DEPARTMENT OF OPERATIONS RESEARCH
UNIVERSITY OF AARHUS

Publication no. 99/4

A Two-Echelon Inventory Model with Lost Sales

Jonas Andersson
Philip Melchiors

ISSN 1398-8964

Department of Mathematical Sciences
Telephone: +45 8942 3188

Building 530, Ny Munkegade
DK-8000 Aarhus C, Denmark

A Two-Echelon Inventory Model with Lost Sales

JONAS ANDERSSON
Division of Production Management,
Department of Industrial Engineering,
University of Lund,
Box 118, SE-22100 Lund,Sweden

PHILIP MELCHIORS*
Department of Operations Research,
University of Aarhus,
Ny Munkegade, Building 530,
DK-8000 Aarhus C, Denmark.

SEPTEMBER 1999

Abstract

Almost all multi-echelon inventory models assume that demand not satisfied immediately can be backordered. In some situations this assumption is not realistic. For example, it may be more representative to model stockouts as lost sales when the retailers are in a competitive market and customers can easily turn to another firm when purchasing the good. Assuming lost sales at the retailers, we consider a one warehouse several retailers inventory system. Using the well-known METRIC-approximation as a framework, we present a heuristic for finding cost effective base-stock policies. In a numerical study we find that the cost of the policies suggested by the heuristic is on average 0.40% above the cost of the $(S-1, S)$-optimal policy.

Keywords : Inventory, Multi-echelon, Lost sales, METRIC

1 Introduction

Consider a two-echelon inventory system with one central warehouse and an arbitrary number of retailers. See Figure 1. The retailers face customer demand and replenish their stocks from the central warehouse. The warehouse, in turn, replenishes its stock from an outside supplier. Evaluation and optimization of control policies for such inventory systems have attracted massive interest in the literature. See, for example, Axsäter [3] for an overview. In the existing literature dealing with multi-echelon inventory control the prevalent assumption is that complete backlogging of orders is allowed in case of stockouts. For example, Axsäter [4] shows how to exactly evaluate the performance for different (R, nQ)-polices when the retailers face compound Poisson demand and inventories

*Corresponding author. Tel.:+45-89 42 35 36; fax: +45-86 13 17 69; e-mail: philip@imf.au.dk

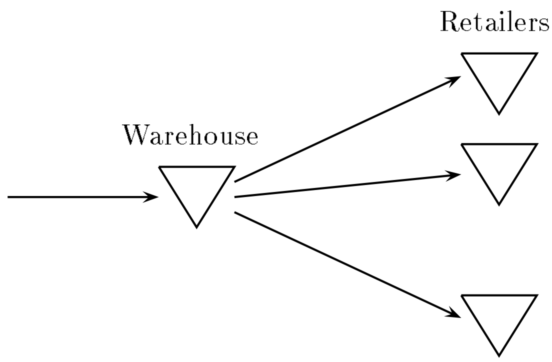

Figure 1: Multi-echelon inventory system

are continuously reviewed. Cachon [5] gives an exact method for the periodic review case with identical retailers.

In some situations the assumption of complete backlogging may not be so realistic. For example, it may be more representative to model stockouts as lost sales when the retailers are in a competitive market and customers can easily turn to another firm when purchasing the good. For some reason the research dealing with multi-echelon inventory models has focused mainly on the backorder case and the number of models dealing with lost sales is rather limited. Anupindi and Bassok [1] consider a periodic review two-echelon inventory system where a part of the unsatisfied sales at the retailers are lost. Since the transportation time between the manufacturer and the retailers is zero, the optimal order policy at each retailer is a base-stock policy. The manufacturer carries linear production cost and no holding cost. The retailers can agree to centralize their stocks and the problem considered is whether or not this will lead to an increase in total expected sales at the manufacturer. Nahmias and Smith [8] also consider lost sales in a multi-echelon environment in a paper more closely related to this paper. However, their model differs from ours in several important aspects. First, they consider periodic review batch order policies. The model is more general since they deal with partial lost sales. This means that, with probability u, demand not satisfied immediately, is lost, and with probability $1-u$, it is satisfied later by a special order. Moreover, for the model to be tractable they assume instantaneous deliveries from the warehouse to the retailers.

