Int. J. Production Economics 157 (2014) 220–237

Contents lists available at ScienceDirect

Int. J. Production Economics

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

More grip on inventory control through improved forecasting: A comparative study at three companies

E. van Wingerden [a], R.J.I. Basten [b,*], R. Dekker [c], W.D. Rustenburg [d]

[a] Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands
[b] University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands
[c] Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands
[d] Gordian Logistic Experts, Groenewoudsedijk 63, 3528 BG, Utrecht, The Netherlands

ARTICLE INFO

Article history:
Received 30 September 2012
Accepted 26 August 2014
Available online 6 September 2014

Keywords:
Inventory
Spare parts
Forecasting

ABSTRACT

Inventory control for parts with infrequent demands is difficult since forecasting their demand is problematic. Traditional forecasting methods, such as moving average and single exponential smoothing, are known not to suffice since they do not cope well with periods with zero demands. Croston type methods and bootstrapping methods are more promising. We propose a new bootstrapping method, which we term empirical plus. The added value of this method lies in the fact that it explicitly takes into account that besides the demand, also the supply lead time is stochastic. We compare its performance with a number of methods from all three above-mentioned categories. Opposite to what is done in most comparative studies, we do not focus on performance metrics that are related directly to the forecasting results (e.g., mean squared error), but we focus on the resulting inventory control policy (achieved fill rate and holding costs). We use in our study large data sets from three companies, which we make publicly available. We find that our empirical plus method outperforms the other methods when the average inter-demand interval is large and the squared coefficient of variation of the demand size is small. This class of parts often consists of the expensive parts, for which forecasting is both difficult, because of the infrequent demands, and important, because of the price. The Syntetos Boylan approximation performs best on the other classes of parts. These findings may be used in practice to use the right forecasting method for each type of part, thus achieving more cost-effective spare parts inventory control.

1. Introduction

Companies making, using, or maintaining capital goods typically have large inventories of spare parts. According to the Aberdeen Group (2005) their value exceeds a trillion dollars; commercial airlines are estimated to have over $40 billion worth of spare parts (Harrington, 2007). Hence, there is a high need to reduce the inventory, also because a lot of these parts may never be used and have to be discarded at some point in time. At the same time, a lack of spare parts can lead to costly unavailability. This is a problem for a user that maintains its own assets, but also for an original equipment manufacturer maintaining its installed base, for example because of penalties that need to be paid.

The management of spare parts stock is a separate discipline, different from finished product inventory management, as demand is often characterized as being infrequent with sudden high demands. This makes the problem of determining the right amount of stocks an important scientific and practical problem. Many methods have been developed for the inventory control of spare parts, see for example the overviews by Kennedy et al. (2002) and Basten and Van Houtum (2014). In many inventory control models a specific demand process, like the Poisson process, is assumed with known parameters. Yet in any real case, neither the demand process nor its parameters are known and they have to be estimated or forecasted.

A recent stream in the forecasting literature investigates models to predict the demand for spare parts. Yet there are few tests of these methods with real data and not always all relevant characteristics of the parts have been revealed. Even fewer papers have addressed the combination of forecasting and spare parts control, as the typical performance criterion in forecasting, the mean squared error, is not the leading characteristic in inventory control.

Another important practical question is whether one approach can be used for all items or whether different approaches should be used for different types of parts. In other words, what is the best classification of spare parts with respect to their inventory

* Corresponding author. Tel.: +31 53 489 4007.
E-mail addresses: e.v.wingerden@tue.nl (E. van Wingerden), r.basten@utwente.nl (R.J.I. Basten), rdekker@ese.eur.nl (R. Dekker), jw.rustenburg@gordian.nl (W.D. Rustenburg).

http://dx.doi.org/10.1016/j.ijpe.2014.08.018
0925-5273/

control? This is particularly relevant when developing a decision support system for spare parts inventory control.

In this paper, we compare several spare part demand forecasting methods using real data from three companies. The number of different spare parts they keep is quite large and ranges from some 4000 to over 250,000 different parts. Still, we consider much fewer parts for several reasons that we discuss in Section 5.1. One company is involved in the maintenance of passenger trains in the Netherlands, the second supplies equipment for trains and other industrial systems, and the third organization is involved in the maintenance of naval ships. Apart from characterizing their spare parts demand situations, we also provide readers access to the data in a complementary file, which allows them to do new analyses.

