., it orders more than 10 (which equals to the MOQ) units of the (s, t) policy. In this example, it is clear that the (s, t) policy is much simpler than the optimal policy because it eliminates the "peak". We also study more general demand distributions with multiple modes. The first demand distribution, denoted by *discrete* 1, has two modes where $D = \{5, 6, 7, 8, 9, 10, 11, 12\}$ and the probability density function is $p = \left\{\frac{3}{10}, \frac{1}{60}, \frac{1}{60}, \frac{1}{60}, \frac{1}{60}, \frac{1}{60}, \frac{1}{60}, \frac{3}{5}\right\}$. The second distribution, denoted by *discrete* 2, has three modes where $D = \{5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15\}$ and the probability density function is $p = \{0.2, 0.0125, 0.0125, 0.0125, 0.0125, 0.2, 0.0125, 0.0125, 0.0125, 0.0125, 0.5\}$. Then numerical study of these examples (not reported here) shows that the optimal policy can have multiple "peaks", and the (s, t) policy essentially eliminates these peaks by ordering exactly the MOQ.

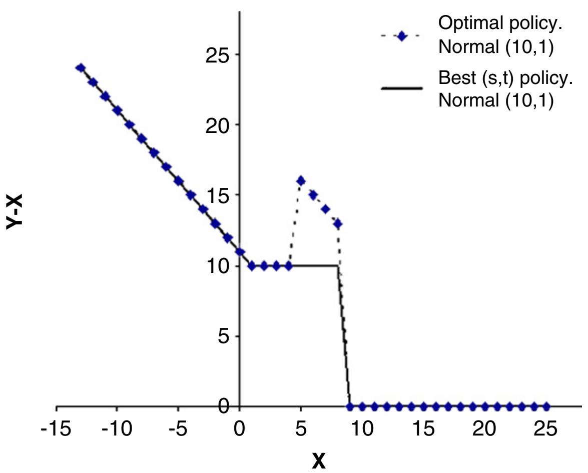

Fig. 2. The ordering quantities under the optimal policy and the best (s, t) policy.

These examples are consistent with the analysis of Zhao and Katehakis (2006) which considers a finite horizon model under the discounted cost criterion. In a two-period example, Zhao and Katehakis (2006) show that the cost functions may have multiple local minimums which result in the "peaks". These examples demonstrate that the "peaks" carry on to the infinite horizon model under the average cost criterion. In such cases, the optimal policy does not have a simple form.

We next compare the best (s, t) policies to the best feasible (s, S) policies with $S - s \geqslant M$. We compute the best feasible (s, S) policy by enumerating all possible combinations of s and S in a sufficiently large region, and evaluating the performance of each combination by Zheng and Federgruen (1991). Table 1 shows that the best (s, t) policies perform better than the best feasible (s, S) policies in all instances. In addition, the % gaps, G_2, between these two policies are substantial. For instance, Table 1 demonstrates that when c.v. $= 0.1$, the maximum G_2 is ranging from 154% to 93.51% for different $\pi/(\pi+h)$ ratios, while the average G_2 varies from 15.88% to 16.87%. As c.v. increases, the gaps between the best (s, t) policies and the best feasible (s, S) policies tend to decrease, however, even at c.v. $= 0.4$ and $\pi/(\pi+h) = 0.90$, the maximum and average G_2 are still as large as 17.19% and 9.15%, respectively. The best (s, t) policies tend to outperform the best feasible (s, S) policies for two reasons: (1) the optimal policies for the single-period problems have the same form as the (s, t) policies (Zhao and Katehakis, 2006). (2) When $S - s = M$, the (s, S) policies are special cases of the (s, t) policies.

To better illustrate the effectiveness of the (s, t) policy with respect to the optimal policy and the (s, S) policy, we plot in Fig. 3 the normalized (by the global minimum value of $L(y)$) long-run average costs of the optimal policy, the best (s, t) policy and the best feasible (s, S) policy as functions of $M/E(D)$, for the "unfavorable" values of c.v. $=$ 0.1 or 0.3. The figure shows that the long-run average costs of the best (s, t) policies are much

