European Journal of Operational Research 214 (2011) 606–615

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Production, Manufacturing and Logistics

Intermittent demand: Linking forecasting to inventory obsolescence

Ruud H. Teunter [a,*], Aris A. Syntetos [b,1], M. Zied Babai [c,d,2]

[a] University of Groningen, Netherlands
[b] Centre for OM, Management Science and Statistics, Salford Business School, University of Salford, Manchester M5 4WT, UK
[c] BEM-Bordeaux Management School, 680 cours de la Libération, 33405 Talence Cedex, France
[d] Princess Fatimah Alnijris's Research Chair for AMT, Industrial Engineering Department, King Saud University, Saudi Arabia

ARTICLE INFO

Article history:
Received 2 July 2010
Accepted 13 May 2011
Available online 23 May 2011

Keywords:
Forecasting
Intermittent demand
Inventory control
Obsolescence

ABSTRACT

The standard method to forecast intermittent demand is that by Croston. This method is available in ERP-type solutions such as SAP and specialised forecasting software packages (e.g. Forecast Pro), and often applied in practice. It uses exponential smoothing to separately update the estimated demand size and demand interval whenever a positive demand occurs, and their ratio provides the forecast of demand per period. The Croston method has two important disadvantages. First and foremost, not updating after (many) periods with zero demand renders the method unsuitable for dealing with obsolescence issues. Second, the method is positively biased and this is true for all points in time (i.e. considering the forecasts made at an arbitrary time period) and issue points only (i.e. considering the forecasts following a positive demand occurrence only). The second issue has been addressed in the literature by the proposal of an estimator (Syntetos–Boylan Approximation, SBA) that is approximately unbiased. In this paper, we propose a new method that overcomes both these shortcomings while not adding complexity. Different from the Croston method, the new method is unbiased (for all points in time) and it updates the demand probability instead of the demand interval, doing so in every period. The comparative merits of the new estimator are assessed by means of an extensive simulation experiment. The results indicate its superior performance and enable insights to be gained into the linkage between demand forecasting and obsolescence.

1. Introduction

Intermittent demand appears sporadically, with some time periods showing no demand at all. When a demand occurs, the demand size may be constant or variable, perhaps highly so. Intermittent demand items may be any Stock Keeping Unit (SKU) within the range of products offered by an organization at any level of the supply chain. Such items may collectively account for up to 60% of the total stock value (Johnston et al., 2003) and are particularly prevalent in the aerospace, automotive, military and IT sectors. They are often the items at greatest risk of obsolescence (Porras and Dekker, 2008).

The occurrence of (many) periods with zero demand renders traditional forecasting techniques such as simple exponential smoothing or simple moving average unsuitable. Those methods ignore the fact that intermittent demand patterns are built from two elements: demand size and demand probability (or its inverse, the demand interval). The Croston (1972) method does differentiate between these two elements. Using exponential smoothing, the demand size and the demand interval are separately updated after each period with a positive demand. The ratio of the former over the latter then provides a forecast for the future demand per period.

The Croston method is often applied in practice (see, for example Fildes et al., 2008) and incorporated in Enterprise Resource Planning (ERP) type solutions such as SAP and specialised forecasting software packages such as Forecast Pro. However, the method has two important disadvantages. First, it is positively biased, as will be discussed further in the next section. Second and most important, by not updating after periods with zero demand, the forecasts become outdated after (many) periods with zero demand and unsuitable for estimating the risk of obsolescence. Obsolescence is an important topic for inventory management, especially for slow moving, intermittent demand. Molenaers (2010) discussed a case study where 54% of the parts stocked at a large petrochemical company had seen no demand for the last 5 years. Syntetos et al. (2009a) evaluated the inventory practices employed in the European spare parts logistics network of a Japanese manufacturer. One case was reported in Sweden where parts in stock had not 'moved' at all over the preceding 10 years. Hinton (1999) described that the US Department of Defence holds 60% excess of spare parts, with 18% of the parts (with a total value of $1.5 billion)

* Corresponding author. Address: University of Groningen, Dept. Operations, P.O. Box 800, 9700 AV, Netherlands. Tel.: +31 50 3638617, fax: +31 50 3632032.

E-mail addresses: r.h.teunter@rug.nl (R.H. Teunter), a.syntetos@salford.ac.uk (A.A. Syntetos), mohamed-zied.babai@bem.edu (M.Z. Babai).

[1] Tel.: + 44 161 2955804; fax: +44 161 2955556.

[2] Tel.: +33 556846351; fax: +33 556845500.

0377-2217/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2011.05.018

having no demand at all. In such situations, decisions to remove 'dead stock' and reduce inventory levels are of great importance. Those decisions rely on accurate demand forecasts and associated risk of obsolescence estimates, especially after long intervals with zero demand.

Despite its importance, obsolescence is ill-researched. We refer interested readers to Teunter (1998, Section 2.2) for a review of obsolescence in the general area of inventory control, and to Kennedy et al. (2002, Section 3.5) for obsolescence in spare parts management. Recent contributions are those by Pinçe and Dekker (2011) and Van Jaarsveld and Dekker (2008). To the best of our knowledge, obsolescence has not been discussed in the forecasting literature. Fildes et al. (2008) and Syntetos et al. (2009b) who recently reviewed the link between forecasting and Operational Research (OR)/inventory planning, also did not discuss inventory obsolescence. Finally, recent attempts to link explicitly the intermittent nature of demand patterns to stock control (Syntetos et al., 2009c; Teunter et al., 2010) also did not take this issue into account.

