Erasmus University Rotterdam
Erasmus School of Economics

Bachelor Thesis: Econometrics and Operations Research

July 4, 2021

Practical comparative analysis for periodic control of intermittent demand items

Author:
Lieke Boogaard
Student Number:
500735

Supervisor:
Oktay Karabag

Abstract

An inventory manager aims to meet demand at all times, when it occurs and as it occurs, but this can be a difficult exercise with intermittent demand items, i.e. items that have a few random demand occurrences and a large proportion non-demand occurrences. This research provides a comparative analysis including various recently introduced Stock Management Policies (SMPs), to see how these policies perform relative to each other in the context of these items, more specifically for intermittent demand items where the lead-time length is shorter than the average inter-demand interval. The empirical performance of the SMPs included is assessed on a large demand data set from the Royal Air Force (RAF, UK). This research eventually recommends inventory managers to use the standard order up-to-level policy in combination with a SBA estimator and a Hurdle Poisson distribution assumption, based on the performance regarding total inventory costs and customer satisfaction. Still, the final choice of an inventory manager should depend on the priorities of the organization.

Contents

1 Introduction

Accurate demand forecasting is a fundamental aspect of supply chain management, therefore the goal of an inventory manager is to regulate Stock Management Policies (SMPs) in such a way that the organization can meet demand at all times. This research will specifically address the under-researched area of forecasting intermittent demand items, for which lead-time length is shorter than the average inter-demand interval. Examples of intermittent demand items are spare parts, duplicates, service parts, or any other Stock Keeping Unit (SKU) which is part of the company's supply chain. These items are quite essential, as they collectively account up for 60% of the total stock value in any industrial sector (Johnston and Boylan (1996); Johnston et al. (2003)).

While the need for efficient SMPs for these items is growing, as a consequence of improvements in logistics which are likely to reduce lead-times in the near future, multiple SMPs have been introduced in academic literature. Accurate demand forecasting for these items, which have an intermittent feature combined with a lumpinous nature, is a difficult exercise, hence some of the introduced policies are not valid when applied to an empirical dataset. This paper will create a broader overview of how different recently introduced SMPs (relatively) relate with each other. This overview will be written both for academic (as comparative analyses on these items are scarce) and practical purposes (as small improvements in inventory management have direct relevance for common business problems). This paper will provide an extensive comparative analysis of multiple SMPs, namely a standard order up-to-level policy in combination with a Syntetos-Boylan Approximation (SBA), the Teunter-Syntetos-Babai method (TSB), and a bootstrapping approach which is called Willemain-Smart-Schwarz method (WSS). Therefore this research is an extension on the useful and prominent comparative analysis given in A. A. Syntetos et al. (2009), where the TSB method was first introduced and compared with standard order up-to-level policy in combination with a SBA estimator. Based on a variety of performance indicators, this paper concluded that the first policy outperformed the second. Since A. A. Syntetos et al. (2009) is very practically relevant, this research decided to fully replicate and validate all the results of the original analysis and expanded the analysis with practical extensions. These extensions are a different distribution assumption for the standard order up-to-level policy in combination with a SBA estimator and adding a new forecasting policy called the WSS method, which was first introduced by Willemain et al. (2004). By adding the first extension to the comparative research, this paper researches if it can improve the performance of the standard order up-to-level policy in combination with a SBA estimator by using a Hurdle Poisson distribution instead of a Negative Binomial distribution. And by adding the second extension, this paper can compare the (relative) performance of more SMPs with each other.

This research uses the same empirical data which was also used in A. A. Syntetos et al. (2009), for practically and academically relevant forecasting evaluations. The utility of this analysis is assessed with the use of large demand data set from the Royal Air Force (RAF, UK). For this purpose, this research aims to answer the following research question; *"Which Stock Management policy (SMP) - a standard* *order up-to-level policy in combination with a SBA estimator, the Teunter-Syntetos-Babai method (TSB) or the Willemain-Smart-Schwarz method (WSS) - performs best for intermittent demand items?"*

Generally, SMPs refer to systems that minimize inventory cost and maximize service level, by using estimators to forecast demand, while keeping inventory constraint in mind (Boylan, Syntetos, et al. (2006)). We subdivided this question into;

- *How the different SMP relate to each other considering total inventory costs and Customer Service Level (CSL: defined as the proportion of demand satisfied directly from stock)?*
- *Is it possible to improve the forecasting evaluations of a standard order up-to-level policy in combination with a SBA estimator, by changing the demand per period distribution assumption from a Negative Binomial distribution to a Hurdle Poisson distribution?*

The answers to questions contribute to the purpose of inventory management, which is ultimately the organization's survival by maximizing CSL, profit, and rate of return on stock items and by minimizing inventory costs. Improvements in the CSL are also beneficial for customers, as fewer people will stand in front of an empty shelf.

Regarding the literature on inventory management, we can conclude that periodic stock control applications have been the focus in most academic work, as they generally reflect to a great extent to real-world practices. Based on the inefficiencies and biases of the existing SMPs observed while using these applications, various SMPs have been introduced to account for these shortcomings. This research compares three recently introduced SMPs. The comparative analysis in this research is also based on these stock control applications. In this way, this research can efficiently show the (relative) performance of the policies included in this research.

The outputs of this study can help inventory managers improve the way they effectively control the inventories they have. With the help of this research, it is possible to choose the most efficient SMP based on a variety of performance indicators. Based on the results found regarding total inventory costs and CSL (usually the two most important aspects considering the perspective of an inventory manager), this research recommends inventory managers to use the standard order up-to-level policy in combination with SBA estimator and a Hurdle Poisson distribution. However, it is always recommended for inventory managers to let their choice depend on the organization's priorities.

In section 2 the main findings of existing research will be discussed. In section 3 this research will examine the dataset and cleaning process used for acquiring the relevant observations. In section 4 the policies that are included in the analysis will be explained in detail, just as the simulation details and the performance evaluations. And in section 5, this research shows the results, or performance evaluations, of every SMP included in the analysis. Section 6 is dedicated to the recommendations which can be obtained from the results. And finally, in section 7, along with the main findings, this research will give some recommendations for future research.

2 Literature

In practice, much inventory management on retail level is done subjectively, both for intermittent demand items and fast-moving items. This research focuses on the first category, specifically on items for which lead-time length is shorter than the average inter-demand interval. Because accurate demand forecasting for these items is especially difficult, given their intermittent demand feature combined with a lumpinous nature.

In the following subsection, first a review will be given on literature of intermittent demand SMPs. Subsequently, the choice of each SMP - the standard order up-to-level policy in combination with a SBA estimator, the TSB method, and the WSS method - will be discussed in the remaining subsections.

2.1 Review of literature on intermittent demand forecasting policies

The analysis of Croston (1972) is considered to be the first influential article on intermittent demand forecasting (Fildes et al. (2008)). He showed that exponential smoothing - a SMP frequently used in stock control systems - almost always produces excessive stock levels for intermittent demand items. The proposed SMP to account for these problems, later referred to as Croston's method, uses stochastic demand arrivals and demand sizes. This SMP was claimed to be unbiased, and Willemain et al. (1994) stated that Croston's method is robustly superior to exponential smoothing and could provide tangible benefits to manufacturers forecasting intermittent demand. However, the unbiasedness was never widely recognized as Sani and Kingsman (1997) provided some empirical evidence that suggested losses in performance for Croston's method compared to other SMPs. Eventually, A. A. Syntetos and Boylan (2001) demonstrated that Croston's method is biased, and since then, various modifications have been proposed and evaluated. Some were more biased than the original approach, while others only worked well under specific constraints. This research decided to compare three recently introduced SMPs with each other, which were all found to outperform Croston's method. In this way, more information becomes available on how recently introduced SMPs compare to each other. The three SMPs used in this paper will be explained in the remaining subsections of this section.

