Forecasting Products with Intermittent Demand

Jack Gallagher, Robyn Goldsmith
j.gallagher4@lancaster.ac.uk, r.goldsmith@lancaster.ac.uk

Lancaster University

STOR-i

1. Introduction

Intermittent demand is a classification of sales data characterised by having several, sporadic, and high variance periods of demand. This, opposed to sales data with more periodic trends (e.g seasonal) or data with smoother sales patterns, makes the demand of products with intermittent sales difficult to accurately forecast using traditional methods. Examples of such products are; machinery, spare parts and jewellery to name a few. The necessity for the accurate forecasting of these products appertains to the expensive nature of holding stock for prolonged periods and the risk of their obsolescence "exacerbated by the greatly reduced product life cycles in modern industry" [1]. In conjunction with examples like the US military having up to a 60% excess of spare parts in its inventory [2], there is a clear need for methods to help combat these issues.

2. Distributing the data

With the aim to estimate optimum stock levels, we must find distributions that best describe intermittent data. Whilst somewhat limiting, the Poisson distribution is a natural choice that for slower moving items in particular, can capture the right skewed nature of the frequency of demand sizes [3]. Further, more flexible models include;

- Negative Binomial Distribution
- Stuttering Poisson

The flexibility of said models arises because unlike Poisson normally, the variance and mean need not be the same and the sporadic nature of these models can lead to high variance low mean outcomes.

3. Simple Exponential Smoothing

To accurately model our data we must forecast the parameters required in the distributions as they vary over time. We choose to focus on the Poisson Distribution hence look to estimate the mean. A fairly primitive method is simple exponential smoothing(SES) given by the formula;

$$\hat{d}_{t+1} = \hat{d}_t + a(d_t - \hat{d}_t),$$

Where the $\hat{d}$ variables are the forecasted demands and d_t is the true demand in the previous time period [1]. "a" serves as an error constant between 0 and 1 controlling how much our future predictions are impacted by the error in the previous forecasts $(d_t - \hat{d}_t)$. Below is an example of SES forecasts for varying alpha values;

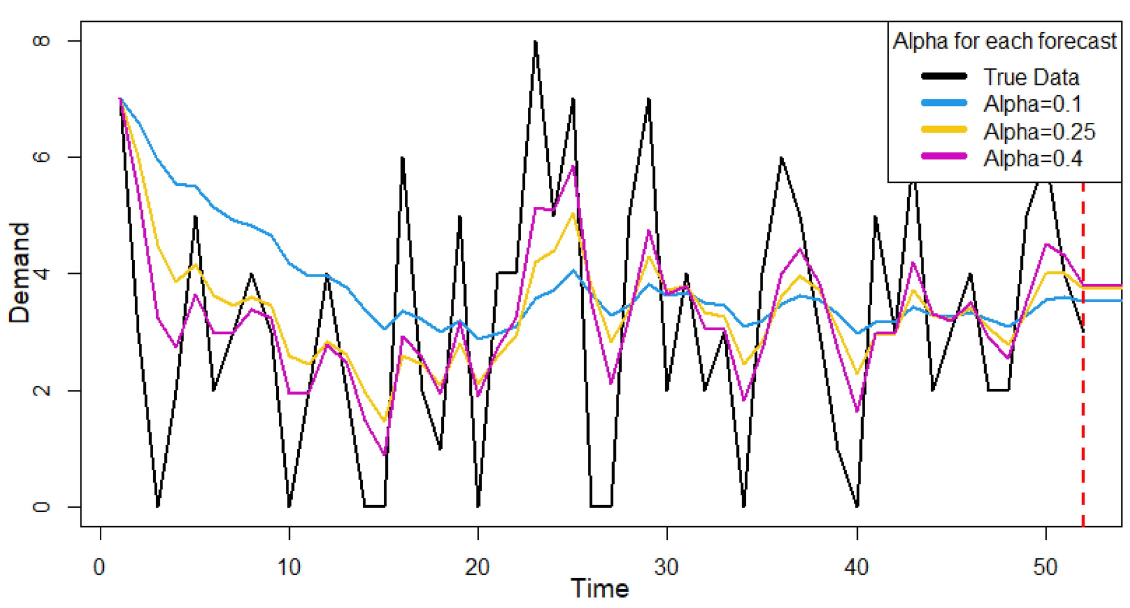

4. Croston's Method

Croston's Method [4] is an improvement upon SES primarily tackling the extremity of the upward bias SES displays. The method forecasts both the mean interval between non-zero demand periods as well as the mean size of non-zero demands in a period. The formula for this method being ;

$$\hat{R}_{t+1} = \hat{R}_t + a(R_t - \hat{R}_t), \qquad \hat{I}_{t+1} = \hat{I}_t + a(I_t - \hat{I}_t) \qquad \hat{D}_{t+1} = \frac{\hat{R}_{t+1}}{\hat{I}_{t+1}}$$

Where the R terms represent the mean demand size and the I terms the mean interval. However whilst Croston's method has less bias than SES it still has an upwards bias over time meaning that as products become obsolete neither of these method capture the right trend. The smoothing constant alpha can be optimised per data set by minimising error measures like the mean squared error [1]. Below you can see the same alpha values as used in SES however we optimised the alpha value by minimising the mean squared error to get alpha $= 0.25$ as our ideal value.

