Int. J. Production Economics 156 (2014) 180–190

Contents lists available at ScienceDirect

Int. J. Production Economics

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

On intermittent demand model optimisation and selection

Nikolaos Kourentzes *

Department of Management Science, Lancaster University Management School, Lancaster, Lancashire LA1 4YX, UK

ARTICLE INFO

Article history:
Received 9 October 2013
Accepted 9 June 2014
Available online 17 June 2014

Keywords:
Intermittent demand
Croston's method
SBA method
TSB method
Forecasting
Optimisation
Model selection

ABSTRACT

Intermittent demand time series involve items that are requested infrequently, resulting in sporadic demand. Croston's method and its variants have been proposed in the literature to address this forecasting problem. Recently other novel methods have appeared. Although the literature provides guidance on the suggested range for model parameters, a consistent and valid optimisation methodology is lacking. Growing evidence in the literature points against the use of conventional accuracy error metrics for model evaluation for intermittent demand time series. Consequently these may be inappropriate for parameter or model selection. This paper contributes to the discussion by evaluating a series of conventional time series error metrics, along with two novel ones for parameter optimisation for intermittent demand methods. The proposed metrics are found to not only perform best, but also provide consistent parameters with the literature, in contrast to conventional metrics. Furthermore, this work validates that employing different parameters for smoothing the non-zero demand and the inter-demand intervals of Croston's method and its variants is beneficial. The evaluated error metrics are considered for automatic model selection for each time series. Although they are found to perform similar to theory driven model selection schemes, they fail to outperform single models substantially. These findings are validated using both out-of-sample forecast evaluation and inventory simulations.

1. Introduction

Spare parts are typically demanded in a sporadic or an intermittent fashion. This makes such time series different from conventional ones, due to the presence of several periods with zero demand. Some examples where intermittent demand can appear are listed by Willemain et al. (2004), including cases of demand of heavy machinery and respective spare parts, aircraft service parts, electronics, maritime spare parts, etc. Johnston et al. (2003) identified that in many cases such items can account for up to 60% of the total stock value. Due to their slow moving nature, such items are at greatest risk of obsolescence. This can have substantial impact on the operations of organisations, which tie resources in stocking items of this nature (Boylan and Syntetos, 2010). Companies often hold more than the necessary stock, to account for poor demand forecasts (Ghobbar and Friend, 2003), thus making accurate forecasts important.

Croston (1972) first argued that traditional time series methods, like exponential smoothing, do not produce reliable forecasts for intermittent demand time series and instead proposed an alternative method. Several studies have verified the good forecasting accuracy and inventory performance of this method; for example see Willemain et al. (1994, 2004), Johnston and Boylan (1996), and Syntetos and Boylan (2006). The literature has identified several improvements of the original method, correcting it for bias (Syntetos and Boylan, 2005), or proposing new methods for intermittent demand which is able to overcome structural limitations of the original method, such as obsolescence issues (Teunter et al., 2011). However, most of the literature uses ad-hoc model parameters in applying these methods on intermittent demand problems.

The lack of a valid and consistent optimisation methodology complicates the application of Croston's method in real problems. The common approach in the literature is to apply a set of different parameters across all time series of a dataset (for some examples see Syntetos and Boylan, 2005; Teunter and Duncan, 2009). However, because of the nature of intermittent data, evaluating which parameter is better is not trivial. Using the conventional time series approach, that is minimising the error of some in-sample fit, is contentious, as the use of conventional errors has been challenged in the literature (Syntetos and Boylan, 2005; Wallström and Segerstedt, 2010). On the other hand, using out-of-sample criteria for identifying which parameter performs better requires availability of adequate sample and may be impractical, if not impossible, for individual time series parameter selection. Additionally using conventional error metrics may be inappropriate, as already mentioned. Nonetheless, in the forecasting literature it is widely accepted that using the optimal model parameters

* Tel.: +44 1524 592911.
E-mail address: n.kourentzes@lancaster.ac.uk

http://dx.doi.org/10.1016/j.ijpe.2014.06.007
0925-5273/

for each time series is best practice, leading to accuracy improvements (Fildes et al., 1998; Hyndman et al., 2002; Gardner, 2006). This paper aims to explore this question: given the limitations of existing error metrics – as optimisation cost functions – for Croston's method and its variants, how should we optimise such models?

In the literature there are some examples of optimised intermittent demand models (Eaves and Kingsman, 2004; Petropoulos et al., 2008; Boylan et al., 2008; Petropoulos et al., 2013; Teunter et al., 2010), however these have not focused on exploring the validity and performance of using conventional time series optimisation for intermittent data. This work explores the performance of different error metrics as cost functions for optimising forecasting models for intermittent demand. Furthermore, different optimisation setup options, such as optimising the initialisation values of the models or not, are evaluated. The conventional error metrics are found lacking and two novel ones are proposed. These are found to produce better results in terms of forecasting and inventory performance and to automatically select parameters with values that are in agreement with the suggestions of the literature.

