Int. J. Production Economics 131 (2011) 69–75

Contents lists available at ScienceDirect

Int. J. Production Economics

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

Derivation of confidence intervals of service measures in a base-stock inventory control system with low-frequent demand

Christian Larsen

Centre for Operations Research Applications in Logistics, Department of Business Studies, Aarhus School of Business, Aarhus University, Fuglesangs Allé 4, DK-8210 Aarhus V, Denmark

ARTICLE INFO

Article history:
Received 1 September 2008
Accepted 12 August 2010
Available online 18 August 2010

Keywords:
Inventory control
Base-stock systems
Confidence intervals of service measures

ABSTRACT

We explore a base-stock system with backlogging where the demand process is a compound renewal process and the compound element is a delayed geometric distribution. For this setting it holds that the long-run average service measures *order fill rate* (*OFR*) and *volume fill rate* (*VFR*) are equal in values. However, though equal *ex ante* one will *ex post* observe differences as actual sample paths are different. By including a low-frequency assumption in the model, we are able to derive mathematical expressions of the confidence intervals one will get if *OFR* and *VFR* are estimated in a simulation using the regenerative method. Through numerical examples we show that of the two service measures it is *OFR* that in general can be estimated most accurately. However, simulation results show that the opposite conclusion holds if we instead consider finite-horizon service measures, namely per-cycle variants of *OFR* and *VFR*.

1. Introduction

It is a well recognized approach in both academia and practice to apply service measures to evaluate the performance of an inventory system. It serves as a substitute for using costs associated with backordering or lost sales. The reason is that these shortage costs in general can be hard to estimate: This holds in particular for elements like loss of goodwill due to not being able to serve a customer immediately. For a thought-provoking account on this; see Gardner (1980). Many various service measures to be applied in inventory management have been proposed, for reviews on the subject; see Schneider (1981) and Tempelmeier (1972). Also, Silver et al. (1998, Section 7.5) provides a good overview and discussion. One of the most commonly used service measures is the fill rate (*FR*), defined as the fraction of demand that can be met immediately from inventory without shortages (Silver et al., 1998, p. 245). Another service measure which focuses on complete fulfillment of an order is the order fill rate (*OFR*). It has been studied in a multi-item context by Song (1998), Hausman et al. (1998) and Hill and Pakkala (2007). Tempelmeier (1972) also addresses this service measure in a periodic review setting, where it is denoted α_{Per} service. As the *OFR* focuses on order fulfillment in contrast to the standard fill rate which focuses on the aggregate demand, we use in the following the specification volume fill rate (*VFR*) instead of *FR*, in order to make the distinction between the two fill rate measures more apparent. Though service measures are good practical tools for assessing the performance of an inventory system a possible drawback is that most service measures, like *OFR* and *VFR*, are expressed as long-run average values. We will in this paper address both characteristics 'long-run' and 'average'.

The latter refers to that service measures state a predicted average value but do not give any information about how much the deviation (from this predicted average value) actually will be. This means that when setting control parameters *ex ante*, based on expected service levels, and then estimating performance *ex post*, say after a year, one might very well observe some deviations between estimated service levels and predicted service levels. It would be of interest to classify various service measures according to how much actual realizations might differ from predicted values. This calls for not solely focusing on predicted values but also on deriving measures of variance of the service measures. This is the principal goal of the paper. We study this issue by comparing *OFR* and *VFR*. As we also believe that it makes most sense to make such a comparison for a case where both measures predict the same value of service. Therefore we explore the issue in a base-stock system with backlogging where the demand process is a compound renewal process and the compound element is a delayed geometric distribution. Because for this setting it is proven in Larsen and Thorstenson (2007) that the long-run average service measures *order fill rate* (*OFR*) and *volume fill rate* (*VFR*) are equal. Also in this reference it is shown in a simulation study (see Table 5 therein) that this result only holds when considering their mean values as actual sample paths are different. By comparing the reported confidence intervals in that table, it can be seen that *OFR* has the narrowest interval. Therefore the table indicates that it seems easier to achieve the

E-mail address: chl@asb.dk

0925-5273/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.ijpe.2010.08.007

pre-specified service level target by using the *OFR* service measure. As can be seen later this preliminary observation gets confirmed in this paper. It is difficult mathematically to derive higher order moments of service measures. However, by including a low-frequency assumption into the base-stock model, we are able to derive a measure of the variance of the two service measures. Specifically, we derive mathematically the confidence interval of *OFR* and *VFR* if these are estimated in a simulation by the regenerative method. The assumption about low-frequency is often coined *intermittent* in the literature see the seminal paper by Croston (1972). This intermittent behaviour of the demand (or lumpiness) is often encountered for inventories of spare parts where base-stock policies in general are recommended; see for instance Feeney and Sherbrooke (1966), Ward (1978), Schultz (1987) and Smith and Dekker (1997). Also, base-stock policies are often used in the retail sector, Hill and Pakkala (2007).

