Int. J. Production Economics 143 (2013) 379–384

Contents lists available at SciVerse ScienceDirect

Int. J. Production Economics

journal homepage: www.elsevier.com/locate/ijpe

Modified base-stock policies for continuous-review, lost-sales inventory models with Poisson demand and a fixed lead time

Søren Glud Johansen *

Department of Mathematics, Aarhus University, Ny Munkegade 118, DK-8000 Aarhus C, Denmark

ARTICLE INFO

Article history:
Received 14 September 2010
Accepted 21 February 2012
Available online 28 February 2012

Keywords:
Base-stock policy
Lost sales
Continuous review
Simulation
Optimization

ABSTRACT

This paper reconsiders the lost-sales inventory system studied by Hill (2007). The commonly assumed policy to apply to the system is a pure base-stock policy (PBSP) for which the best base stock is easily found. Hill shows that his simple delay policy (SDP) and full delay policy (FDP) perform better. The SDP is a (s,d) policy where s is the base stock of the best PBSP and d is a common lower bound on the delay between the placement of successive replenishment orders. We show by simulation that the d value suggested by Johansen (2001) outperforms Hill's suggestion and that the performance often can be further improved by optimizing d. For the test bed investigated by Hill, we show that, for some parameter settings, an additional improvement is achieved when s and d are optimized simultaneously. The policy suggested by Johansen performs better than the FDP in all settings where the former policy reduces the average cost of the best PBSP by at least 1%.

1. Introduction

We reconsider the lost-sales inventory system studied by Hill (2007). The system has Poisson demand with rate λ, continuous review and a fixed lead time L. There is a holding cost h per unit per unit time and a penalty p per unit lost. The objective is to minimize the long-run average cost per unit time subject to the condition that all replenishments are unit sized. This condition is met without loss of optimality when, as assumed by Hill, there is no fixed order cost because then economies of scale are lacking. However, unless λ is relatively big, a good replenishment policy satisfying the condition for a positive fixed order cost can be found if p is computed as the difference between the lost sales cost per unit and the fixed order cost.

The considered system can describe slow-moving but important and possibly expensive spare parts for which the replenishment lead time is relatively long. When demand for such parts occurs during a stockout, the demand is lost to the regular control system because it is satisfied (at an extra cost) by some other means. For the retail sector, the system can describe high-value goods for which a customer demand is lost if the item is not in stock. Both applications often have Poisson demand with a rate λ which is not relatively big. Then a continuous review model provides a reasonable representation of the system.

The commonly assumed policy to apply to the system is a pure base-stock policy (PBSP). It prescribes to maintain the inventory position (the sum of the stock in hand and the stock on order) at some base-stock level s. Hence, if the initial inventory position equals the chosen s, then a new replenishment order for one unit is placed immediately whenever a demand is satisfied. As explained in Section 2 it is straightforward to find the best base stock s_{PBSP} for the PBSP. However, Hill (1999) has shown that a PBSP can never be optimal if $s_{PBSP} > 1$, which applies for most realistic parameter settings. Johansen (2001) and Hill (2007) offer better solutions by suggesting modified base-stock policies (MBSPs) which impose some minimum delay between the placement of successive replenishment orders. The intuitive reason for such a minimum is that, if we shortly after placing an order place another one, then we are very likely to be in stock when the second order arrives and therefore we would be increasing stock with a very small likelihood of that unit being immediately needed.

In this paper we investigate by simulation how the considered system performs when it is controlled by different MBSPs specified by a pair (s,d), where s is the base stock and the lower delay bound is fixed as d. The investigated policies are related to (and some of them improve) the MBSPs suggested by Johansen and Hill.

The paper is organized as follows. Section 2 provides a brief review of related literature. Our simulation models of the MBSPs are presented in Section 3. Numerical results obtained by the simulation models are reported and discussed in Section 4 and Section 5 contains our conclusion.

* Tel.: +45 871 55806; fax: +45 861 31769.
E-mail address: sgj@imf.au.dk

0925-5273/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.ijpe.2012.02.021

2. Related literature

Bijvank and Vis (2011) provide a good review of lost-sales inventory theory. The authors present a classification scheme for the replenishment policies most often applied in literature and practice. They notice that their scheme does not include policies that are optimal in a lost-sales setting. For systems with continuous review, which can be implemented by using transactions reporting (Hadley and Whitin, 1963, p. 159), they conclude that there are hardly any comparisons with an optimal policy. The reason for this conclusion is that such systems are very difficult to analyze in general. The difficulties can be overcome for systems which have an optimal policy prescribing at most one order outstanding at any time. However, the inventory system considered by us only has this property if $s_{\text{PBSP}}=1$, which is a very restrictive condition.