Table 1
The effectiveness of the (s, t) policy

Demand			Best (s,t) policy (G_1)				Best (s,S) policy (G_2)			
$E(D)=10$	$\frac{\pi}{\pi+h}$	$\frac{M}{E(D)}$	Max G_1	Avg G_1	$\frac{M}{E(D)}=3$	$=5$	Max G_2	Avg G_2	$\frac{M}{E(D)}=3$	$=5$
c.v. $=0.1$	0.80	[0, 5]	16.65	1.16	0.01	0.00	154.12	16.87	25.80	9.77
	0.85	[0, 5]	18.17	1.49	0.00	0.00	130.01	16.44	30.17	8.44
	0.90	[0, 5]	22.37	1.84	0.67	0.02	111.36	15.88	28.31	11.43
	0.95	[0, 5]	24.51	2.15	1.75	0.10	93.51	16.01	24.67	12.17
c.v. $=0.2$	0.80	[0, 5]	0.72	0.03	0.00	0.00	73.31	16.22	16.95	8.63
	0.85	[0, 5]	1.08	0.05	0.00	0.00	80.21	17.01	17.64	8.57
	0.90	[0, 5]	1.03	0.08	0.00	0.00	83.31	16.79	17.52	8.45
	0.95	[0, 5]	1.87	0.13	0.00	0.00	81.54	16.61	16.94	8.58
c.v. $=0.3$	0.80	[0, 5]	0.01	0.00	0.00	0.00	28.03	10.47	9.97	8.92
	0.85	[0, 5]	0.02	0.00	0.00	0.00	28.84	10.41	9.95	6.97
	0.90	[0, 5]	0.04	0.00	0.00	0.00	28.34	10.02	9.59	6.78
	0.95	[0, 5]	0.06	0.00	0.00	0.00	28.22	9.90	9.36	6.53
c.v. $=0.4$	0.80	[0, 5]	0.00	0.00	0.00	0.00	17.57	9.24	10.13	7.24
	0.85	[0, 5]	0.00	0.00	0.00	0.00	17.62	9.34	10.07	7.19
	0.90	[0, 5]	0.00	0.00	0.00	0.00	17.19	9.15	9.60	7.01
	0.95	[0, 5]	0.00	0.00	0.00	0.00	17.26	8.75	8.92	6.79

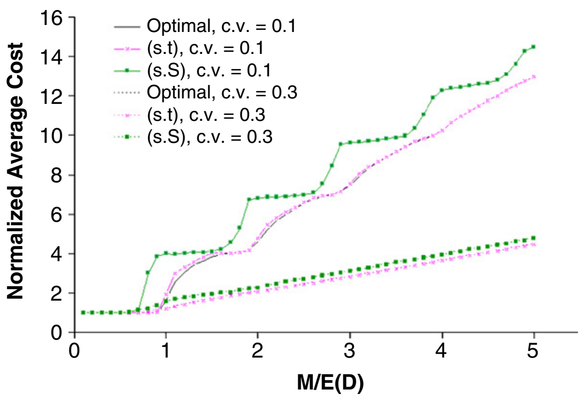

Fig. 3. The best (s, t) policy vs. the optimal policy vs. the best (s, S) policy.

closer to those of the optimal policies than the best feasible (s, S) policies. Indeed, there are no significant differences between the best (s, t) policies and the optimal policies, while the differences between the best feasible (s, S) policies and the optimal policies are quite substantial.

It is clear from Figs. 1 and 3 that the gaps among the optimal, (s, t) and (s, S) policies are negligible when $M/E(D)$ is relatively small ($\leqslant 0.5$). As $M/E(D)$ increases, the gaps may increase but not always. As demonstrated by Fig. 3 and Table 1 columns "$M/E(D) = 3$" and "$= 5$", the gaps vary for different MOQs.

For comparison, we study the case of deterministic demand where $P\{D = 10\} = 1$ (i.e., c.v. $= 0$). As before, the MOQ varies from 0 to 50, and penalty cost ratio takes values of 0.80, 0.95, 0.90 or 0.95. We compute the average costs under the optimal policy, the best (s, t) policy, and the best feasible (s, S) policy, respectively. Our numerical results show that the maximum and average gaps between the optimal policy and the best (s, t) policy continue to increase as c.v. decreases from 0.1 to 0. However, the gaps between the best (s, t) policy and the best feasible (s, S) policy are smaller at c.v. $= 0$ relative to c.v. $= 0.1$. This can be explained as follows: at c.v. $= 0.1$, the gap between the best (s, t) policy and the best feasible (s, S) policy reaches its maximum when $M/E(D)$ is an integer. Whereas at c.v. $= 0$, i.e., deterministic demand, the gap between the best (s, t) policy and the best feasible (s, S) policy is zero when $M/E(D)$ is an integer. Finally, we observe that gaps are always positive which implies that the best (s, t) policies outperform the best feasible (s, S) policies.

To summarize, our computational study reveals that

- In general, the optimal policies of the inventory systems with MOQ have complex structures that make them difficult to identify and implement in practice.

- The best (s,t) policies have close to optimal performance under reasonable conditions.
- Although the (s,t) policy can be substantially inferior to the optimal policy in the special cases of small demand c.v., it always and significantly outperforms the best feasible (s,S) policy.