The other characteristic mentioned above: 'long-run' refers to that in principle the performances of the inventory concerns an infinite time horizon. The study of finite-horizon service measures has not received much attention in the literature. There are some analyses of the finite-horizon volume fill rate service measure in Chen et al. (2003) and Thomas (2005) however for very special cases. Also the minimal service level constraints studied in Chen and Krass (2001) can be interpreted as a finite-horizon volume fill rate service measure. In the paper we also explore what happens if we instead of long-run average service measures look at per-cycle service measures (denoted respectively OFR_{Cycle} and VFR_{Cycle}); these can be considered a sort of finite-horizon service measure. As these measures are harder to analyze mathematically we use here simulation. Our simulation experiments show that for a given base-stock parameter, the per-cycle service levels are higher than their long-run counterparts. This result is in accordance with Chen et al. (2003), which prove their result under more restrictive assumptions. Furthermore, of the two per-cycle service measures, it is VFR_{Cycle} that can be estimated most accurately.

In Section 2 we first give mathematical derivations of the service measures *OFR* and *VFR*. These derivations explicitly exploit the low-frequency assumptions and therefore appear as alternative derivations compared to Larsen and Thorstenson (2007). Then in Section 3 we mathematically derive the confidence intervals of our service measures. In Section 4 we present numerical examples from which it can be concluded that it in general it is *OFR* that can be estimated most accurately compared to *VFR* and that it is VFR_{Cycle} that can be estimated most accurately compared to OFR_{Cycle}. Finally, Section 5 contains some concluding remarks.

2. Derivation of the (long-run) average service measures

We consider a base-stock inventory control system with parameter S which is a positive integer. All replenishment orders are issued instantaneously at the time point of an order request and all replenishments have a constant lead time L. All unfilled demand is backlogged and serviced as soon as possible. The positive integer valued random variable X denotes the size of a customer order. We assume successive customer orders are independent and generated by X. The inter-arrival times between order requests are also assumed to be independent and to be generations of a positive continuous random variable T. We also assume that T and X are independent. We will assume that the demand process is *low-frequent*. Specifically, we assume that

$$P(T > L/2) = 1. \tag{1}$$

This means that at most one new customer order will appear during the replenishment lead time of any order. As an example consider that $L=5$ and T is uniformly distributed between 4 and 9. Then (1) holds and $P(T<L)=0.2$. Thus, during any replenishment there is a probability of 0.2 that one (and only one) new customer order will appear. Assumption (1) also implies that upon the arrival of any customer order there is at most one replenishment order not yet received at the stocking point. Furthermore, assumption (1) makes it possible to apply renewal-reward theory as mathematical tool for analysis. For a good overview of renewal-reward theory, see Tijms (2003, Chapter 2). Basically, when using renewal-reward theory one must identify a *regeneration point*, where any future development of the stochastic process is independent of its past. The time interval between two successive regeneration points is then denoted a *cycle*. It holds that the long-run behaviour of a regenerative stochastic process can be studied in terms of the behaviour of the process during a single cycle. Thus one is left with a somewhat easier task of studying what happens during a cycle. In our setting a regeneration point is the time point where the net inventory (the net inventory is the physical on-hand inventory minus the amount backlogged, see the definition in Silver et al., 1998, p. 233) is S, but a new customer order has just appeared. We call this order *the triggering order* and represent its size by the random variable X_0. Denote by N the number of orders (except the triggering order) that occurs between two successive regeneration points (which is the cycle). The random variable N has the following probability distribution:

$$P(N=i) = P(T \geq L)(P(T<L))^i, \quad i=0,1,2\ldots \tag{2}$$

Thus, N is geometrically distributed with mean $E[N]=P(T<L)/P(T\geq L)$. In Fig. 1 we have illustrated a cycle in the case where $N=2$.

Define the random variables (the random variable $\delta_{\{A\}}$ is 1 if the event A occurs and 0 otherwise):

$$O_0 = \delta_{\{X_0 \geq S+1\}}, \tag{3}$$

$$O_i = \delta_{\{X_{i-1}+X_i \geq S+1\}}, \quad i=1,2,\ldots, \tag{4}$$

$$B_0 = \max\{X_0 - S, 0\} \tag{5}$$

and

$$B_i = \max\{\min\{X_{i-1}, S\} + X_i - S, 0\}, \quad i=1,2,\ldots. \tag{6}$$

The random variable O_0 is 1 if the triggering order does not receive complete fulfillment and the random variable B_0 quantifies how much of the triggering order that is backordered. Similarly, the random variables O_i records if order number i in the cycle does not

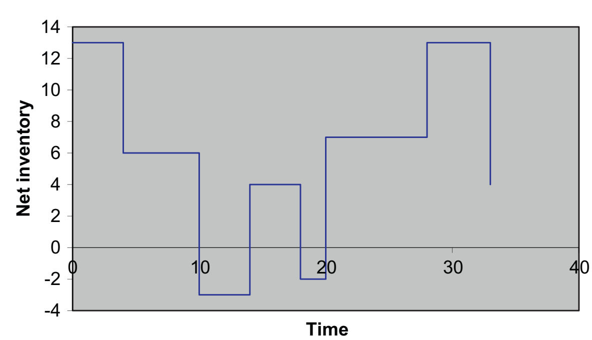

Fig. 1. *Illustration of a cycle*. In the example $S=13$ and $L=10$. The cycle starts in time point 4 and the triggering order is of size $X_0=7$ (replenishment received at time point 14). At time point 10 a new order, of size $X_1=7$, arrives (replenishment received at time point 20) causing 3 units to be backlogged. At time point 18 a new order of size, $X_2=6$, arrives (replenishment received at time point 28) causing 2 units to be backlogged. Because the next order arrives after time point 28, at time 28 all replenishment orders are received and the net inventory is S. The next order (arriving at time point 33) causes the start of a new cycle.