In this paper, we propose a new forecasting method for intermittent demand that is always up-to-date and can deal with obsolescence as well as other inventory decisions. Different from the Croston method, which updates the demand interval, the new method updates the demand probability. This may appear as a small difference, since the demand interval is the inverse of the underlying probability of demand occurrence. However, the demand probability can be updated in every period, whereas the demand interval can only be updated after a positive demand. Moreover, as we will show in Section 3, using demand probability instead of demand interval eliminates the forecasting bias when considering an arbitrary point in time.

The new method achieves a high flexibility by using different smoothing constants for demand size and demand probability. We will derive an exact expression for the variance of the new method that is useful in setting the smoothing constants. Higher values make the new method more adaptable, but also increase its variance.

The performance of the new method (TSB for short – after Teunter, Syntetos and Babai) is evaluated through a simulation experiment on theoretically generated data. The experimental structure allows the evaluation of TSB's performance both on stationary and non-stationary demand data, the latter being generated in such a way so as to capture possible obsolescence-related scenarios. Performance is reported through the consideration of the Mean signed Error (ME – that expresses bias) and the Mean Squared Error (MSE – that is directly associated with the issue of variance).

Of course, the new method or in fact any other forecasting method can not prevent obsolescence completely. Indeed, obsolescence may be the outcome of bad management, where changes in the demand rate go unnoticed (for too long). However, with thousands of slow moving items in stock all too often is not easy at all to determine when an item should be discontinued. Having a forecasting method that automatically adjusts a forecast downward after (long) periods without demand, as the new method does, is certainly preferable with respect to (identifying) obsolescence. The same could be achieved by using a simple moving average (MA) or simple exponential smoothing (SES). Both methods are also unbiased when considering an arbitrary point in time. However, and as discussed above, the relevant forecasts are not built from constituent elements (sizes of demand and probability of demand occurrence). The explicit consideration of the probability of demand occurrence is a major determinant of the intuitive appeal and operational implications of the new estimator. With regards to the latter issue, a number of studies have established the fact that compound stock control models that rely upon both constituent elements perform considerably better than traditional stock control formulations that rely solely on estimates of mean demand (see, for example, Syntetos et al., 2009c; Teunter et al., 2010).

The remainder of this paper is organized as follows. In the next section, we review the most relevant literature. Section 3 introduces notations and formally defines the new forecasting method along with the demonstration of its unbiased nature and the derivation of its variance under stationary demand. A discussion on setting the smoothing constants for the new method is also included. In Section 4 the details governing our simulation experiment are outlined, after which the results are presented and discussed in Section 5. We end with conclusions and directions for further research in Section 6.

2. Research background

The standard method used in industry to forecast intermittent demand requirements is the Croston method (Croston, 1972). The method is incorporated in statistical forecasting software packages (e.g. Forecast Pro), and demand planning modules of component based enterprise and manufacturing solutions (e.g. Industrial and Financial Systems – IFS AB). It is also included in integrated real-time sales and operations planning processes (e.g. SAP Advanced Planning and Optimisation-APO 4.0).

Croston suggested separating intermittent demand series into their constituent elements, i.e. demand sizes and inter-demand intervals and separately updating their estimates through simple exponential smoothing (SES). He further suggested using the same value for the smoothing constant of both sizes and intervals, although the method may also be applied through the consideration of different smoothing constants (see also Schultz, 1987). The ratio of those estimates (sizes/intervals) may then be used as an estimate of the demand per period. The reciprocal of the inter-demand interval estimate is used essentially as the estimate of the probability of demand occurrence (that was modelled as following a Bernoulli sequence). However, Syntetos and Boylan (2001) showed Croston's method to suffer from a positive 'inversion bias' (for a random variable X: $1/E(X) \neq E(1/X)$) resulting in over-forecasting mean demand. Subsequently, they proposed a modification, the so-called Syntetos–Boylan Approximation or SBA for short (Syntetos and Boylan, 2005). The modified method applies a deflating factor $(1 - \alpha/2)$ that is linear in the smoothing constant (used for updating the inter-demand interval) to the Croston method. The method has been shown in a number of independent empirical studies to compare very favourably to Croston's estimator (see, for example, Eaves and Kingsman, 2004; Gutierrez et al., 2008). However, some (negative) bias remains and there are indeed cases when the SBA method may be more biased than the original Croston method (Wallström and Segerstedt, 2010; Teunter and Sani, 2009). Also, the deflating factor renders the method less intuitive, which may hinder its implementation in practice. The latter disadvantage applies even stronger to a more complicated modification proposed by Syntetos (2001) (see also Teunter and Sani, 2009), although that does remove the bias completely. Levén and Segerstedt (2004) also proposed a modified Croston method. Their method does not update the inter-demand interval and demand size separately, but does use the distinction under concern to directly update the demand per period. By doing so, this method avoids the original bias-causing issue. However, it is still biased and in fact even more heavily so than the original estimator (Boylan and Syntetos, 2007; Teunter and Sani, 2009).

Another disadvantage of the Croston method and all the above discussed variants is that the forecast is only updated after periods with positive demand. As a result, these methods are not up-to-date after (many) periods with zero demand and therefore cannot be used to estimate the risk of obsolescence and deal with the removal of excess/dead stock.