2.2 Standard order up-to-level policy in combination with a SBA estimator

A correction factor that is gaining prominence is the SBA, this factor was designed to correct for the bias in Croston's method and introduced by A. A. Syntetos and Boylan (2005). The SBA estimator applies a deflating factor to Croston's method in which the factor is linear in the smoothing constant used for updating the demand intervals. This correction factor is independently verified to be useful in a number of empirical studies (Gutierrez et al. (2008); Eaves and Kingsman (2004)). Therefore it will be the first SMP we evaluate in this research. Usually, this policy is used with the assumption that demand per period follows a Negative Binomial distribution (A. A. Syntetos et al. (2009); A. A.

Syntetos and Boylan (2006)). This distribution is chosen because it has quite a high flexibility regarding variance to mean ratio, and there is also empirical evidence proving that this distribution assumption is suitable in the context of intermittent demand items (Kwan (1991)). But considering the characteristics of the demand per period distribution, which are a large proportion of zero observations and a lumpinous nature of demand, a Hurdle Poisson distribution may also be suitable. This distribution allows for a separate probability for zero-observations, and in addition, it assumes zero- and non-zero-observations to be two separate processes. These characteristics are likely to be suitable since non-zero-observations are so scarce. Therefore this research added this distribution assumption to the comparative analysis of this paper. By adding this distribution assumption, this research extends the existing literature by researching if this assumption influences the forecasting evaluations of an order up-to-level policy combined with a SBA estimator.

2.3 TSB method

Another prominent SMP is the TSB method, which was first introduced in A. A. Syntetos et al. (2009). The TSB method is a SMP that relies upon the employment of two separate estimates, one for intervals and one for sizes, which are used directly for stock control purposes. This works differently from earlier introduced SMPs, which first focus on parameters for the demand per period distribution, and then feed these results in stock control procedure. By including this SMP in our comparative analysis, this paper researches if a SMP that allows interaction between demand forecasting and inventory control procedures outperforms SMPs that neglect this interaction. Because both demand forecasting and stock control are important given the perspective of an inventory manager, this out-performance is expected. But as most SMPs don't allow for such an interaction, empirical evidence for such a hypothesis is scarce. Therefore this research decided to include this under-researched area in the comparative analysis. The relation between forecasting, inventory rules, and performance indicators is well displayed in a framework introduced by Boylan, Syntetos, et al. (2006), given in Figure 1. This policy is used with the assumption that demand sizes follow a Lognormal distribution for representing sizes, as Willemain et al. (1994) found evidence in real-world data that suggested the use of this assumption. This research doesn't recognize that for discrete demand, the Lognormal distribution is only an approximation. Considering the flexibility of this distribution, and the empirical evidence that exists in its support, this research decided to use the distribution after all, just like in A. A. Syntetos et al. (2009).

Figure 1: Inventory system performance measurement

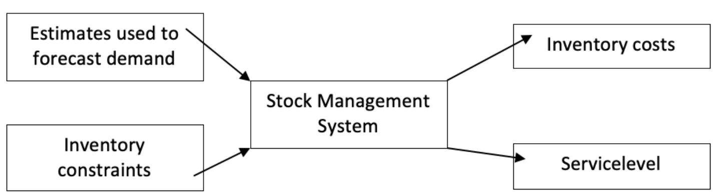

2.4 WSS method

Lastly, this research includes a bootstrap approach called WSS in the comparative research. Both the standard order up-to-level policy in combination with SBA estimator and the TSB method rely on assumptions for distributions. Although sufficient empirical evidence exists for such assumptions, the performance of these SMPs could still be heavily affected by a small skewness in the true distribution. This raises the question, if distribution assumptions contribute to the performance of an SMP. To research this phenomenon, this paper decided to include a 'distribution assumption-free approach' in the comparative research, namely a bootstrapping approach which is called WSS and is introduced by Willemain et al. (2004). This approach is very suitable in the context of this research, as it is responsive to autocorrelation, frequently repeated values, and relatively short series. This approach is also included because it is attractive to inventory managers, as it is different than earlier SMPs discussed in this paper. The previous SMPs discussed in this paper are abstract because they work with statistical inference. The WSS method, on the other hand, has a concrete and algorithmic nature as it is based on computational inference, which makes it an easier method to understand and to implement (Willemain et al. (2004)). A well-known disadvantage of bootstrap approaches, among which WSS, is that there doesn't exist one complete procedure to overcome an underestimation regarding service level performance evaluations (Fricker Jr and Goodhart (2000); Willemain et al. (2004); Zhou and Viswanathan (2011)). There are procedures developed to account for this underestimation for specific SMPs, including WSS method. But still, since Hasni et al. (2019) refers to WSS as a well-performing bootstrapping method to deal with the inventory forecasting of intermittent demand items, this research decided to neglect these procedures.

3 Data

The dataset used in this research consists of the individual monthly demand histories of 5000 Stock Keeping Units (SKUs) over seven years from the Royal Air Force (RAF). It contains 84 monthly demand observations, from January 1996 to December 2002, of items varying from oil filters to dust-caps. For each SKU, along with the historical monthly demand observations, information about the lead-time length and unit cost are also available. Not all SKUs were used in this research, to assess the empirical utility of the stock control approaches proposed in this paper, we excluded all SKUs characterized by an average inter-demand interval less than the lead-time plus one review period. This process resulted in 2455 remaining SKUs that this paper uses in their research, which is similar to the remaining observations in A. A. Syntetos et al. (2009). Because almost 50 % of the SKUs met this (self-defined) condition, we can conclude that these intermittent demand items are indeed of frequent occurrence, which validates our relevance of the focus on these items. Some descriptive statistics related to the lead times and unit cost information are described in Table 1. The first row can be interpreted as follows, the minimum lead-time of a SKU is 0 months, while the mean lead-time across 2455 SKUs is 5.167 months, and the maximum lead-time is 20 months. The variation across the different SKUs is captured in the standard deviation (St. dev.) and equal to 3.975 months, and lastly the median lead-time is 5 months. The second row can be interpreted similarly. After Table 1 is the lead-time length distribution displayed in Figure 2, the x-axis indicates the lead-time length and the y-axis the frequency in which this lead-time occurs.