For single-echelon inventory models the lost sales assumption is more common. The exact cost for a single level inventory system facing Poisson demand and fixed leadtimes was first given by Hadley and Whitin [6]. Smith [11] demonstrates how to evaluate and find optimal $(S-1,S)$-policies for an inventory system with zero replenishment costs and generally distributed stochastic leadtimes. Recently Hill [7] showed that for the lost sales case the $(S-1,S)$-policy is not necessarily optimal.

In this paper we analyze a model for a one warehouse, multiple retailer inventory system. Demand occurs only at the retailers and follows independent Poisson processes. All leadtimes are assumed to be constant. All installations use $(S-1, S)$-policies with continuous review. It is assumed that backlogging of customer demand is not allowed. The analysis departs in one of the most widely known multi-echelon inventory models, the METRIC-model developed by Sherbrooke [10]. In its original setting, it is assumed that stockouts at the retailers are completely backlogged. We demonstrate how the METRIC-model can be modified to handle the lost sales case. Our approach gives an approximate model which is quite simple and efficient from a computational point of view. Simulation experiments indicate that the performance is very good.

The outline of this paper is as follows: In Section 2 we give a detailed problem formulation and pose all assumptions. Section 3 gives the solution procedure. The numerical results are given in Section 4, and in Section 5 we give some conclusions and point out some possible directions for future research.

2 Problem Formulation

The inventory system under consideration consists of one central warehouse and an arbitrary number of retailers. The retailers face Poisson customer demand. No backlogging is allowed at the retailers. Consequently, the customers that arrive to a retailer that is out of stock will become lost sales for the retailer. When stockouts occur at the warehouse, all demands from the retailers are fully backlogged and the backorders are filled according to a FIFO-policy. The transportation time between the warehouse and a given retailer is assumed to be constant as well as the transportation time from the external supplier to the warehouse. The cost of a replenishment is assumed to be zero or negligible compared to the holding and stockout costs. The external supplier is assumed to have infinite capacity, which means that the replenishment leadtime for the central warehouse is constant. All installations use $(S-1, S)$-policies with continuous review. Units held in stock both at the warehouse and at the retailers incur holding costs per unit and time unit. Moreover, a fixed penalty cost per lost customer is incurred at the retailers. In this paper we present a model for the considered inventory system, which can be used to evaluate the long-run average cost for different policies within the class of $(S-1, S)$-policies. The objective is to find the policy that minimizes the long-run average cost for the inventory system. Let us introduce the following notation:

$N =$ the number of retailers,

$\lambda_i =$ demand intensity at retailer i, $i = 1, 2, \ldots, N$,

$L_i =$ transportation time for the deliveries from the warehouse to retailer i, $i = 1, 2, \ldots, N$,

$L_0 =$transportation time for the deliveries from the external supplier to the warehouse,

S_0 = order-up-to level at the warehouse,
S_i = order-up-to level at retailer i, $i = 1, 2, \dots, N$,
h_0 = holding cost rate at the warehouse,
h_i = holding cost rate at retailer i, $i = 1, 2, \dots, N$,
π_i = penalty cost for a lost sale at retailer i, $i = 1, 2, \dots, N$.

We want to determine the total cost for the inventory system in steady state. Define
TC = total cost for the inventory system per time unit in steady state,
C_0 = cost per time unit for the warehouse in steady state,
C_i = cost per time unit for retailer i in steady state, $i = 1, 2, \dots, N$.

Obviously,

$$TC = C_0 + \sum_{i=1}^{N} C_i. \tag{1}$$

Our objective is to determine a control policy, $S_0, S_1, .., S_N$ that minimizes the total cost, TC.