A closely related topic to parameter selection is model selection, as both are conventionally based on evaluating the quality of the fit of models to time series. Recently there has been a surge in researching alternative intermittent demand forecasting methods, raising the issue of model selection. Assuming that for each method optimal parameters can be identified, an important problem is which model to use. Although there have been some theoretical advancements for selecting between a subset of intermittent demand methods, depending on the time series characteristics (Syntetos et al., 2005; Kostenko and Hyndman, 2006; Heinecke et al., 2013), these cannot provide generic model selection guidance. This paper investigates whether the different error metrics explored here can facilitate and automate model selection for intermittent demand data. The findings are contrasted with the literature, highlighting problems in model selection for intermittent demand time series.

The rest of the paper is structured as follows: Section 2 frames the parameter and model selection problem for intermittent demand drawing from the literature. Section 3 presents the forecasting methods that this paper focuses on. Section 4 discusses the applicability of established metrics as optimisation cost functions and introduces two novel ones that overcome the limitation of existing ones. Section 5 presents the experimental setup and the results of the empirical evaluation, while Section 6 concludes with a short discussion of the findings.

2. Background research

Croston (1972) proposed a forecasting method specific for intermittent demand problems. Croston's method derives a non-zero demand and an inter-demand interval time series from the original intermittent data. These new time series are then smoothed and forecasted independently using single exponential smoothing, employing the same smoothing parameter α for both. Dividing the resulting estimates produces the final forecast, which is used to predict the average future demand per time period. Since then, Croston's method has been widely researched, applied in practice, and attracted some criticism. The theoretical grounding of the method has been questioned by Snyder (2002), Shenstone and Hyndman (2005) and Snyder et al. (2012). Furthermore, the method assumes that the demand size and the inter-demand intervals are independent, something that has been questioned by Willemain et al. (1994). Kourentzes (2013) showed that modelling such dependence explicitly is beneficial. Nonetheless, several studies have verified the good forecasting accuracy and inventory performance of this method (for example see Willemain et al., 1994; Johnston and Boylan, 1996; Willemain et al., 2004), with Shenstone and Hyndman (2005) arguing that these limitations and issues do not "*mean that Croston's method itself is not useful*". The availability of the method in several established forecasting packages, such as ForecastPro, SAS and SAP APO, is indicative of its widespread use in practice by organisations. These factors make it important to explore how to optimally parametrise and use such models.

Syntetos and Boylan (2001) showed that the original Croston's method is biased and proposed a modified version that corrected the problem (Syntetos and Boylan, 2005), demonstrating improved accuracy. Teunter and Sani (2009b) and Wallström and Segerstedt (2010) provide evidence that this modification can still be biased when the intermittency of a series is quite low. Shale et al. (2006) showed that if the orders arrive as a Poisson process then a different modification is more appropriate. Levén and Segerstedt (2004) proposed an alternative modification to Croston's method to avoid the bias of the original method, yet this was shown to be biased in a different manner (Boylan and Syntetos, 2007). Teunter et al. (2011) observed that although the previous work dealt with the bias of Croston's method, it was still unsuitable to deal with obsolescence issues. To address this, they proposed a new method that updates the probability of demand continuously, in contrast to Croston's method that updates its estimates only when non-zero demand is observed. Recently, several other methods have appeared in the literature, such as bootstrapping based methods (Willemain et al., 2004; Porras and Dekker, 2008), neural networks (Gutierrez et al., 2008; Kourentzes, 2013) and aggregation based methods (Nikolopoulos et al., 2010; Babai et al., 2012); however, these have not been widely used yet. The reader is referred to Boylan and Syntetos (2010) and Bacchetti and Saccani (2012) for a more detailed review of the area.

Focusing on Croston's method and its variants, it is evident from the literature that the selection of the parameters is mostly done ad-hoc (for example see Syntetos and Boylan, 2005; Teunter and Duncan, 2009; Wallström and Segerstedt, 2010; Romeijnders et al., 2012). Syntetos and Boylan (2005) argue that optimising the parameters of Croston's method is not straightforward, due to the limited non-zero observations. The short demand history causes the initialisation of the method to be carried forward into the forecasts. In practice intermittent demand datasets are often short. Arguably, this problem may be overcome by optimising both model parameters and initial values, which is a standard practice for exponential smoothing on fast moving items (Hyndman et al., 2002; Ord and Fildes, 2012). Croston (1972) suggested that the parameters should be between 0.1 and 0.3, while Syntetos and Boylan (2005) advised for values between 0.05 and 0.2. Babai et al. (2011) investigated parameters up to 0.3 and found that depending on the levels of intermittency and lumpiness different smoothing parameters achieved minimum bias, thus demonstrating the need to tune the parameters according to the series at hand.

Similarly, the method proposed by Teunter et al. (2011) was not optimised, although they provide guidelines on how to select the parameters. Romeijnders et al. (2012) employed that the same method used ad-hoc parameters, although they provide additional details on how to initialise the method.

Conventional exponential smoothing optimisation is typically done on squared or absolute errors (Gardner, 2006). Wallström and Segerstedt (2010) point out that in the case of intermittent demand, an optimiser will focus on the many zero demand periods, thus biasing the forecast to be lower than the actual demand. They found that mean absolute error was particularly vulnerable to this, something highlighted by Teunter and Duncan (2009) as well. This issue has been investigated in the literature in a related context, when trying to identify an appropriate evaluation metric for intermittent series. Such