Table 1: Lead-time and unit cost information

	Mean	Median	St. dev.	Minimum	Maximum
Lead-time (months)	5.167	5.000	3.975	0.000	20.000
Unit costs (£)	62.553	6.000	257.046	0.000	7959.881

Figure 2: Leadtime length distribution

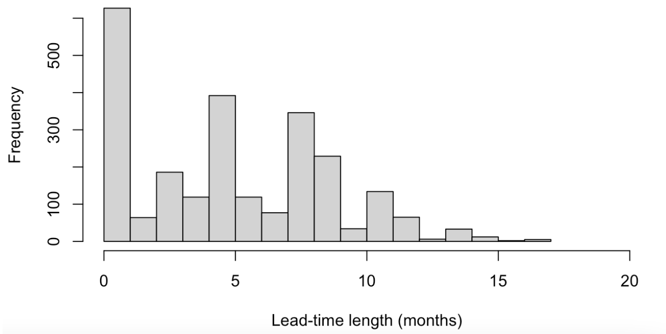

Detailed descriptive statistics on the demand data series characteristics are presented in Table 2. The information on demand sizes can be interpreted as follows, for every SKU, the mean and standard deviation regarding demand size are calculated and ordered. The minimum mean demand size is 1.000 with a corresponding standard deviation equal to 0.000, the maximum mean demand size is 668.000 with a corresponding standard deviation equal to 874.420. The median, or middle observation in the ordering, has a mean demand size of 4.500 with a corresponding standard deviation equal to 3.354. When considering the quantiles of the ordering, we can conclude that the first quantile has a mean demand size equal to 1.750 with a corresponding standard deviation equal to 0.954, and the third quantile has a mean demand size is 13.375 with a corresponding standard deviation equal to 11.166. The other columns, demand intervals and demand per period, can be interpreted similarly.

Table 2: Demand data descriptive statistics

	Demand sizes		Demand intervals		Demand per period	
2455 SKUs	Mean	St. dev	Mean	St. dev	Mean	St. dev
Minimum	1.000	0.000	4.267	0.000	0.036	0.187
1^{st} Quantile	1.750	0.954	8.400	5.990	0.155	0.575
Median	4.500	3.354	10.571	7.616	0.369	1.560
3^{rd} Quantile	13.375	11.166	13.500	9.411	1.173	4.978
Maximum	668.000	874.420	24.000	16.460	65.083	275.706

For the remainder of this paper, this research will use the notation introduced in Table 3. This notation follows the notation used in A. A. Syntetos et al. (2009), but some additions were made regarding the notation used in the extensions.

Table 3: Description of the notation used in this paper

Notation	Description	Notation	Description
L	Lead time (months)	$\Phi_z(.)$	Cumulative distribution function of the demand size
S	Order up-to-level	$\phi_z(.)$	Probability density function of the demand size
t	Time period subscript, t = 1, 2, ..., $\mathcal{T}$	$\Phi_{LTD}(.)$	Cumulative distribution function of the demand during the lead-time and the review period
D_t	Actual demand recorded for an item in period t	$\phi_{LTD}(.)$	Probability density function of the demand during the lead-time and the review period
Z_t	Demand size in period t	u	Unit costs (£)
$\hat{Z}_t$	Estimate of demand size produced at period t for period t + 1	h	Inventory holding charge per unit per period (ie hu: the corresponding cost figure)
T_t	Inter-demand interval recorded in period t (ie a backward looking measurement which records time that has passed since last demand)	b	Backorder charge per unit per period (ie bu: the corresponding cost figure)
$\hat{T}_t$	Estimate of the inter-demand interval produced at period t for period t+1	R	Review period (1 month)
F_t	An estimate of mean demand produced by SBA (at the end of period t for period t + 1)		
α, β & λ	smoothing constants ($0 \leq \alpha, \beta, \lambda \leq 1$)		
X_t	$\begin{cases} 1, & \text{if a demand occurs in period t} \\ 0, & \text{otherwhise} \end{cases}$		
$p_{X_t,X_{t+1}}$	transition probability for the two-state Markov process		

4 Methodology

This research aims to conclude which SMP performs best for intermittent demand items, for which lead-time length is shorter than the average inter-demand interval. Therefore this research considers how different SMPs relate to each other considering total inventory costs and CSL. The SMPs used in this comparative analysis are; a standard order up-to-level policy in combination with a SBA estimator, the TSB method, and the WSS method. The goal of SMP is to estimate the entire distribution of the sum of the demands over the lead time and the review period, called the lead time demand (LTD), defined as; $LTD = \sum_{t=1}^{L+R} Z_t$. With the use of this distribution, the performance of every policy can be evaluated using out-of-sample generations. A comparative analysis for the different SMPs will be done with CSL and total inventory costs for a combination of various h/b charge ratios.

The remainder of this section will be structured as follows, in the first subsection, the different SMPs will be explained and discussed. The second subsection will include some simulation details for the different SMPs included, the inventory control procedures, and performance evaluations.

4.1 Policies for comparative analysis

In order to create a broader overview of how recently introduced methods react with each other, this paper will provide an extensive comparative analysis of the performance evaluations of three SMPs. Namely, a standard order up-to-level policy in combination with a SBA estimator, the TSB method, and finally the WSS method. Below is information given about the theoretical background and procedures of every SMP.

4.1.1 Standard order up-to-level policy in combination with a SBA estimator

Considering different SMPs, the analysis of Croston (1972) has been very influential, especially in the context of intermittent demand forecasting (A. A. Syntetos et al. (2009)). But since A. A. Syntetos and Boylan (2001) demonstrated that Croston's method is biased, various modifications have been proposed and evaluated. The standard order-up-to-level policy is included in this research in combination with a correction factor that is supposed to account for this bias, called the SBA. The SBA estimator was introduced by A. A. Syntetos et al. (2009) and applied a deflating factor to Croston's method in which the factor is linear in the smoothing constant used for updating the demand intervals. Further, this SMP assumes that intermittent demand intervals and demand sizes are both independent and identically distributed (i.i.d.) random variables. And more specifically, this research assumes the demand per period distribution to be Negative Binomial or potentially Hurdle Poisson, as they are both likely to be suitable when dealing with a large proportion of zero-observations. As both of these assumption include the distribution of the demand to be discrete, the optimal order up-to-level policy (S) is given by:

$$\sum_{x=0}^{S-1} \phi_{LTD}(x) \quad \leq \quad \frac{b}{b+h} \quad \leq \quad \sum_{x=0}^{S} \phi_{LTD}(x) \tag{1}$$

In order to calculate this S, one needs to know the parameters of the demand distribution for the demand during lead-time and the review period. For the Negative Binomial distribution, these are estimated as follows;

$$\mu_{L+R} = F_t(L+R) \tag{2}$$

$$\text{and} \qquad \sigma_{L+R} = \sqrt{(L+R)MSE_t} \tag{3}$$

$$\text{where} \qquad F_t = (1-\frac{\alpha}{2})\frac{\hat{Z}_t}{\hat{T}_t} \qquad \text{(the SBA estimator)} \tag{4}$$

$$\text{with} \qquad \hat{T}_t = \hat{T}_{t-1} + \alpha(T_t - \hat{T}_{t-1}) \tag{5}$$

$$\text{and} \qquad \hat{Z}_t = \hat{Z}_{t-1} + \beta(Z_t - \hat{Z}_{t-1}) \tag{6}$$

$$\text{and where} \qquad MSE_t = \lambda(F_{t-1} - D_t)^2 + (1-\lambda)MSE_{t-1} \tag{7}$$

The estimates of demand interval and demand size are both updated at the end of a demand occurring period by means of simple exponential smoothing (SES).