3 Solution Procedure

In this section we first demonstrate how the total cost for different control policies can be evaluated. For the backorder case the exact cost of the system can be derived by observing that any unit ordered by retailer i is used to fulfill the S_ith demand. The cost can then be derived by conditioning on the arrival time of the S_ith demand (which is Erlang distributed) and the arrival of the ordered unit (see Axsäter [3]). In a lost sales environment the corresponding observation is that any unit ordered by retailer i is used to fulfill the $S_i + X_i$th demand, where X_i is a random variable denoting the number of lost sales incurred at the retailer during the replenishment lead time. X_i is obviously very hard to characterize and we have therefore chosen to focus on a heuristic rather than on the exact solution.

The analysis has many similarities with the analysis in Sherbrooke [10]. However, our assumption of lost sales at the retailers destroys some of the nice properties valid for the backorder model. The analysis of the warehouse, e.g., becomes more complex for the lost sales case. In the backorder case, all customers arriving at the retailers generate demands at the warehouse immediately at the arrival epoch, since all retailers use continuous review $(S-1, S)$-policies. Consequently, the warehouse faces a Poisson process with intensity $\lambda_0 = \lambda_1 + \lambda_2 + \cdots + \lambda_N$. For the lost sales case this is not true. When backordering is not allowed, customer demands can be lost due to stockouts at the retailers. Therefore the demand at the warehouse is not Poisson process anymore.

Another important difference compared with the backorder case is that the order-up-to level S_i at retailer i, affects the costs at all retailers and at the warehouse. In the backorder case S_i only affects the local cost at retailer i, since the warehouse demand process is unaffected by the order-up-to levels at the retailers. For the lost sales case the order-up-to level affects the number of lost sales and consequently, the demand process at the warehouse is not independent of the policies at the retailers. Therefore the order-up-to level at a certain retailer affects the costs at all installations in the inventory system.

We will first show how to evaluate the costs at the retailers given a certain replenishment leadtime provided by the warehouse. We then show how to calculate the cost at the warehouse given the demand intensity from the retailers. Finally we introduce an iterative procedure from which we obtain the total cost for the inventory system.

3.1 Approximate retailer cost

As Sherbrooke [10] we use a queueing system analogy when evaluating the costs for the retailers. For a retailer where backlogging is allowed, the number of outstanding orders towards the central warehouse is the same as the occupancy level in an $M/G/\infty$ queue. Recall that the customer demand is Poisson and the replenishment leadtimes are stochastic, since orders can be delayed due to stockouts at the central warehouse. For this type of queue a famous theorem by Palm [9] states that the steady state occupancy level is Poisson distributed with mean λL, where λ is the arrival rate and L is the mean service time. Palm's theorem holds for i.i.d. service times. The stochastic leadtimes in our case are evidently not independent, but if we disregard this correlation we can approximate the number of outstanding orders with a Poisson distribution. This is the idea behind the METRIC-approximation.

When demand is lost, the queueing system of interest is an $M/G/S/S$ queue, with S servers, each with generally distributed service times and no queueing allowed. If the service times are independent random variables with mean $\bar{L}$, Erlang's loss formula states the steady-state distribution for the occupancy level as

$$q^S(j) = \frac{(\lambda \bar{L})^j / j!}{\sum_{n=0}^{S} (\lambda \bar{L})^n / n!} \text{ for } 0 \leq j \leq S$$

where $q^S(j) =$ the probability that j servers (out of S) are occupied in steady state. Following METRIC we approximate the number of outstanding orders with this distribution.

