Int. J. Production Economics 133 (2011) 481–485

Contents lists available at ScienceDirect

Int. J. Production Economics

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

Comparison of a new bootstrapping method with parametric approaches for safety stock determination in service parts inventory systems

Chenxi Zhou [a], S. Viswanathan [b,*]

[a] Department of Marketing, Warrington College of Business Administration, University of Florida, Gainesville, FL 32611-7155, USA
[b] Nanyang Business School, Nanyang Technological University, Singapore 639798, Singapore

ARTICLE INFO

Article history:
Received 7 September 2008
Accepted 9 September 2010
Available online 25 September 2010

Keywords:
Forecasting
Intermittent demand
Safety stock determination
Bootstrapping
Inventory management

ABSTRACT

In this paper, we address the problem of forecasting and managing the inventory of service parts where the demand patterns are highly intermittent. Currently, there are two classes of methods for determining the safety stock for the intermittent item: the parametric and bootstrapping approaches. Viswanathan and Zhou (2008) developed an improved bootstrapping based method and showed through computational experiments that this is superior to the method by Willemain et al. (2004). In this paper, we compare this new bootstrapping method with the parametric methods of Babai and Syntetos (2007). Our computational results show that the bootstrapping method performs better with randomly generated data sets, where there is a large amount of (simulated) historical data to generate the distribution. On the other hand, with real industry data sets, the parametric method seems to perform better than the bootstrapping method.

1. Introduction

In this paper, we address the problem of forecasting and managing the inventory of items for which the demand patterns are highly intermittent. An intermittent demand pattern is defined as one where a sequence of demand data contain a large percentage of zero values. One common example of items with intermittent demand is service (or spare) parts. It is believed that if this intermittent demand can be forecasted properly, it can result in significant improvement in managing the service parts inventories.

The Croston (1972) method is considered to be the standard method for forecasting intermittent demand. This method has been shown to perform better than the single exponential smoothing (SES) technique and has been widely incorporated in commercial forecasting packages. Recently, many researchers have linked this method with inventory management. Some corrections or variations of Croston's method are used to determine the average per period demand. Thereafter, the order quantity or order-up-to-levels are determined by calculating the safety stock required over the lead time (LT). Based on separate forecasts for the demand size of non-zero demands and the time-interval between non-zero demands, many researchers (e.g. Willemain et al., 1994, Johnston and Boylan, 1996; Levén and Segerstedt (2004), and Babai and Syntetos, 2007) have used a parametric approach to derive the formula for the safety stock. See also chapter 16 of a recent book by Hyndman et al. (2008).

Parametric methods typically assume that the demand per period follows a particular probability distribution. The parameters of the assumed distribution (such as mean and variance) are estimated using historical data and updated based on the latest demand value in every period. The safety stock or order-up-to-levels are then calculated using the estimated demand distribution parameters. The parametric approach has low computational overheads and is quite accurate.

On the other hand, the bootstrapping method, a non-parametric approach, generates a large number of data points for the lead time demand by sampling from the historical demand data and demand interarrival data. By repeated sampling, a lead time distribution can be built (or at least the inventory position required for a certain service rate can be determined). Some researchers (e.g. Efron and Gong, 1983, and Willemain et al., 2004) have found bootstrapping to be efficient in dealing with intermittent demand items. The best known bootstrapping based method for determining the safety stock is by Willemain et al. (2004). In their method, the periods of positive demands are generated using a two-state Markov model and the actual demand sizes using past demand history. Their results show that the bootstrapping method produces more accurate forecasts of the distribution of demand over a fixed lead time than exponential smoothing and Croston's method. Recently,

* Corresponding author. Tel.: +65 6790 4798; fax: +65 6792 2313.
E-mail addresses: zhou.1985@ufl.edu (C. Zhou), vish@pmail.ntu.edu.sg (S. Viswanathan).

0925-5273/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.ijpe.2010.09.021

Viswanathan and Zhou (2008) developed an improved bootstrapping based method and showed through computational experiments that this is superior to the method by Willemain et al. (2004) both in the context of computer generated demand data and industrial data. The key highlight of this improved bootstrapping method is that instead of a two-state Markov model, the positive demand arrivals are generated using the historical distribution of the inter-demand intervals (or intervals between non-zero demand).

Viswanathan and Zhou (2008) compared their bootstrapping method with that of Willemain et al. (2004) using a computational study that was designed similar to the one used in this paper. Based on the distribution of the lead time demand generated using bootstrapping, the safety stock required was calculated under both methods for various service levels. The resulting total cost (comprised of holding and penalty cost) was calculated through a simulation experiment. For both the computer generated demand data and the real industrial data, the new method of Viswanathan and Zhou (2008) always obtained a lower average total cost.

In this paper, we compare this new bootstrapping method with the parametric methods of Babai and Syntetos (2007) in two ways. First, we conduct a simulation study to compare the two approaches on randomly generated demand data sets with long demand history. We then conduct a numerical study to compare the two approaches with industrial demand data sets, where the length of the historical demand data available is relatively short. Our computational results show that the bootstrapping method of Viswanathan and Zhou (2008) performed better than the parametric approach with the randomly generated demand data sets. With industrial demand data sets which have only limited historical information, the parametric approach performed better than the bootstrapping method. The managerial insight from this research is that when sufficiently large amount of historical demand information is available, bootstrapping works well, but with limited historical information (as is most often the case when scientific forecasting approaches are initially applied in a company), parametric approaches such as the method of Babai and Syntetos produce better forecasts and are better at controlling the total inventory related costs.

