Journal of Business Research 68 (2015) 1746–1752

Contents lists available at ScienceDirect

Journal of Business Research

Forecasting intermittent inventory demands: simple parametric methods vs. bootstrapping☆

Aris A. Syntetos [a,*], M. Zied Babai [b,1], Everette S. Gardner Jr. [c,2]

[a] Cardiff Business School, Cardiff University, Aberconway Building, Colum Drive, Cardiff CF10 3EU, UK
[b] Kedge Business School, 680 cours de la Libération, 33400 Talence, France
[c] Bauer College of Business, University of Houston, Houston, TX 77204-6021, USA

ARTICLE INFO

Article history:
Received 1 October 2013
Received in revised form 1 February 2014
Accepted 1 February 2015
Available online 8 April 2015

Keywords:
Inventory management
Operations forecasting
Time series methods
Intermittent demand

ABSTRACT

Although intermittent demand items dominate service and repair parts inventories in many industries, research in forecasting such items has been limited. A critical research question is whether one should make point forecasts of the mean and variance of intermittent demand with a simple parametric method such as simple exponential smoothing or else employ some form of bootstrapping to simulate an entire distribution of demand during lead time. The aim of this work is to answer that question by evaluating the effects of forecasting on stock control performance in more than 7,000 demand series. Tradeoffs between inventory investment and customer service show that simple parametric methods perform well, and it is questionable whether bootstrapping is worth the added complexity.

1. Introduction

1.1. The intermittent demand forecasting problem

Traditionally, the relevant literature treats inventory management and demand forecasting as independent problems. Most inventory papers ignore forecasting altogether and simply assume that the distribution of demand and all its parameters are known, while most forecasting papers do not evaluate the stock control consequences of employing different forecasting methods. The interactions between forecasting and stock control are analyzed in this paper for items with intermittent demand. Such demand series are characterized by zero demand occurrences interspersed by positive demands. The choice of forecasting method is shown to be an important determinant of the customer service that can be obtained from a given level of inventory investment.

Since the early work of Brown (1959), the problem of forecasting for fast moving inventory items has attracted an enormous body of academic research. However, forecasting for items with intermittent demand has received far less attention, even though such items typically account for substantial proportions of stock value and revenues. Intermittent demand items dominate service and repair parts inventories in many industries (including the process industries, aerospace, automotive, IT and the military sector), and they may constitute up to 60% of total stock value (Johnston, Boylan, & Shale, 2003). A survey by Deloitte (2011) benchmarked the service businesses of many of the world's largest manufacturing companies with combined revenues reaching more than $1.5 trillion; service operations accounted for an average of 26% of revenues. Thus small improvements in management of intermittent demand items may be translated to substantial cost savings; research in this area has direct relevance to a wide range of companies and industries.

In addition, intermittent items are at the greatest risk of obsolescence, and case studies have documented large proportions of dead stock in many different industrial contexts (Hinton, 1999; Molenaers, Baets, Pintelon, & Waeyenberg, 2010; Syntetos, Keyes & Babai, 2009). Improvements in forecasting may be translated to significant reductions in wastage or scrap with further environmental implications.

Intermittent demand series are difficult to forecast because they usually contain a (significant) proportion of zero values, with nonzero values mixed in randomly. When demand occurs the quantity may be highly variable (Cattani, Jacobs, & Schoenfelder, 2011). One critical research question is whether one should make point forecasts of the mean and variance of intermittent demand with a simple parametric method or else employ some form of bootstrapping to simulate an

☆ The authors thank Ruud Teunter, University of Groningen, Nikos Kourentzes, Lancaster University, Thomas Willemain, Rensselaer Polytechnic Institute, and Powell Robinson, University of Houston, for comments on earlier drafts of this paper. At least one of the authors has read each reference in this paper. We contacted Ruud Teunter and Thomas Willemain to ensure that our study summarizes their work accurately.

* Corresponding author. Tel.: +44 29 2087 6572; fax: +44 29 2087 4301.
E-mail addresses: SyntetosA@Cardiff.ac.uk (A.A. Syntetos), Mohamed-Zied.Babai@kedgebs.com (M. Zied Babai), EGardner@uh.edu (E.S. Gardner).
[1] Tel.: +33 5 56 84 63 51; fax: +33 5 56 84 55 00.
[2] Tel.: +1 713 743 4744; fax: +1 713 743 4940.

http://dx.doi.org/10.1016/j.jbusres.2015.03.034
0148-2963/

entire distribution of demand during lead time. Is bootstrapping worth the added complexity? The aim of this study is to answer that question in an empirical investigation of forecasting more than 7,000 inventory demand series.

1.2. Research background

Two parametric methods, simple exponential smoothing (SES) and Croston's (1972) method with corrections by Rao (1973), are widely used to forecast intermittent demand. SES forecasts the mean level of demand for both non-zero and zero demand periods, treating them in the same way, while Croston makes separate forecasts of the mean level of non-zero demand and the mean inter-arrival time (time between demand occurrences). Croston assumes that the distribution of nonzero demand sizes is normal, the distribution of inter-arrival times is geometric, and that demand sizes and inter-arrival times are mutually independent. Shenstone and Hyndman (2005) challenge these assumptions and show that Croston's method is inconsistent with the properties of intermittent demand data. The primary problem is that Croston's method assumes stationarity, while any possible model underlying the method must be non-stationary. Furthermore, the underlying model must be defined on a continuous sample space that can take on either negative or positive demand values, something that is inconsistent with the reality that demand is always non-negative.