The Hurdle Poisson distribution belongs to a class of 'two-parts distribution,' as it specifies different probability processes for zero and non-zero observations. The probability density function of the Hurdle Poisson distribution is given as follows;

$$\mathbb{P}(Y = y) = \begin{cases} \pi_{L+R}, & \text{for y = 0} \\ (1 - \pi_{L+R})\mu_{L+R}\frac{\mu_{L+R}{}^{y}}{y!}, & \text{for y = 1, 2, 3, ...} \end{cases} \tag{8}$$

A Hurdle model assumes a latent Bernoulli random variable that determines whether a count will be zero or positive. This results in the two described processes: a constant probability for attaining a non-demand occurrence and a separate truncated Poisson random variable for all positive counts. For this distribution, one needs to know two parameters, namely the mean of the underlying Poisson distribution, which is defined as μ_{L+R}, and the Hurdle-crossing probability defined as π_{L+R}. For the calculation of μ_{L+R} we refer to (2), and this parameter is also updated by using SES. In this research, the Hurdle-crossing probability is defined as the probability of having a zero demand observation. Since crossing the threshold of a zero-truncated Poisson distribution can be interpreted as a demand occurrence. Therefore π_{L+R} is equal to the proportion of no demand occurrences. This is defined as follows:

$$\pi_{L+R} = \frac{\text{t} - \text{\# demand occurrences}}{\text{t}} \tag{9}$$

The Hurdle-crossing probability adjusts the probability of no demand occurrence upwards, while it still allows for lumpinous demand occurrences. These features support the fact that the Hurdle Poisson distribution may be very suitable in the context of intermittent demand items.

It should also be noted that the standard order-up-to-level policy was originally not developed for intermittent demand items, as it assumes that more than one demand occurs during lead-time. However, by looking at many demand histories of intermittent demand items, it can be concluded that that is often not the case. Given these limitations, this research also considers an alternative SMP that relies on the direct utilization of the inter-demand interval and demand size estimates in the inventory control derivations, just as A. A. Syntetos et al. (2009). In the next section, such an SMP is discussed in detail.

4.1.2 TSB method

This paper includes a modified version of the standard order up-to-level policy, which we refer to as the TSB method, as it was first introduced in A. A. Syntetos et al. (2009), where it showed auspicious results for the context of intermittent demand items. The goal of A. A. Syntetos et al. (2009) was to find a way to compute the optimal order up-to-level S that minimizes total inventory costs under the constraint that L $\leq E[T_t]$. The optimal order up-to-level S can be calculated by minimizing the expected cost of the described system for S. Therefore, this research assumes that demand sizes follow a Lognormal distribution, because data-supported evidence was found in literature. The system description assumes that after every demand occurrence, the following demand arrives on average after a time equal to the average inter-demand interval just as in A. A. Syntetos et al. (2009). This results in the following calculation of the optimal order up-to-level S, as we assume that demand sizes follow a continuous distribution:

$$S = \Phi_Z^{-1}\left(1 - \frac{h}{(h+b)} \frac{E[T_t]}{(L+R)} \right) \tag{10}$$

The exact derivation of this equation is explained in A. A. Syntetos et al. (2009). While running this policy, the average inter-demand interval ($E[T_t]$) is updated after every demand occurrence while making use of equation (5). The mean and standard deviation of the demand size distribution are similarly updated as in the standard order up-to-level policy.

4.1.3 WSS method

Lastly, this paper includes a bootstrap approach called the WSS method, which is expected to be especially efficient while forecasting intermittent demand items. It is the approach introduced by Willemain et al. (2004), who developed a modified bootstrap responsive to three difficult features of intermittent demand: auto-correlation, frequently repeated values, and relatively short series. A concise summary of the necessary steps is given in Table 4, more detailed information about this algorithm can be found in Willemain et al. (2004). Every step of the algorithm will be discussed below.

Table 4: WSS method

Algorithm:	WSS method
1	Estimate transition probabilities for two-state (zero vs. non-zero) Markov model
2	Conditional on last observed demand, use Markov model to generate a sequence of zero/non-zero values over forecast horizon
3	Replace every nonzero state marker with a numerical value sampled at random with a replacement from the set of observed nonzero demands.
4	Jitter the nonzero demand values.
5	Do inventory control performances.

Since intermittent demand data can have positive and negative auto-correlation, demand forecasts should capture and exploit this auto-correlation. This is done in the first step by using a two-state, first-order Markov process. A history demand series can contain a positive auto-correlation, as literature has concluded that intermittent demand items contain longer sequences of zero or nonzero values than one would expect from a simple Bernoulli process. But sometimes, we also observe more frequent alternation between zero and nonzero values than one would expect, and this means that it is also possible to observe a negative auto-correlation (Willemain et al. (2004)). The two states of the Markov process refer to a demand-occurring state (a period with a positive demand) and a non-demand-occurring state (a period with no demand). Consequently, four transition probabilities need to be calculated for every SKU, as the two states are transient and communicate with each other. These transition probabilities are calculated in a defined in-sample period, in the following way:

$$p_{X_t,X_{t+1}} = \frac{\#\text{ transitions from } X_t \text{ to } X_{t+1}}{\#\text{ total number of transitions beginning from } X_t} \quad \text{for } X_t \in \{0,1\} \text{ and } X_{t+1} \in \{0,1\} \tag{11}$$

Given these transition probabilities, we move to the second step of the algorithm, where we start with a forecast of the sequence of zero and nonzero values over the out-of-sample period. These forecasts are conditional on the state of the period before, X_{t-1}. The third step of the algorithm is to give specific numerical values to the nonzero forecasts. Therefore this research selects one of the historical nonzero demand values at random. We jitter this value (i.e. add some random variation) in step four of the algorithm, in this way, some chance is allowed for a nearby value to be used. Without the jittering process, demand sizes different from those observed in the in-sample period would never appear in the out-of-sample period. This result would give unrealistic bootstrap samples for the out-of-sample period, and especially gives poor estimates of the tails of the LTD distribution when working with short historical demand series (Willemain et al. (2004)). This issue has been recognized by Taylor and Thompson (1992), and is a weakness of the bootstrap approaches when applied to real-world practices. To avoid these issues, Willemain et al. (2004) has developed a jittering process. His jittering process is designed to allow for greater variation around larger demands, and this corresponds with the common empirical phenomenon, which is also observable in Table 2 that variances often increase when means increase. The jittering process goes as follows, let X^* be one of the historical demand values randomly selected, and Z^* be standard normal random deviate, then:

$$\text{Jittered} = 1 + \text{round}\{X^* + Z^*\sqrt{X^*}\} \text{ if Jittered} \leq 0, \text{ then Jittered} = X^* \tag{12}$$

These jittered values are considered the optimal order up-to-level values for S, and will therefore be used for inventory control performances in the last step of the algorithm.

4.2 Simulation details

This research simulates the performance of the various SMPs, and evaluates what happens in terms of stock control. This subsection will be dedicated to the simulation details of this research, which will be presented for the data-cleaning process, every policy included in the comparative analysis, and the inventory control procedure combined with the performance evaluations. This research uses the *'tsintermittent'* package in R while cleaning the data, as it is useful for analysing and forecasting intermittent demand items. The average inter-demand interval used for cleaning the data can be obtained by the *cros.decomp* function in this package. In order to evaluate what happens in terms of stock control procedures with different SMPs, this research splits the available demand history which consists of 84 periods per SKU. The simulation details for every SMP will be explained below, followed by an explanation of the inventory control applications.