The methods that we are comparing are simple moving average, single exponential smoothing, double exponential smoothing, three Croston type methods, a so-called MSE method, and two bootstrapping methods. The MSE method is a method that picks the best traditional or Croston type method per part, based on the resulting MSE (mean squared (forecasting) error). The key reason to consider bootstrapping methods is that they immediately forecast the demand over the lead time, opposed to traditional and Croston type methods that forecast the demand per period and then multiply that by the lead time. Hence the bootstrapping methods can incorporate all kind of hidden correlations in demand. One of the two bootstrapping methods that we use is a new method, in which we explicitly incorporate the fact that not only demand, but also the supply lead time is stochastic. We call this method the empirical plus method.

Using all data, we compare the performance of these methods combined with an inventory control method by establishing trade-off curves between inventory holding costs and service levels. This gives a better view of their performance than comparing costs for one service level.

Summarizing, we contribute in four ways:

- We propose a new bootstrapping method that explicitly incorporates a stochastic supply lead time.
- We perform a comparative study, focusing on the resulting inventory control, in which we compare traditional, Croston type, and bootstrapping methods.
- We further incorporate in our comparative study the so-called MSE method, through which we show that focusing on the MSE when forecasting may not result in the best inventory control policy.
- For the study we use an extensive data set from three companies, which we make publicly available.

We find that for the class of parts for which accurate forecasting is most important in practice, since demands are infrequent and the parts have high prices on average, our new bootstrapping method performs best.

The set-up of the remainder of this paper is as follows. We discuss the relevant literature in Section 2, including detailed descriptions of the forecasting methods that we use in our study. We next describe the three companies of which we use data sets in Section 3. In Section 4, we propose the empirical plus method. We explain the set-up of the simulation study in Section 5, and we discuss the results in Section 6. Finally, we give conclusions and recommendations in Section 7.

2. Literature review

The literature on forecasting of spare parts, or slow moving items, is quite extensive, starting with the paper by Croston (1972). A lot of progress has been made from the mid 1990s onwards and the topic still receives a lot of attention. There exist two relatively recent review papers that include forecasting methods: Boylan and Syntetos (2010) focus on spare parts forecasting, including forecast support systems, and Syntetos et al. (2009) review the literature on forecasting for general inventory control; the sections on forecasting methods for spare parts in these two papers overlap to some extent. Here, we discuss only the most relevant papers in the context of our study. First, since we perform a comparative study using a number of existing forecasting methods and a newly developed forecasting method, we refer to literature on the existing methods (we give the exact calculations for the methods in the appendix). Second, we discuss a couple of recent and related comparative studies, especially those that consider the resulting inventory control, as we do.

We discuss a number of forecasting methods that are commonly used in industry (the traditional methods) and some methods that are reported in the literature to perform well (the Croston type and bootstrapping methods). Our aim is not to compare (and thus discuss) all forecasting methods, but to compare methods from each of the three categories. The first type of forecasting methods are the *traditional methods* that are not designed specifically for slow moving items. These methods are discussed in various text books (e.g., Axsäter, 2006) and we consider the *simple moving average* (SMA; by some authors just called moving average), *single exponential smoothing* (SES; by some authors just called exponential smoothing), and *double exponential smoothing* (DES). The downside of these models is that they do not explicitly take into account that there may be periods with zero demands, whereas this is quite common for spare parts.

To cope with this problem, Croston (1972) proposes *Croston's method*, an exponential smoothing method that updates forecasts only in periods with positive demands. It forecasts both the demand size in periods with positive demands and the average number of periods between two periods with positive demands, the inter-demand interval. These forecasts are combined to come up with a forecast of the average demand per period. Since Syntetos and Boylan (2001) find that Croston's method is biased, Syntetos and Boylan (2005) propose an improved version of Croston's method, the *Syntetos Boylan approximation* (SBA). A recently proposed new Croston type method is that by Teunter et al. (2011): *Teunter Syntetos Babai* (TSB). Like the other Croston type methods, this method adapts the expected demand size only in periods with positive demand. However, the expected inter-demand interval is adapted in each period, so that inventory levels may be decreased when demand decreases (e.g., when a component has become obsolete). Together, we call these three methods the *Croston type methods*.

The third and final category of forecasting methods are the *bootstrapping methods*. Probably the most well known method in this category is that of Willemain et al. (2004). The authors construct a two-state Markov chain, with one state representing a period without demand and the other a period with positive demand. The transition probabilities are determined from the historical data. A lead time demand sample is then constructed by taking L steps in the Markov chain, with L being the lead time, and filling each non-zero demand with a random historical demand. The samples are used to construct an empirical distribution of the lead time demand. Porras and Dekker (2008) propose another method that slightly outperforms that of Willemain et al. (2004) on a data set from a refinery. Since the method of Porras and Dekker (2008) is also easier to understand and implement, we discuss (and use) the latter method only, in Section 4. The method is called the *empirical method*, because it determines an empirical demand distribution by sampling from historical demands. Our method is based on that of Porras and Dekker (2008) and it is