Suppose that the mean leadtime for retailer i is $\bar{L}_i$ and let $q_i^{S_i}(j)$ be the steady state probability of j outstanding orders given a desired base-stock level S_i. The expected number of lost sales per time unit is clearly $\lambda_i q_i^{S_i}(S_i)$ and the expected number of units in stock is

$$\sum_{j=0}^{S_i}(S_i - j)q_i^{S_i}(j) = S_i - [1 - q_i^{S_i}(S_i)]\lambda_i \bar{L}_i. \tag{2}$$

The total relevant cost for retailer i is therefore

$$C_i(S_i, \bar{L}_i) = \lambda_i \pi_i q_i^{S_i}(S_i) + h_i\Big(S_i - [1 - q_i^{S_i}(S_i)]\lambda_i \bar{L}_i\Big)$$

and the rate of demand from retailer i which is not lost is $(1 - q_i^{S_i}(S_i))\lambda_i$.

The derivation of the exact cost of a $(S-1, S)$ lost sales single stage inventory system with generally distributed leadtimes was first presented by Smith [11]. He also proves that $C_i(S_i, \bar{L}_i)$ is convex in S_i for fixed $\bar{L}_i$, which means that the optimal value can be found by a local search routine.

3.2 Approximate warehouse cost

In the backorder case the demand process at the warehouse is a Poisson process. In the lost sales case this is not the case. If, for example, the base-stock level at a retailer is one, the smallest interval between two successive demands from that retailer will be the retailer leadtime. We will ignore this and approximate the demand process at the warehouse with a Poisson process with mean Λ. Λ depends on how much demand is lost at the retailers and is determined as

$$\Lambda = \sum_{i=1}^{N} \lambda_i(1 - q_i^{S_i}(S_i)) \tag{3}$$

Since we have a fixed deterministic leadtime L_0, we can find the average holding cost incurred at the warehouse as a function of Λ and S_0.

$$C_0(S_0, \Lambda) = h_0 \sum_{j=0}^{S_0}(S_0 - j)\frac{(\Lambda L_0)^j}{j!}\exp(-\Lambda L_0)$$

We can also derive the mean delay due to stockouts at the warehouse by first calculating B_0, the average number of backorders at the warehouse.

$$B_0 = \sum_{j=S_0+1}^{\infty}(j - S_0)\frac{(\Lambda L_0)^j}{j!}\exp(-\Lambda L_0), \tag{4}$$

We then apply Little's formula to obtain the average delivery delay, B_0/Λ. The mean leadtime for retailer i is then

$$\bar{L}_i = L_i + B_0/\Lambda \tag{5}$$

Finally, we obtain the total cost from (1).

3.3 Overall solution procedure

We can now establish the solution procedure. The procedure enumerates over S_0. It can be shown that for a cost minimizing solution, S_0 can not be negative. See, for example, Axsäter [2]. Consequently our procedure starts with S_0=0. Moreover, S_0 is bounded from above by an abortion criteria. We need the following new notation:

$C_i^{\min}$	=	$\min_{S_i} C_i(S_i, L_i)$ = minimum cost per time unit for retailer i in steady state when the leadtime, $\bar{L}_i$, is equal to the transportation time, L_i, $i = 1, 2, \ldots, N$.
$S_i(k)$	=	order-up-to level at retailer i in iteration k.
$TC^*(S_0)$	=	minimum value of TC given a fixed value of S_0.

Let us first consider two simple lemmas. The proofs can be found in the Appendix. The first lemma gives a lower bound for the retailer costs, and the second establishes two important properties for the warehouse cost.

Lemma 1. *$C_i^{\min}$ is a lower bound for the retailer cost, $C_i(S_i, \bar{L}_i)$ for all S_i and any $\bar{L}_i > L_i$.*

Lemma 2. *$C_0(S_0, \lambda_0) \leq C_0(S_0, \Lambda)$, for all S_0 and all $\Lambda \leq \lambda_0$. Moreover, $C_0(S_0, \lambda_0)$ is convex in S_0.*

To construct an abortion criteria for the procedure, consider the cost function

$$TC_{lb}(S_0) = C_0(S_0, \lambda_0) + \sum_{i=0}^{N} C_i^{\min}.$$

By Lemma 1 and Lemma 2, $TC_{lb}(S_0)$ is a lower bound for the cost function, $TC^*(S_0)$. Moreover, since the cost function $TC_{lb}(S_0)$ is convex in S_0 the search over S_0 can be aborted when S_0 satisfies

$$\min_{x \leq S_0} TC^*(x) \leq TC_{lb}(S_0).$$

The abortion criteria is illustrated in Figure 2.