4.2.1 Standard order up-to-level policy and the TSB method

As the simulation details for the standard order up-to-level policy and the TSB method are pretty similar we will discuss them together in this paragraph. The following section will be dedicated to the simulation details of the WSS method. Both for the standard order up-to-level policy and the TSB method, we split the available demand history into three parts; one for initialization, one for optimization, and one for evaluation of out-of-sample forecasts. The first part (24 periods) is used for initialization of the average demand size and the average demand interval, both of which can be obtained by the *cros.decomp* function in R. If a SKU doesn't have any demand in their first 24 periods, the initial demand size estimate is set to 1, and the initial inter-demand interval estimate is set to 24. For employing the standard order up-to-level policy in combination with SBA estimator, the initial MSE estimate is initialized at the end of 48^{th} period as the squared error between the initial average demand size and the actual demand in period 48. This is different than the TSB method, where the initial MSE estimate is initialized at the end of 48^{th} period as the squared error between the first estimated demand size and the actual size of the last demand that occurs during the first 48 periods. The second part (24 periods) is used to optimize the smoothing parameters, which are employed for updating purposes. The smoothing constants applied to demand sizes and inter-demand intervals are separately optimized while using MSE as a performance indicator. This research aims to find optimal smoothing constants by minimizing their MSE in the realistic range of [0.05, 0.25], with step increases equal to 0.01. Lastly, the third part (36 periods) is used for the out-of-sample generation of results and performance evaluation. For the standard order up-to-level policy, we need to assume a distribution for demand per unit time period. This research chooses both Negative Binomial and Hurdle Poisson distribution. If a clear percentage (%) can be achieved by employing one distribution over the other, we can conclude that one distribution outperforms the other. For the TSB method, we assume a Lognormal distribution for demand sizes.

4.2.2 WSS method

For the WSS method, we only split the available demand history in two parts, an in-sample period (first 48 periods) and an out-of-sample period (last 36 periods). Initialization and optimization are not necessary, as this is a 'distribution-free' approach. After we've calculated the transition probabilities by using the in-sample data, we forecast the state of the first out-of-sample period conditional on the state of the last in-sample period. For this, we draw a random value from a uniformly distributed distribution by using the *rnorm*-function in R. This function is specifically designed for attaining random variates from the standard normal distribution. Given the state of the last in-sample period, we have two transition probabilities, one corresponding with the probability of changing the state and one corresponding with having no state change. Depending on the exact transition probability, one uses this random variate to depend if the next state changes relative to the previous state. For instance, if the transition probability of changing state is equal to 0.5, this means that any random variate value bigger than 0.5 corresponds with changing the state of the next period, and any random variate value smaller than 0.5 corresponds with a similar state in the next period. This process of forecasting a state of the next period conditional on the last period repeats itself until the last out-of-sample period. Next, we replace every nonzero state in the out-of-sample period with a numerical value sampled at random from the set of observed nonzero demands in the in-sample period and jitter these values. This jittering process is designed to allow greater variation around larger demands.

4.2.3 Inventory control procedure and performance evaluation

The inventory control applications are similar for all the SMPs included in the paper to make a useful comparative analysis. After the initialization and optimization of the relevant parameters of the SMPs, the values of these parameters can be calculated for the out-of-sample generations. This research first calculates the optimal order up-to-level S in period 48. It initializes the net stock to be equal to this optimal order up-to-level in period 48 minus the number of requested purchases in period 48. From period 49 onwards, this research calculates the net stock to be equal to the optimal order up-to-level minus the number of requested purchases in that period plus the number of orders delivered in that period. If the net stock in a certain period is negative, a backlog order is placed, and this backlog order can be added to the net stock at the delivery time of that item. Therefore this research calculates the size of the backlog orders as the (positive) difference between the optimal order up-to-level and the net stock plus the orders that will be delivered within the delivery time of the current order placement. The last important aspect of the comparative analysis is the performance evaluation of the different SMPs, to assess which SMP outperforms the others. This research uses holding volume, backlog volume, total inventory costs consisting of inventory holding cost and backlog costs, the number of orders placed in the out-of-sample period, and CSL. This research didn't incorporate ordering costs, as no relevant information related to ordering charges were provided. Important to note is that some SKUs have a unit cost equal to zero pounds, the unit costs of these items are set 0.01 pounds, to guarantee the calculation of the holding and backlog cost per unit per time period. These are namely defined as a multiplication of the charges by the unit costs.

With these evaluations, we can see which SMP has the nominal costs and which SMP ensures that relatively fewer people will stand for an empty shelf. In this way, this research measures both the accuracy of forecasting as well as the efficiency of the stock control operations, which are both important considering the perspective of an inventory manager. Results are presented for different combinations of inventory holding and backlog charges. In particular, the inventory holding cost is held equal to unity, while the backlog charge is varied, as shown in Table 5. The first combination used in this research is a holding charge equal to 1.00 £, a backlog charge equal to 33.33 £, which results in an *h/b-ratio* of 3%, the second, third, and fourth columns correspond with the remaining combinations used in this research.

Table 5: Holding and backlog charge ratios

Holding Charges	1	1	1	1
Backlog Charges	33.33	20.00	14.28	11.11
h/b ratios	3%	5%	7%	9%

5 Results

The output of our research is displayed in Table 6. The holding volume and backlog volume of a SKU are calculated respectively as the average holding volume across 2455 SKUs in the number of units, which is displayed by *'Holding (volume)'*, and the average negative stock across 36 out-of-sample periods across 2455 SKUs in number of units, which is displayed by *'Backlog (volume)'*. Every time a negative stock occurs, an inventory manager places a backlog order equal to the difference between the order up-to-level and the net stock, *'No. of orders'* displays the amount of placed orders during the 36 out-of-sample periods averaged across 2455 SKUs. The *'Total cost (£)'* represents the total inventory cost averaged for 2455 SKUs, consisting of both the holding and the backlog costs. Lastly, CSL is defined for each SKU as the number of periods during out-of-sample with a positive net stock, divided by the 36 out-of-sample periods, *'CSL'* represents the CSL for each policy averaged across 2455 SKUs. Table 6 displays the results for different combinations of inventory holding and backlog charges for all the SMPs discussed in the methodology section. Important to note is that this research refers to the standard order up-to-level policy in combination with SBA estimator in Table 6 as $St.\ outl_{NBD}$ and $St.\ outl_{hpois}$, for respectively the Negative Binomial distribution assumption and the Hurdle Poisson distribution assumption. For the Hurdle Poisson distribution, the percentage difference for every performance evaluation between this SMP and the SMP that assumes a Negative Binomial distribution is displayed between brackets. A negative percentage difference implies that the Hurdle Poisson distribution causes a reduction for the specific performance evaluation compared to the Negative Binomial distribution. In contrast, a positive percentage difference implies that the Hurdle Poisson distribution causes an increase for the specific performance evaluation compared to the Negative Binomial distribution. Finally, this research indicated the best-performing policy for every performance evaluation (which is the highest value for CSL and the lowest value for all the other performance evaluations), by making this value bold.