The rest of the paper is organized as follows. In Section 2, we provide a brief description of the methods being compared in this study. The design of the simulation experiments and the results are discussed in Sections 3 and 4. In Section 3, we discuss the simulation with randomly generated data and in Section 4, we discuss the simulation using industrial data sets. Finally, some concluding remarks are presented in Section 5.

2. Methods for safety stock determination

We now describe the parametric and bootstrapping methods compared in this study. The most recent parametric method is the one proposed by Babai and Syntetos (2007). In this paper, the average demand per period is first estimated using a variation of the Croston's method. This and an estimate of the demand variance per period is used to determine the order up to level.

The Croston's original method estimates the mean demand per period by applying exponential smoothing separately to the intervals T_t between consecutive non-zero demands and their demand sizes Z_t (of non-zero demands). Let $\hat{T}_{t-1}$be the forecast of the demand interval. And let $\hat{Z}_{t-1}$be the forecast of the demand size for period t. Croston's method works as follows: in period t, if no demand occurs, then the estimates of the demand interval and the demand size remain unchanged, namely, $\hat{T}_t=\hat{T}_{t-1}$and $\hat{Z}_t=\hat{Z}_{t-1}$. If a demand does occur, i.e., $Z_t>0$, then the estimates are updated to

$$\hat{T}_t=\hat{T}_{t-1}+\alpha(T_t-\hat{T}_{t-1})$$

$$\hat{Z}_t=\hat{Z}_{t-1}+\beta(Z_t-\hat{Z}_{t-1})$$

in which parameters α and β are smoothing constants. This follows the adaptation by Schultz (1987) who specified different smoothing parameters for demand intervals and sizes, instead of identical parameters as specified by Croston (1972). Hence, combining the estimates of demand size and demand interval, one can obtain the estimates of mean demand per period after period t, $\hat{D}_t=\hat{Z}_t/\hat{T}_t$.

Syntetos and Boylan (2005) adopted a different set of calculations from that of Croston (1972) for determining the estimate of the mean demand per period. First of all, in an earlier work, Syntetos and Boylan (2001) show that there is a bias in the Croston's (1972) expression for the forecast of the mean demand per period. The corrected estimate of the mean demand per period proposed by Syntetos and Boylan (2005) is $F_t=(1-\alpha/2)(\hat{Z}_t/\hat{T}_t)$. Babai and Syntetos (2007) used this to determine the mean of the lead time demand. The variance (or mean square error) of the mean demand per period is estimated by Babai and Syntetos (2007) as $MSE_t=\alpha(D_t-F_{t-1})^2+(1-\alpha)MSE_{t-1}$. The lead time demand is assumed to follow a negative binomial distribution with mean and variance values calculated as above. The optimal order-up-to-level S_t at period t, is the smallest S_t that satisfies $\sum_{x=0}^{S_t}\Phi_D(x)\geq CSL$, where CSL represents cycle service level and $\Phi_D(x)$ is the probability density function of the total lead time demand.

In this paper, the two parametric methods we study are both based on Babai and Syntetos (2007), where the Syntetos and Boylan (2005) estimator is used to determine the mean demand per period and the variance of the demand per period is estimated using the formula of Babai and Syntetos (2007) mentioned earlier. One method (referred to as BS1 in the rest of paper) determines the order up to level by assuming a negative binomial lead time demand distribution while the other (referred to as BS2) assumes a normal lead time demand.

Recently, Viswanathan and Zhou (2008) developed an improved bootstrapping based method and showed through computational experiments that this is superior to the method by Willemain et al. (2004). The key highlight of the newly proposed bootstrapping method is that instead of a two-state Markov chain, the positive demand arrivals are generated using the historical distribution of the inter-demand intervals. A histogram of the distribution of the lead time demand is built by repeatedly generating all the demand arrivals during the lead time. The demand arrivals and demand sizes are generated by sampling from the historical data of the inter-arrival time and demand sizes. The optimal order-up-to-level can then be calculated based on the required CSL. The method is summarized in Table 1.

Table 1
Steps of the bootstrapping method of Viswanathan and Zhou (2008)

Step 1: Obtain histogram of the historical demand data (including both demand size data and demand interval data) in a chosen time bucket.
Step 2: Randomly generate demand interval according to the corresponding histogram. Update the time horizon, which is used to count the time passing by.
Step 3: If the time horizon is equal to or less than the lead time, randomly generate demand size according to the demand size interval histogram. Then go to Step 2. Else, sum the generated demand sizes over the lead time and get one predicted value of the lead time demand. Then go to Step 4.
Step 4: Repeat Steps 2 and 3, 1000 times.
Step 5: Sort and generate the resulting distribution of the lead time demand. Using this distribution, obtain the order-up-to-level according to the *CSL*.