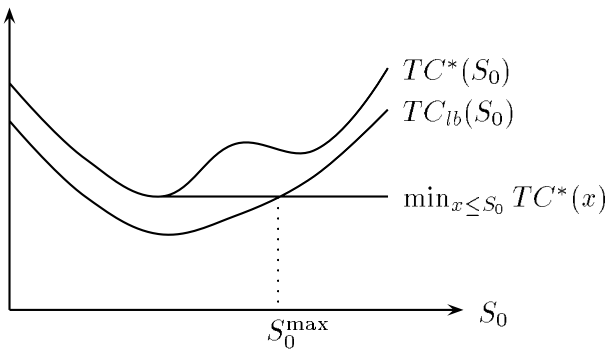

Figure 2: Illustration of the abortion criteria. The search for the optimal S_0 is aborted at $S_0^{\max}$.

The solution procedure can now be established as:

STEP 0: Set $S_0 = 0$ and $TC_{\min} = \infty$.

STEP 1: Set $k = 0$ and $\Lambda = \lambda_0$.

STEP 2: For each $i = 1, 2, \ldots, N$ calculate $\bar{L}_i$ by (4) and (5)
Let $z^* = \min_S C_i(S, \bar{L}_i)$ and set
$S_i(k) = \min\{S | C_i(S, \bar{L}_i) = z^*\}$

STEP 3: If $k > 0$ and $S_i(k) = S_i(k-1)$ for all $i = 1, 2, \ldots, N$ then goto STEP 4, else calculate Λ by (3), set $k := k + 1$ and goto STEP 2.

STEP 4: Set $TC^*(S_0) = C_0(S_0, \Lambda) + \sum_{i=1}^{n} C_i(S_i(k), \bar{L}_i)$.
If $TC^*(S_0) < TC_{\min}$ then set $TC_{\min} = TC^*(S_0)$
and let $S_0^{opt} = S_0$ and $S_i^{opt} = S_i(k)$ for $i = 1, 2, \ldots, N$.
If $TC_{\min} < TC_{lb}(S_0)$ then STOP, else set $S_0 = S_0 + 1$ and goto STEP 1.

4 Numerical Results

In order to examine the effectiveness of the presented methodology we have performed a small numerical study. In total we consider 36 different test problems with five identical retailers. For each test problem we find the best order-up-to levels according to our method. We also obtain the approximate total holding and stock out costs for the inventory system. The accuracy of these results are then evaluated by simulation. Each simulation consists of 10 runs, each with a run length of 100 000 time units. The result is a confidence interval for the exact cost. We express the confidence limits on a 95% significance level. A comparison between the total cost given by our method and the total cost for the simulation gives an indication of how accurate our method is when estimating the total cost for the inventory system.

We also use simulation to determine the optimal policy for the system. The cost for this policy can then be compared with the cost for the policy determined by our technique, to obtain an estimate for the performance of the method when optimizing the ordering policies. The policy that we report as the optimal policy is the policy with the lowest average cost. However, this policy does not necessarily dominate all the other policies when taking confidence intervals into consideration. Moreover, we only search within policies where the order-up-to levels are identical for the retailers.

The problem data and results can be found in Table 1. We only report the optimal policy when it is different from the one obtained from our algorithm. From Table 1 we can see that our method performs rather well for all the considered problems. It seems that we mostly tend to underestimate the total cost, especially in the problems with high stockout costs at the retailers. This is due to the METRIC-approximation, where the stochastic leadtimes are replaced by their averages when evaluating the costs for the retailers. On average the method underestimates the costs with 1.1 %.