Overall, this research observes that the inventory holding volume decreases for most of the SMPs, and the backlog volume increases as the back ordering charge decreases. This is expected because as the b decreases and the h stays the same, it becomes less economical to hold inventories and more economical to backorder. The cost reduction that is visible for all SMPs as b decreases supports this fact, just as the increase in the number of orders. A decrease in CSL for most SMPs as b decreases also follows from this phenomenon as less inventory holding volume means that more people will stand for an empty shelf. For the WSS method, the holding volume, backlog volume, number of orders, and CSL stay the same for every h/b-ratio. This is observed because this SMP is heavily dependent on the usage of simulation. When considering the steps of the algorithm, we can conclude that the calculation of the transition probabilities are independent and thus the same for every h/b-ratio, just as the generated sequence of zero/nonzero values over the forecast horizon. For the third and fourth step in Table 4 in Table 4, one need to use simulation. To have as little variation between the h/b-ratios, this research utilizes the same seed for every ratio. This results in a similar holding volume, backlog volume, number of orders, and CSL but different total costs for every h/b-ratio. Just as the other SMPs included in this comparative analysis, the WSS method displays a pattern of decreasing total costs as the backorder charge decreases.

Before moving to the other findings in this research, we need to consider that this research is written as an extension on the valuable and prominent comparative analysis given in A. A. Syntetos et al.

(2009), because it fully replicates and validates the found results and expanded the research with practical extensions. The results that belong to the replication part of the analysis displayed in Table 6 and Table 7. These results don't match precisely with the results found in A. A. Syntetos et al. (2009). Various reasons can cause the deviations, for instance, the evaluation for free items. This research has set the unit costs of a free item equal to 0.01, otherwhise it is impossible to calculate the holding- and backorder charges. It is unclear if A. A. Syntetos et al. (2009) uses the same or another heuristic for these items. In addition, this research sets the smoothing constant λ equal to 0.25, as no explanation is given in A. A. Syntetos et al. (2009), on how to optimize this parameter. Furthermore, it is important to note that this research uses the equation of Prichard et al. (1965) for the Negative Binomial demand assumption, as it is efficient to implement in R. Still, it is unclear if this is also incorporated in the software packages used by A. A. Syntetos et al. (2009). In addition, it is also unclear how A. A. Syntetos et al. (2009) deals with order placements when there are some orders not yet delivered. This research orders always the (positive) difference between the optimal order up-to-level and the net stock plus the orders that will be delivered within the delivery time of the current order placement, it is unclear if A. A. Syntetos et al. (2009) uses a similar heuristic for inventory control or something different that causes other results. Despite the deviations, the overall results of the replication part correspond with A. A. Syntetos et al. (2009), as the previous paragraph indicates. Therefore the is no reason to reject the results of A. A. Syntetos et al. (2009), and this section continues with the findings of the extensions this paper added.

Table 6: Empirical results

2455 SKUs		*Holding (volume)*	*Backlog (volume)*	*Total cost (£)*	*No. of orders*	*CSL*
$h/b = 3\%$	$St.\ outl_{NBD}$	36.9382	0.4411	6.6971	2.7967	**0.9812**
	$St.\ outl_{hpois}$	12.4930	0.7502	**4.6809**	2.5059	0.9724
		(-66.2%)	(70.1%)	(-30.1%)	(-10.4%)	(-0.9%)
	TSB	33.9914	**0.3853**	5.2488	**2.3988**	0.9805
	WSS	**8.2272**	0.9838	4.8458	3.2900	0.9422
$h/b = 5\%$	$St.\ outl_{NBD}$	27.8068	0.5171	5.0359	3.1259	0.9760
	$St.\ outl_{hpois}$	11.9232	0.7659	3.7505	2.5381	0.9707
		(-57.1%)	(48.2%)	(-25.5%)	(-18.8%)	(-0.5%)
	TSB	27.2247	**0.4547**	4.0830	**2.4224**	**0.9765**
	WSS	**8.2272**	0.9838	**3.2304**	3.2900	0.9422
$h/b = 7\%$	$St.\ outl_{NBD}$	23.0245	0.5723	4.1465	3.3654	0.9719
	$St.\ outl_{hpois}$	11.5425	0.7771	3.2708	2.5165	0.9690
		(-49.9%)	(35.8%)	(-21.1%)	(-25.2%)	(-0.3%)
	TSB	23.2438	**0.5138**	3.4575	**2.4293**	**0.9728**
	WSS	**8.2272**	0.9838	**2.5373**	3.2900	0.9422

$h/b = 9\%$	$St.\ outl_{NBD}$	20.0342	0.6155	3.5974	3.5401	0.9685
	$St.\ outl_{hpois}$	11.2467	0.7862	2.9639	2.5299	0.9679
		(-43.9%)	(27.8%)	(-17.6%)	(-28.5%)	(-0.1%)
	TSB	20.5348	**0.5629**	3.0744	**2.4448**	**0.9696**
	WSS	**8.2272**	0.9838	**2.1531**	3.2900	0.9422

When this research looks specifically to the percentage difference for the standard order up-to-level policies with a Hurdle Poisson distribution assumption and a Negative Binomial distribution assumption, it can be concluded that the Hurdle Poisson distribution leads to significant cost reductions, with only small penalties in CSL. The percentage reduction in total cost varies from approximately 30% (for $h/b = 3\%$) to approximately 18% (for $h/b = 9\%$), while the CSL penalties vary from approximately 0.1% (for $h/b = 9\%$) to approximately 0.9% (for $h/b = 3\%$). These results indicate that the demand per period distribution features are well described in this specific 'two-parts distribution.' It is difficult to estimate the demand per period while using one distribution for all observations because real-world intermittent demand history indicates that demand occurrences are scarce and very varying in size. Therefore, these results indicate that zero- and non-zero-observations should be considered two separate processes, for instance, using a Hurdle Poisson distribution.

Table 7: Percentage difference (%) associated with the TSB method compared to the standard order up-to-level policies in combination with a SBA estimator

		Total cost	CSL
$h/b = 3\%$	$St.\ outl_{NBD}$	-21.6	-0.07
	$St.\ outl_{hpois}$	12.1	0.83
$h/b = 5\%$	$St.\ outl_{NBD}$	-18.9	0.05
	$St.\ outl_{hpois}$	8.9	0.59
$h/b = 7\%$	$St.\ outl_{NBD}$	-16.6	0.09
	$St.\ outl_{hpois}$	5.7	0.39
$h/b = 9\%$	$St.\ outl_{NBD}$	-14.5	0.11
	$St.\ outl_{hpois}$	-3.7	0.17

Another interesting aspect of Table 6 is the performance of the TSB method compared to the standard order up-to-level policy in combination with a SBA estimator. The percentage difference (%) associated with the TSB method compared to the standard order up-to-level policies in combination with a SBA estimator for total costs and CSL are displayed in Table 7. A positive percentage difference implies the TSB method causes an increase for the specific performance evaluation compared to the standard order up-to-level policies in combination with a SBA estimator and a Negative Binomial distribution or a Hurdle Poisson distribution, and a negative percentage difference implies the TSB method causes a decrease for the specific performance evaluation compared to the standard order up-to-level policies in combination with a SBA estimator and a Negative Binomial distribution or a Hurdle Poisson distribution.