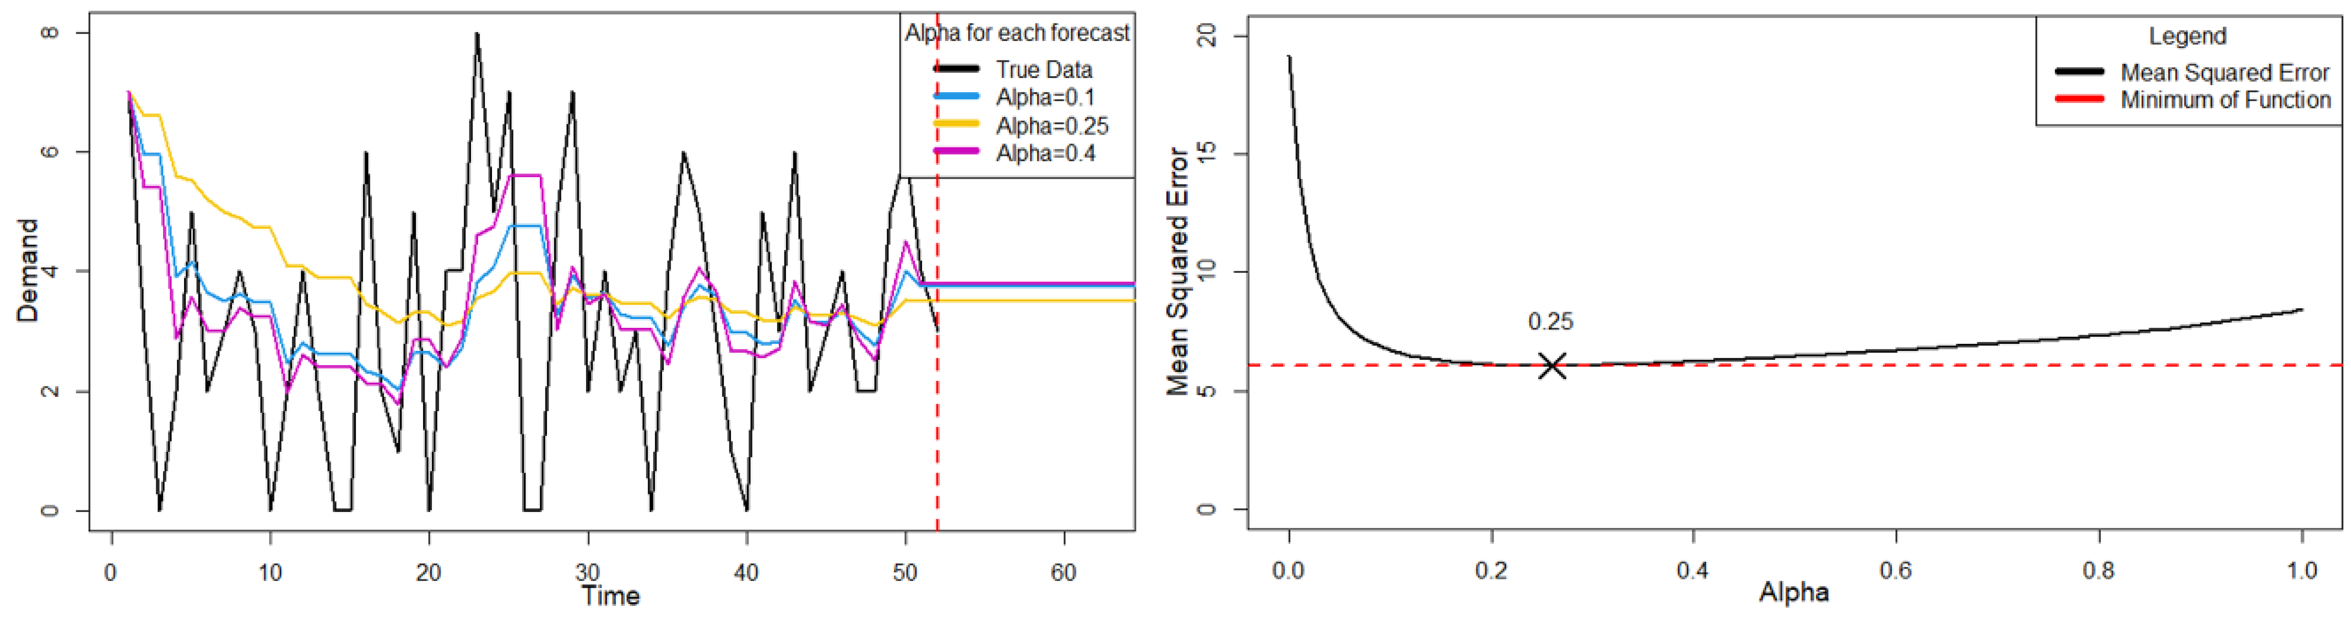

5. Temporal Aggregation

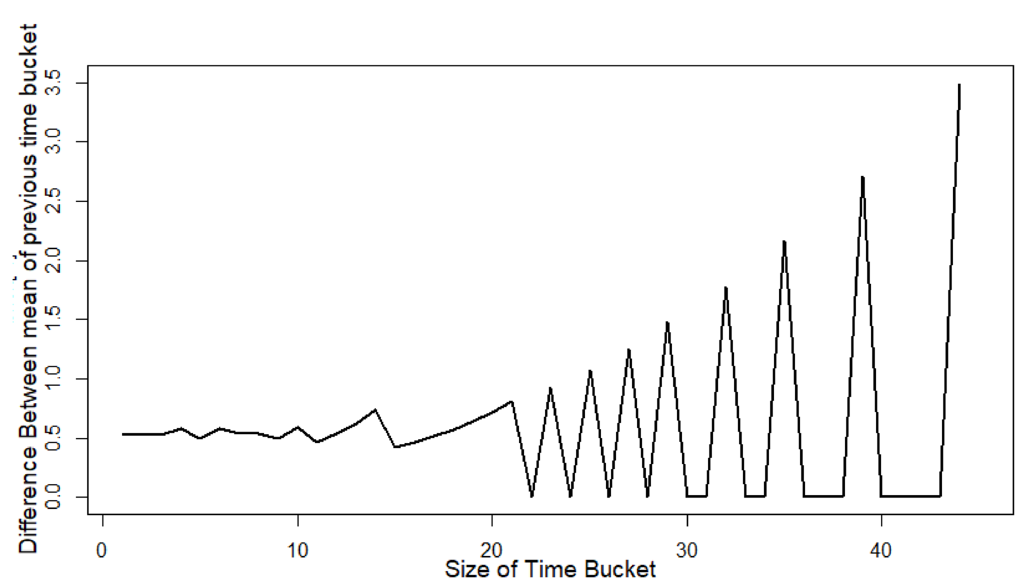

Whilst we see the application of these methods to data with clear intermittent patterns, by manipulating the data we have through temporal aggregation, we can make data seem less or more intermittent. Take for example bread being sold in a store, as a staple item we would expect consistently high sales when looked at across a day, however changing the time bucket (time period) to hours, the data may seem intermittent.
More usefully however, we can group time periods together to reduce the amount of 0 demand periods in our data in the hopes that our forecasting methods improve.
The time bucket we choose is very reliant upon how long it takes for the product to be produced and what stocking strategy you use. In a market where spare parts may be niche and significant in processes that cannot happen without them, temporal aggregation may lead to over estimation of how much is really needed. We may also see the benefit of service being out weighed by the cost of storing such items for much longer periods. So a balance must be struck between the two.
The graph shows the difference in the mean of the data as the time bucket is changed. We see that at first we have consistent growth in the mean of our sample however as this continues, the time bucket changes become less significant and we are losing the amount of data we can input into our methods for longer time buckets.

6. Further Research

1. Exploring the forecasting of variance and looking further into the benefits of the compound Poisson distributions.
2. More emphasise into looking at how we can combat the upward bias of these models to help forecast products becoming obsolete.
3. Make a more specialised model that breaks products down into sections of their life cycle and gives a more specialised model.

7. References

[1] J.E. Boylan and Aris A Syntetos. Intermittent demand forecasting: Context, methods and applications. 2021.

[2] H.L. Hinton Jr. Defense inventory: Continuing challenges in managing inventories and avoiding adverse operational effects. 1999.

[3] F. R. et.al Johnston. Forecasting for items with intermittent demand. *JORS*, 1996.

[4] J. D. Croston. Forecasting and stock control for intermittent demands. *The Journal of the Operational Research Society*, 1972.