5.2. Sensitivity studies

Fig. 1 illustrates that the long-run average costs at different demand c.v.s increase as $M/E(D)$ increases, indicating that higher MOQ requirement leads to higher average cost. However, the costs are in general not convex functions of M. The convex relationship between the cost and M was observed first by Hammond and Raman (1995) for the single-period problem. In the multi-period case studied herein, the optimal cost-to-go functions of the multiple period problem are typically not convex, and therefore, the convex relationship in general does not hold.

We next quantify the impact of demand variability on the effect of MOQ. To this end, we refer to Fig. 1 for the normalized average cost as a function of $M/E(D)$ at different levels of c.v. $= 0.1, 0.2, 0.3$, and 0.4. The results based on both the (s,t) policy and the optimal policy reveal that

- When M is relatively large, e.g., greater than $E(D)$, a higher level of demand variability leads to a lower rate of cost increase as $M/E(D)$ increases.
- When M is relatively small, e.g., less than $E(D)$, the demand variability does not have a significant impact on the normalized marginal costs.

We finally study the long-run average cost as a function of $M/E(D)$ under different levels of $\pi/(\pi+h)$ in Fig. 4. The costs are those corresponding to the best (s,t) policy, and are normalized by the global minimum value of $L(y)$. Fig. 4 shows that

- When $M/E(D)$ is relatively small, e.g. less than or equal to 1, $\pi/(\pi+h)$ does not have a significant impact on the normalized marginal costs.
- When $M/E(D)$ is relatively large, the average cost at a higher value of $\pi/(\pi+h)$ increases at a slower rate than that at a lower value of $\pi/(\pi+h)$.

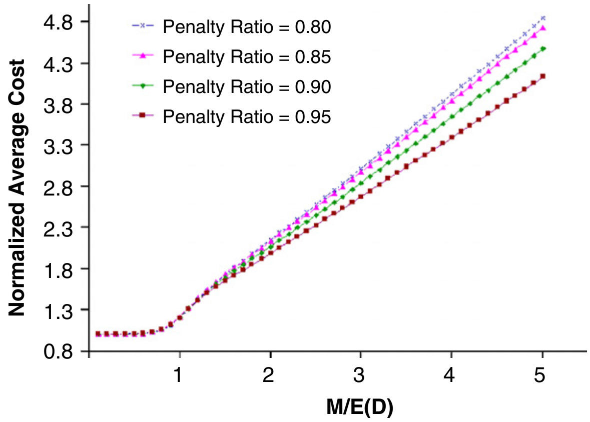

Fig. 4. The impact of $M/E(D)$ and $\pi/(\pi+h)$. c.v. $= 0.3$.

6. Conclusion

In this paper, we analyze a class of simple heuristic policies, the (s,t) policies, to control stochastic inventory systems with minimum order quantities. Policies in this class are easily implementable in practice. In the case of the average cost criterion, we demonstrate the effectiveness of the optimal (s,t) policy within the class with respect to the optimal policy as well as the widely used (s,S) policies. We also develop insights into the impact of MOQ on repeatedly ordered items.

Acknowledgments

We thank Peter Kelle and three anonymous referees for constructive comments and suggestions that have allowed us to improve the paper.

References

Bertsekas, D.P., 1995. Dynamic Programming and Optimal Control, vol. 2. Athena Scientific, Belmont, MA.

Chan, E.W., Muckstadt, J.A., 1999. The effects of load smoothing on inventory levels in a capacitated production inventory system. Technical Report, School of Operations Research and Industrial Engineering. Cornell University, Ithaca, NY.

Chen, F., 1998. On (R,NQ) policies in serial inventory systems. In: Tayur, S., Ganeshan, R., Magazine, M.J. (Eds.), Quantitative Models for Supply Chain Management (International Series in Operations Research & Management Science, vol. 17). Kluwer Academic Publishers, Dordrecht (Chapter 4).

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

Hammond, J.H., Raman, A., 1995. Teaching note for the case of "Sport Obermeyer, Ltd". Harvard Business School.

Hammond, J.H., Raman, A., 1996. Sport Obermeyer, Ltd., Harvard Business School Case Study.

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

Nahmias, S., 2001. Production and Operations Analysis, fourth ed. McGraw-Hill Irwin, Boston.

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

Veatch, M.H., Wein, L.M., 1996. Scheduling a make-to-stock queue: index policies and hedging points. Operations Research 44, 634–647.

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

Veinott, A.F., Wagner, H., 1965. Computing optimal (s,S) inventory policies. Management Science 11, 525–552.

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

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

Zipkin, P., 2000. Foundations of Inventory Management. McGraw Hill, Boston.