In 13 problems we can observe (on a significance level of 95%) that the method fails to find the optimal policy. In 9 more problems the policy suggested by our method does not have the lowest average cost according to the simulation runs. However, in these cases the deviations are not significant on a 95% confidence level. In comparison to the optimal policies obtained by simulation, the increase in costs by using the policies obtained by our method is only 0.40 %, on average. In 16 of the 22 problems where we fail to find the true optimal policy, the method merely underestimates the order-up-to level at the warehouse by a single unit. In one problem the warehouse order-up-to level is underestimated by two units. In the other 5 problems where the optimal policy is not found, the method tends to allocate more stock to the retailers and less stock to the warehouse than what is optimal from a cost perspective.

Finally it seems that our methodology performs better if the warehouse leadtime is small compared to the transportation time from the warehouse to the retailers. In the 12 problem instances with $L_i = 1.5$ the average cost increase, SC/CC is only 0.07%, whereas in the problems with $L_i = 0.5$, the corresponding figure is considerable higher, 0.67%. This behavior is due to the METRIC-approximation, where the stochastic replenishment leadtime facing a retailer is replaced by its mean value. If the constant transportation time to the retailers is large compared to the warehouse leadtime, the stochastic delivery delays tend to have less relative variation and consequently the impact of the METRIC-approximation will be smaller.

5 Conclusions and directions for future research

This paper presents a heuristic method for evaluation and optimization of $(S-1, S)$-policies for a one warehouse, multiple retailers inventory system. The evaluation technique uses the well-known METRIC-approximation as a framework. From a computational point of view the presented technique is very efficient and simple. Numerical results also indicate that the performance is quite good.

Up to our knowledge, no paper is yet published, which deals with lost sales in a continuous review multi-echelon inventory setting. Moreover, the original backorder METRIC-model [10] is one of the most widely used multi-echelon inventory models. Our lost sales generalization makes the policy evaluation a bit more complex, since we have to use an iterative procedure to obtain the cost. Still, the model is rather simple and easy to implement. Moreover, in many practical situations lost sales is a reasonable way to model stockouts. Therefore our technique is also relevant for practitioners.

In a research perspective our model can form a framework in which different generalizations can be considered as options for future research. For example, batch ordering policies and more general demand processes may be analyzed, still using the ideas presented in this paper. Generalizations to periodic review policies is also important. The derivation of an exact evaluation of costs seems to be a very difficult problem to solve. This is a real challenge for future research.

Appendix

Proof for Lemma 1

We need to show that

$$\min_{S_i} \; C_i(S_i, L_i) \leq \min_{S_i} \; C_i(S_i, \bar{L}_i) \text{ for } L_i \leq \bar{L}_i. \tag{6}$$

Let l_i be an arbitrarily chosen leadtime, where $L_i \leq l_i \leq \bar{L}_i$. Consider the cost $C_i(S_i, l_i)$, where S_i is set to its optimal value for each l_i. Obviously, $C_i(S_i, l_i) \leq C_i(S_i - 1, l_i)$ for each l_i such that $L_i \leq l_i \leq \bar{L}_i$. Start with $l_i = \bar{L}_i$ and let l_i be continuously lowered until we reach $l_i{=}L_i$, while S_i is set to its optimal value for each l_i. Since $C_i(S_i, l_i)$ is a continuous function of l_i for fixed S_i, it also is a continuous function of l_i when S_i is optimally chosen. Moreover, the fact that S_i minimizes the cost $C_i(S_i, l_i)$, implies that $C_i(S_i, l_i) \leq C_i(S_i - 1, l_i)$. Consequently, (6) follows if

$$C_i(S_i, l_i) \leq C_i(S_i - 1, l_i) \Rightarrow \frac{\partial C_i(S_i, l_i)}{\partial l_i} \geq 0 \text{ for } L_i \leq l_i \leq \bar{L}_i. \tag{7}$$