Theoretically and based on the results found in A. A. Syntetos et al. (2009), this research expects that the TSB method outperforms the standard order up-to-level policy in combination with a SBA estimator and a Negative Binomial Distribution. And indeed, the employment of the TSB method leads to significant cost reductions compared to the standard order up-to-level policy in combination with a SBA estimator and a Negative Binomial distribution, with almost no penalties in the CSL achieved. When we consider Table 7, the percentage reduction in total cost varies from approximately 15% (for h/b = *9%*) to approximately 22% (for h/b = *3%*), while the CSL penalties vary from approximately 0.07% (for h/b = *3%*) to a reward of approximately 0.11% (for h/b = *9%*). This out-performance is not consistent when comparing the TSB method to the standard order up-to-level policy in combination with a SBA estimator and a Hurdle Poisson distribution. Here the employment of TSB method compared to the standard order up-to-level policy in combination with a SBA estimator and a Hurdle Poisson distribution causes significant cost increases (varying from approximately 4% for h/b = *9%* to approximately 12% for h/b = *3%*) and small rewards for CSL (varying from 0.17% for h/b = *9%* to approximately 0.83% for h/b = *3%*). These results indicate that the level of performance for the standard order up-to-level policy combined with a SBA estimator is heavily dependent on the accuracy of the demand per period distribution assumption. This becomes clear as we compare different demand per period distribution assumptions with each other, but also when we use the TSB method as benchmark. Standard order up-to-level policies first estimate demand, and then they feed their demand estimates into the inventory control procedures. This is different than the TSB method, which relies on the performance of two separate estimates for intervals and sizes, which are used directly in inventory control procedures. An important consequence of this difference is that the TSB method allows for interaction between demand forecasting and inventory control procedures. At the same time, a standard order up-to-level policy neglects this interaction. If a standard order up-to-level policy estimates demand very accurately, the policy outperforms the TSB method. However, this out-performance is not observed for every demand assumption. These findings indicate that considering forecasting and inventory control as two separate domains doesn't always lead to the best performance evaluations. When we consider the perspective of an inventory manager, we conclude that both domains, but also their relation, are very important for meeting demand at all times. Consequently, the performance of an SMP doesn't only depend on the accuracy of demand forecasting or the decisions in inventory control procedures, but also on how these aspects relate with each other. Suppose the interaction between these domains is well described in the SMP, it improves the performance evaluations for an inventory manager significantly compared to a SMP that neglects this interaction.

When we look at the performance of the WSS compared to the other SMPs, we first conclude that the CSL achieved for every h/b-ratio is equal to approximately 0.94, which is significantly lower than the CSL achieved in all the other SMPs in the analysis. Therefore, one should only consider switching to the WSS method and accept these penalties in CSL achieved if significant cost reductions are accomplished. However this is not always the case, for h/b = *5%*, h/b = *7%* and h/b = *9%*, we observe that the employment of the WSS method respectively results in a cost reduction of approximately 13.8%, 22.4% and 27.4% compared to a standard order up-to-level policy in combination with a SBA estimator and a Hurdle Poisson distribution (the policy that produces the second-lowest-cost for every h/b-ratio), but for $h/b = 3\%$ we observe an increase in total cost when we use the WSS method of approximately 3.5% in a similar comparison. This indicates that the WSS method always indicates a decrease in CSL achieved and often (but not in all cases) also a decrease in total costs compared to the other SMPs in the analysis. This underachievement regarding CSL may be caused by underestimating the mean and/or the variance of the lead time demand, this is a well-established issue when analyzing the performance of various parametric and bootstrapping approaches. The other SMPs in this research correct for this underestimation by assuming more appropriate demand distributions for intermittent demand items, and giving consideration for the interaction between demand forecasts and inventory control procedures. But for bootstrap approaches, there doesn't exist one complete procedure to overcome this underestimation.

6 Recommendations

The results explained in the previous section provide a broad overview of how different recently introduced SMPs (relatively) relate with each other. In the first place, this is relevant for academic purposes, as comparative analyses on intermittent demand items, for which lead-time length is shorter than the average inter-demand interval, are scarce. But it is also relevant for practical purposes, as small improvements in inventory management have direct relevance for common business problems. It is good to explain what the learned insights from this analysis mean for organizations for practical purposes. In the first place, these results indicate that a Hurdle Poisson distribution can describe the demand per period distribution for intermittent demand items better than a Negative Binomial distribution. Therefore, it is recommended to use a Hurdle Poisson distribution assumption when we consider a standard order up-to-level policy in combination with a SBA estimator. When the organization finds a SMP that allows for interaction between demand forecasting and inventory control procedures, one can consider using the TSB method. However, it is always recommended to compare this TSB method to various standard order up-to-level policies in combination with a SBA estimator, as it is possible that the latter SMP has a better demand accuracy and therefore better performance evaluations. The TSB method is therefore only recommended if the incorporated interaction between demand forecasting and inventory control procedures improves the performance evaluations compared to a method that neglects this interaction. If an organization wants to avoid abstract SMP based on statistical inference, it is definitely recommended to use a bootstrap approach such as the WSS method. The concrete and algorithmic nature of this method is easy to understand and implement. Although the fact that significant cost reductions often accompany the WSS method, it isn't usually recommended as it has quite a low service level.

Furthermore, it is always recommended for an organization to let their choice of SMP correspond with their priorities and capabilities. An organization which finds service level extremely important may accept some higher costs that can be passed on to customers who value this high CSL. This may not be the case for an organization located in an extremely competitive market and therefore needs their prices and costs to be as low as possible. In addition, an organization should also consider if they have the time and personnel to backlog more frequently and if there are any economic advantages for having less inventory volume.

7 Conclusions and extensions

After the comparative analysis of various Stock Management Policies, the following research question can be answered; *"Which Stock Management Policy performs best for intermittent demand items?"* To answer this question, this research analyzed how the standard order up-to-level policy in combination with a SBA estimator, the TSB method, and the WSS method relate to each other considering a range of performance indicators. And at the same time, this research analyzed if the performance indicators of the standard order up-to-level policy in combination with a SBA estimator could be improved by changing the demand per period distribution assumption. This is done because the need for efficient forecasting policies is growing, especially for items where the lead-time length is shorter than the average inter-demand interval. And at the same time this also done because information on how newly proposed methods relate to each other is scarce. Hence, this research provides a comparative analysis both for academically and practically relevant purposes. The analysis in this research was assessed using a large data set from the RAF, where actual lead-times, unit cost, and demand history was available for 2455 SKUs. Results are presented for different combinations of inventory holding and backlog charges.

This research concludes that it is possible to improve the forecasting evaluations of a standard order up-to-level policy combined with a SBA estimator, by changing the demand per period distribution. This is because the Hurdle Poisson distribution leads to significant cost reductions, with only small penalties in the CSL achieved compared to the Negative Binomial distribution. It can also be concluded that the employment of the TSB method leads to significant cost reductions compared to the standard order up-to-level policy in combination with a SBA estimator and a Negative Binomial distribution, with almost no penalties in the CSL achieved. In addition, the employment of a standard order up-to-level policy in combination with a SBA estimator and a Hurdle Poisson distribution causes significant cost reductions and small CSL penalties compared to the employment of the TSB method. Lastly, the employment of the WSS method always indicates a decrease in CSL achieved and often also a reduction of total costs compared to the other SMPs in this research. These findings may be subject to change when incorporating ordering costs, this research neglected these costs since no relevant information about them was available.