Despite its theoretical shortcomings, Croston's method has been successful in empirical research (see the review in Gardner, 2006) and is widely used in practice. Both Croston and SES are available in demand planning modules of component based enterprise and manufacturing solutions (e.g. Industrial and Financial Systems – IFS AB) and in integrated real-time sales and operations planning processes (e.g. SAP Advanced Planning and Optimisation – APO 4.0).

Many improvements to Croston's original method appear in the literature including Johnston and Boylan (1996), Snyder (2002), Syntetos and Boylan (2005), Shale, Boylan, and Johnston (2006), and Teunter, Syntetos, and Babai (2011). The Syntetos and Boylan method (known as the SBA method for Syntetos-Boylan approximation), is the only Croston improvement that has substantial empirical support. Although Croston claims that his method is unbiased, Syntetos and Boylan (2001) show that the opposite is true and present an improved method that corrects for bias (Syntetos & Boylan, 2005). The SBA method was tested by Eaves and Kingsman (2004) using a sample of more than 11,000 monthly repair parts demand series from Royal Air Force (RAF) inventories. The results varied somewhat depending on the degree of aggregation of the data (weekly, monthly, quarterly) and the type of demand pattern (ranging from smooth to highly intermittent). However, in general the SBA method was more accurate than SES and the original Croston method. Another study by Gutierrez, Solis, and Mukhopadhyay (2008) reaches similar conclusions. In the empirical study below, all three parametric alternatives are tested: SES, Croston's original method, and the SBA method.

Given the parametric point forecasts, a demand distribution is necessary to set inventory levels. Both the Poisson and Bernoulli processes fit demand arrivals, that is, the probability of demand occurring (Dunsmuir & Snyder, 1989; Eaves, 2002; Janssen, 1998; Willemain, Smart, Shockor, & DeSautels, 1994). Regarding the size of demand when it occurs, various suggestions have been made for distributions that are either monotonically decreasing or unimodal positively skewed. With Poisson or Bernoulli arrivals of demands and any distribution of demand sizes, the resulting distribution of total demand over a fixed lead time is compound Poisson or compound Bernoulli, respectively. Compound Poisson distributions are simpler and have empirical evidence in their support (e.g., Boylan & Syntetos, 2008). In this empirical study, demand is modeled with the Negative Binomial Distribution (NBD), which performed well in the empirical study by Snyder, Ord, and Beaumont (2012). The NBD is a compound distribution in which the number of demands in each period is Poisson distributed, with random demand sizes governed by a logarithmic distribution.

As the data become more erratic, the true demand size distribution may not conform to any standard theoretical distribution, and it may be that non-parametric approaches (that do not rely upon any underlying distributional assumption) may improve stock control. Numerous bootstrapping methods are available to randomly sample (with or without replacement) observations from demand history to build a histogram of the lead-time demand distribution. Alternative bootstrapping methods are found in Efron (1979), Snyder (2002), Willemain, Smart, & Schwarz (2004, hereafter WSS), Porras and Dekker (2008), Teunter and Duncan (2009), Zhou and Viswanathan (2011), and Snyder et al. (2012). The most robust bootstrapping method appears to be that of WSS, a method patented earlier by Willemain and Smart (2001). WSS is tested in this paper; further discussion on the justification for excluding other bootstrapping alternatives follows in the next section.

In a large empirical study, WSS claim significant improvements in forecasting accuracy over both SES and Croston's estimator. However, Gardner and Koehler (2005) criticize this study because the authors do not use the correct lead time demand distribution for either SES or Croston's method, and they do not consider published improvements to Croston's method, such as the SBA method (see Willemain et al., 2005, for a rejoinder). These mistakes are corrected in this empirical study.

The Teunter and Duncan (2009) empirical study is similar to the study in the present paper. Using a sample of demand series for military spare parts, Teunter and Duncan compare the inventory and service tradeoffs that result from forecasting with the same parametric methods tested below. They also test a simple bootstrapping method in which they sample lead time demand with replacement to estimate mean and variance, which are then fed into a normal distribution to set stock levels. Reliance on the normal distribution defeats the purpose of bootstrapping, which does not require a distributional assumption.

1.3. Organization of the paper

Section 2 explains the parametric and bootstrapping methods. Section 3 discusses the data tested, performance measurement, and simulation procedures. Empirical results are given in Section 4; in contrast to most previous research in intermittent demand forecasting, results are presented in terms of stock control performance rather than forecast accuracy. Section 5 discusses implications of the results followed by conclusions and opportunities for further research in Section 6.

2. Forecasting intermittent demand

2.1. Parametric forecasting

Simple exponential smoothing (SES) is written:

$$S_t = \alpha X_t + (1-\alpha) S_{t-1}, \tag{1}$$

where α is the smoothing parameter, X_t is the observed value of both zero and nonzero demand, and S_t is the smoothed average as well as the forecast for next period. Although SES is widely used to forecast intermittent demand, the method has important limitations. Exponential smoothing weights recent data more heavily, which produces forecasts that are biased high just after a demand occurs and biased low just before a demand. Replenishment quantities are likely to be determined by forecasts made just after a demand, resulting in unnecessarily high stock levels most of the time.

In an attempt to compensate for these problems, Croston's (1972) method forecasts two components of the time series separately, the observed value of nonzero demand (D_t) and the inter-arrival time of