For notational reasons we omit the index i from all variables. It can be shown that

$$\frac{\partial C(S, l)}{\partial l} = -h\lambda(1 - q^S(S)) + \lambda(h\lambda l + \pi\lambda)(q^S(S-1) - q^S(S) + q^S(S)^2). \tag{8}$$

Moreover, $C(S, l) \leq C(S - 1, l)$ implies that

$$\frac{h}{h\lambda l + \pi\lambda} \leq q^{S-1}(S-1) - q^S(S) \tag{9}$$

Let

$$A = \frac{1}{\lambda(h\lambda l + \pi\lambda)} \frac{\partial C_i(S,l)}{\partial l}. \qquad (10)$$

From (8) we have that

$$A = \frac{h}{h\lambda l + \pi\lambda} \big(q^S(S) - 1\big) + q^S(S-1) - q^S(S) + q^S(S)^2). \qquad (11)$$

(9) and (11) now give

$$\begin{aligned} A &\geq (q^{S-1}(S-1) - q^S(S))(q^S(S) - 1) + q^S(S-1) - q^S(S) + \big(q^S(S)\big)^2 \\ &= q^{S-1}(S-1)q^S(S) - q^{S-1}(S-1) + q^S(S-1) \\ &= q^{S-1}(S-1)q^S(S) - (q^{S-1}(S-1)q^S(S)) \\ &= 0. \end{aligned}$$

Note that $q^S(S) \leq 1$. Consequently, $A \geq 0$ and therefore (7) holds and the proof is complete.

Proof for Lemma 2

Since $\Lambda \leq \lambda_0$, we only need to show that $\frac{\partial C_0(S_0,\Lambda)}{\partial \Lambda} \leq 0$. The convexity of $C_0(S_0, \lambda_0)$ in S_0 follows, for example, from Axsäter [3].

$$\begin{aligned} \frac{\partial C_0(S_0,\Lambda)}{\partial \Lambda} &= -h_0 L_0 \exp(-\Lambda L_0) \left(S_0 + \sum_{j=1}^{S_0} (S_0 - j) \cdot \left(\frac{(\Lambda L_0)^j}{j!} - \frac{(\Lambda L_0)^{(j-1)}}{(j-1)!} \right) \right) \\ &= -h_0 L_0 \exp(-\Lambda L_0) \sum_{j=0}^{S_0-1} \frac{(\Lambda L_0)^j}{j!} \\ &\leq 0. \end{aligned}$$

References

[1] R. Anupindi and Y. Bassok. Centralization of stocks: Retailers vs. manufacturer. *Management Science*, 45:178–191, 1999.

[2] S. Axsäter. Simple solution procedures for a class of two–echelon inventory problems. *Operations Research*, 38:64–69, 1990.

[3] S. Axsäter. Continuous review policies for multi–level inventory systems with stochastic demand. In S.C. Graves, A.H.G. Rinnooy Kan, and P. Zipkin, editors, *Handbooks in OR & MS, Vol. 4*, pages 175–197. Elsevier Science Publishers B.V., North–Holland, 1993.

[4] S Axsäter. Exact analysis of continuous review (R,Q)–policies in two–echelon inventory systems with compound Poisson demand. *Lund University, Sweden*, 1997.

[5] G.P. Cachon. Exact evaluation of batch-ordering inventory policies in two–echelon supply chains with periodic review. *Fuqua School of Business, Duke University*, 1995.

[6] G. Hadley and T.M. Whitin. *Analysis of Inventory Systems*. Prentice–Hall, Englewood Cliffs, NJ, 1963.

[7] R.M. Hill. On the suboptimality of $(S-1,S)$ lost sales inventory policies. *International Journal of Production Economics*, 59(1-3):377–385, 1999.

[8] S. Nahmias and S.A. Smith. Optimizing inventory levels in a two-echelon retailer system with partial lost sales. *Management Science*, 40:582–596, 1994.