For the PBSP with base stock s, the stock on hand equals the difference between s and the stock J on order. When the lead times are mutually independent with mean L (and, in particular, constant as assumed in this paper), it is well known (Zipkin, 2000, Section 7.2.3) that J is specified as the number of busy servers in an $M/G/s/s$ system with traffic intensity $\rho=\lambda L$ and that the long-run fraction of demand lost is given by Erlang's loss formula

$$\overline{B}(s)=\frac{\frac{\rho^s}{s!}}{\sum_{j=0}^{s}\frac{\rho^j}{j!}}.$$

The average stock on order is

$$\overline{J}(s)=\rho[1-\overline{B}(s)],$$

and the long-run average cost per unit time is

$$\overline{C}(s)=p\lambda\overline{B}(s)+h[s-\overline{J}(s)]=[p\lambda+h\rho]\overline{B}(s)+h[s-\rho].$$

It is straightforward to find the base stock s_{PBSP} which minimizes this cost function because it is convex in s. Convexity follows from the fact that $\overline{B}(s)$ is convex in s (Jagers and Van Doorn, 1986).

Hill (2007) explores three alternative policies which might offer better solutions than the PBSP. Two of his alternative policies are MBSPs, which he refers to as the simple delay policy (SDP) and the full delay policy (FDP), with base stock s_{PBSP}. His third policy is a fixed order policy (FOP). The FOP places orders at fixed and regular intervals of time τ. Hence, the stock on hand is modeled as a $D/M/1$ queue with arrivals occurring at fixed intervals of length τ and exponential service times having mean $1/\lambda$. Hill shows how to find the best τ for the FOP and he reports that there may be some situations in which the FOP with the best τ offers a realistic alternative to the PBSP. We do not investigate the FOP further in this paper because it is not a good policy for the parameter values investigated in our numerical study.

Hill's SDP is a $(s_{\text{PBSP}},d_{\text{H}})$ policy where

$$d_{\text{H}}=\frac{-\ln(1-h/(\lambda p))}{\lambda}. \tag{1}$$

His FDP requires that a suggested delay is computed each time a sale occurs or an order is placed. The procedure for computing each suggested delay of the FDP is a myopic heuristic. When a sale occurs or an order is placed, the suggested delay T is computed by this heuristic based on information about the actual state of the system. If the next sale occurs within the time interval of length T, then a new value of T is computed at the sale epoch. If not, the FDP places the next order after the suggested delay and then a new T value is computed. We agree with Hill in the conclusion that the FDP is "less likely to be operationally acceptable". Moreover, his SDP and FDP are dominated by better policies as we demonstrate in our numerical study.

The periodic review lost-sales model where L is an integral number of review periods and there is no fixed order cost, has been studied by many authors. Zipkin (2008) provides several earlier references and he investigates for short lead times (L is at most four review periods) how various policies specified by plausible heuristics perform relative to the optimal periodic-review policy which he computes by value iteration. He concludes, based on a numerical study, that several of the investigated policies perform reasonably well, but that the base-stock policy, with base stock computed as for a backlog system, performs poorly. However, the PBSP specified by the best base stock for the lost-sales system performs better. For a large class of demand distributions, Huh et al. (2009) have shown that the relative difference between the cost of the optimal policy and the best PBSP converges to zero when the lost-sales penalty p becomes large compared to the holding cost rate h.

Johansen (2001) presents a policy-iteration algorithm for finding the policy which is optimal for the periodic review model with Poisson demand subject to the condition that s, defined as the largest sum of the stock on hand and the number of units under delivery for the policy found by the algorithm, does not exceed some integer S. He suggests initially to apply the algorithm with S equal to an easily computed base stock $\bar{s}$, say, of a good PBSP and to repeat applying the algorithm for S incremented by one until the same policy is found by the algorithm for the last two S values. This stop rule ensures for the final S that the optimal value of s is $S-1$. For $L=10$ and various values of λ and the parameters h and p, he reports the long-run average cost per review period for the found optimal policy and three other policies. The first of these is the PBSP with base stock $\bar{s}$, whereas the others are MBSPs. The second policy is the MBSP with base stock $\bar{s}$ and lower delay bound d specified as $L/\bar{s}$ round off to an integer number of review periods (so that each replenishment order is placed at a review epoch). The third policy is the best MBSP. The reported results (and numerous others) illustrate that the second and third policies provide most of the cost reductions which can be obtained by replacing the best PBSP by the optimal policy. The second policy is easy to compute, whereas extensive computations are needed to find the third policy. Because the second policy is often a good choice, Johansen recommends to implement it. If the length of the review period is decreased and approaches zero, the periodic review model approaches the continuous review model and no round off is needed when the lower delay bound of the second policy is specified as

$$d_{\text{J}}=L/s_{\text{PBSP}}. \tag{2}$$

Therefore, Johansen suggests to control the continuous review system by the $(s_{\text{PBSP}},d_{\text{J}})$ policy.

3. Simulation models

3.1. Performance of a (s,d) policy

For an MBSP specified by (s,d), we redefine J as the sum of the number N of delayed units and the number of single-unit orders outstanding at the supplier. Hence, the stock on hand remains specified as $s-J$ and we can evaluate the performance of the policy based on information about how J and N evolve over time. A demand occurring when $J=s$ is lost. When $J<s$ at a demand epoch, the demand becomes a sale and the policy prescribes to increment N by one if either N is already positive or the time since the last placement of an order is shorter than d. If not, a replenishment order is placed immediately at the supplier. During periods where $N>0$, the policy prescribes to place an order at