This research eventually recommends inventory managers to use the standard order up-to-level policy in combination with a Hurdle Poisson distribution, as it produces lower cost and only small or no penalties in achieved CSL compared to both the standard order up-to-level policy in combination with a SBA estimator and a Negative Binomial distribution and the TSB method. The WSS method isn't recommended because although it oftentimes leads to significant cost reductions, it is also accompanied by low CSL.

Finally, it is always recommended for inventory managers to let their choice of SMP correspond to the priorities and capabilities of the organization.

The specific consideration of organizational priorities and capabilities can lead to various extension ideas. For instance, this research neglected the ordering costs, as no information for these costs was provided, and it is difficult to come up with some values per SKU as orders are usually placed for multiple SKUs in the same period. But it could be very well possible that the reported benefits or losses would have been more significant if the ordering costs were incorporated. Therefore it would be a good suggestion for further research. It can also be noted that environmental concerns have become more important in every aspect of an organization, even for inventory control management. Less backlog volume can lead to environmental benefits because backlogs are defined as demand which you cannot satisfy immediately from your inventories. Customers who stand in front of an empty shelf need to revisit the store after the lead-time length of their demand. Reducing the number of backlogs will reduce the number of double trips for customers, which can lead to substantial environmental benefits. This applies in particular to stores that are located far away from the customer. But since the dataset in this paper provides no information on this aspect, no methodologies can be applied to this aspect, and it will therefore be a good suggestion for further research. The extension suggestions of this research are not limited to incorporating more performance indicators, such ordering or environmental costs. Since intermittent demand items are so complex and so important nowadays and in the future, many modifications or variations have been developed for the SMP included in this research. For instance, one can consider the performance of standard order up-to-level policies combined with other correction factors, such as the SY introduced in A. Syntetos (2001), or other demand per period distribution assumptions, such as Zero-Inflated Poisson or Compound-Poisson distributions. For the TSB method, it may apply that other assumptions for the demand size distribution better describe the interaction between demand forecasting and inventory control procedures, given the fact that Lognormal is only an approximation of discrete demand. And lastly, one can consider adding an adjusted WSS method, described in Hasni et al. (2019), to account for the observed underestimation. The extension ideas are also not limited to modifications and variations of the SMP that are already included in this research. Many other approaches have also been developed in academic literature. This research only looked at some parametric approaches and a specific bootstrap approach, but other suggestions to extend the comparative analysis given in this research are machine learning approaches such as artificial neural networks (Kourentzes (2013)), nearest neighbors (Nikolopoulos et al. (2016)), fuzzy logic-based techniques (Vasumathi and Saradha (2015)) or some of their adaptations such as ANFIS (Badrzadeh et al. (2018)).

References

Badrzadeh, H., Sarukkalige, R., & Jayawardena, A. (2018). Intermittent stream flow forecasting and modelling with hybrid wavelet neuro-fuzzy model. *Hydrology Research*, *49*(1), 27–40.

Boylan, J. E., Syntetos, A. A. et al. (2006). Accuracy and accuracy-implication metrics for intermittent demand. *Foresight: The International Journal of Applied Forecasting*, *4*, 39–42.

Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Operational Research Quarterly (1970-1977)*, *23*(3), 289–303.

Eaves, A. H., & Kingsman, B. G. (2004). Forecasting for the ordering and stock-holding of spare parts. *Journal of the Operational Research Society*, *55*(4), 431–437.

Fildes, R., Nikolopoulos, K., Crone, S. F., & Syntetos, A. A. (2008). Forecasting and operational research: A review. *Journal of the Operational Research Society*, *59*(9), 1150–1172.

Fricker Jr, R. D., & Goodhart, C. A. (2000). Applying a bootstrap approach for setting reorder points in military supply systems. *Naval Research Logistics (NRL)*, *47*(6), 459–478.

Gutierrez, R. S., Solis, A. O., & Mukhopadhyay, S. (2008). Lumpy demand forecasting using neural networks. *International Journal of Production Economics*, *111*(2), 409–420.

Hasni, M., Aguir, M., Babai, M., & Jemai, Z. (2019). On the performance of adjusted bootstrapping methods for intermittent demand forecasting. *International Journal of Production Economics*, *216*, 145–153.

Johnston, F. R., & Boylan, J. E. (1996). Forecasting for items with intermittent demand. *The Journal of the Operational Research Society*, *47*(1), 113–121.

Johnston, F. R., Boylan, J. E., & Shale, E. A. (2003). An examination of the size of orders from customers, their characterisation and the implications for inventory control of slow moving items. *The Journal of the Operational Research Society*, *54*(8), 833–837.

Kourentzes, N. (2013). Intermittent demand forecasts with neural networks. *International Journal of Production Economics*, *143*(1), 198–206.

Kwan, H. W. (1991). *On the demand distributions of slow-moving items.* (Doctoral dissertation). University of Lancaster.

Nikolopoulos, K. I., Babai, M. Z., & Bozos, K. (2016). Forecasting supply chain sporadic demand with nearest neighbor approaches. *International Journal of Production Economics*, *177*, 139–148.

Prichard, J. W., Eagle, R. H., & Eagle, R. H. (1965). *Modern inventory management.* Wiley.

Sani, B., & Kingsman, B. G. (1997). Selecting the best periodic inventory control and demand forecasting methods for low demand items. *The Journal of the Operational Research Society*, *48*(7), 700–713.

Syntetos, A. A., Babai, M. Z., Dallery, Y., & Teunter, R. (2009). Periodic control of intermittent demand items: Theory and empirical analysis. *The Journal of the Operational Research Society*, *60*(5), 611–618.

Syntetos, A. (2001). *Forecasting of intermittent demand* (Doctoral dissertation). Brunel University Uxbridge.

Syntetos, A. A., & Boylan, J. E. (2001). On the bias of intermittent demand estimates. *International Journal of Production Economics*, *71*(1-3), 457–466.

Syntetos, A. A., & Boylan, J. E. (2005). The accuracy of intermittent demand estimates. *International Journal of Forecasting*, *21*(2), 303–314.

Syntetos, A. A., & Boylan, J. E. (2006). On the stock control performance of intermittent demand estimators. *International Journal of Production Economics*, *103*(1), 36–47.

Taylor, M., & Thompson, J. (1992). A nonparametric density estimation based resampling algorithm. *Exploring the Limits of Bootstrap*, 397–404.

Vasumathi, B., & Saradha, A. (2015). Enhancement of intermittent demands in forecasting for spare parts industry. *Indian Journal of Science and Technology*, *8*(25).

Willemain, T. R., Smart, C. N., & Schwarz, H. F. (2004). A new approach to forecasting intermittent demand for service parts inventories. *International Journal of Forecasting*, *20*(3), 375–387.

Willemain, T. R., Smart, C. N., Shockor, J. H., & DeSautels, P. A. (1994). Forecasting intermittent demand in manufacturing: A comparative evaluation of croston's method. *International Journal of Forecasting*, *10*(4).

Zhou, C., & Viswanathan, S. (2011). Comparison of a new bootstrapping method with parametric approaches for safety stock determination in service parts inventory systems. *International Journal of Production Economics*, *133*(1), 481–485.