[9] C. Palm. Analysis of the Erlang traffic formula for busy signal assignment. *Ericson Technics*, 5:39–58, 1938.

[10] C.C. Sherbrooke. METRIC: A multi-echelon technique for recoverable item control. *Operations Research*, 16:122–141, 1968.

[11] S.A. Smith. Optimal inventories for an $(S-1,S)$ system with no backorders. *Management Science*, 23:522–528, 1977.

	λ_i	π_i	L_i	S	Calc cost	Sim cost	Spread	Opt pol	Cost	Spread	Cost dev
1	1	5	0.5	4,2	10.82	10.74	0.01				
2	1	5	1.0	2,3	11.99	12.03	0.02				
3	1	5	1.5	3,3	12.51	12.46	0.01				
4	1	25	0.5	5,3	15.86	16.15	0.03				
5	1	25	1.0	4,4	18.18	18.43	0.04	5,4	18.41	0.03	0.1%
6	1	25	1.5	3,5	19.94	20.16	0.06	4,5	20.12	0.04	0.2%
7	1	125	0.5	5,4	20.27	21.27	0.05	6,4	20.90	0.06	1.7%
8	1	125	1.0	5,5	23.64	24.16	0.06	6,5	24.09	0.06	0.3%
9	1	125	1.5	5,6	26.15	26.45	0.08				
10	2	5	0.5	8,3	15.43	15.41	0.02	9,3	15.33	0.02	0.5%
11	2	5	1.0	6,5	17.11	17.27	0.02	8,4	17.16	0.02	0.6%
12	2	5	1.5	5,6	18.23	18.34	0.02	7,5	18.29	0.03	0.3%
13	2	25	0.5	8,5	21.56	22.52	0.04	9,5	22.30	0.04	1.0%
14	2	25	1.0	9,6	24.96	25.35	0.06	10,6	25.33	0.05	0.1%
15	2	25	1.5	7,8	27.49	27.92	0.08	8,9	27.90	0.04	0.1%
16	2	125	0.5	9,6	26.82	28.56	0.12	10,6	28.15	0.07	1.5%
17	2	125	1.0	9,8	31.84	32.75	0.08				
18	2	125	1.5	10,9	35.49	36.07	0.11				
19	1	5	0.5	4,1	16.96	16.41	0.02				
20	1	5	1.0	2,2	17.55	17.51	0.02				
21	1	5	1.5	2,2	18.14	18.04	0.03				
22	1	25	0.5	4,3	27.50	27.99	0.03	5,3	27.96	0.02	0.1%
23	1	25	1.0	5,3	30.81	30.72	0.06	6,3	30.65	0.05	0.2%
24	1	25	1.5	4,4	33.07	33.11	0.05				
25	1	125	0.5	7,3	36.86	37.21	0.12	8,3	36.92	0.08	0.8%
26	1	125	1.0	4,5	42.75	43.64	0.14	7,4	43.19	0.17	1.1%
27	1	125	1.5	6,5	46.83	46.92	0.10				
28	2	5	0.5	6,3	24.21	24.42	0.02	8,2	24.14	0.03	1.2%
29	2	5	1.0	5,4	26.48	26.60	0.02	7,3	26.18	0.03	1.6%
30	2	5	1.5	6,4	27.61	27.43	0.02				
31	2	25	0.5	10,4	36.44	37.10	0.06	11,4	36.95	0.05	0.4%
32	2	25	1.0	10,5	42.85	42.91	0.12	11,5	42.85	0.07	0.1%
33	2	25	1.5	8,7	46.74	47.11	0.08				
34	2	125	0.5	11,5	47.24	48.92	0.08	13,5	48.21	0.12	1.5%
35	2	125	1.0	10,7	56.53	57.73	0.16	11,7	57.37	0.11	0.6%
36	2	125	1.5	11,8	63.77	64.19	0.19	12.8	64.05	0.10	0.2%

Table 1: Numerical results. The optimal policy is only reported when it is different than the policy suggested by our algorithm.