Demand Forecasting: Investigating Effects of TSB and Neural Networks in Philips's Demands and an Experiment on the Demand Category's Threshold Change.

By Qinyu ZHANG

Student 641309

Supervisor: Prof. dr. ir. R. Dekker

Second assessor: dr. V. Avagyan

Erasmus School of Economics

Erasmus University Rotterdam

Master Thesis for programme: Economics and Business

Final version: August 2024

Abstract

This thesis completed two tasks: the first one is proposing a new demand threshold to accommodate the need of Philips's future experiments; the second task includes examining five forecasting methods and comparing their performances and effectiveness. The potential ADI threshold for optimizing demand categories is proposed using statistical counts and machine learning methods. Among the forecasting methods, the weighted average shows superior performance in both accuracy and volatility. TSB, a focus of Philips's supply chain planning, reduces volatility but has relatively low-accuracy forecasts compared to the weighted average and single exponential smoothing. The NN model indicates the lowest volatility and relatively high accuracy compared to the TSB and weighted average. Finally, suggestions for Philips are brought up for consideration and further research.

Chapter 1 Introduction

1.1 Introduction

Service spare parts are essential in keeping machines or equipment running, and the management of service spare parts is critical for maintaining efficient operations. As efficient inventory holding can reduce downtime and holding costs, companies are seeking ways of accurately predicting demand. Inventory holding costs can range from 5 to 45 percent of the cost price of the inventory per year, with an often-used average of 25 percent (Durlinger and Paul, 2012). Koninklijke Philips N.V. (hereafter Philips) is a healthcare technology company improving people's health and well-being through meaningful innovation (inner files, Philips Mission, 2024). Philips emphasizes the resilience of its supplier network for critical parts and prioritizes risk mitigation throughout the supply and demand process. Forecasting demands stands out as a crucial step in minimizing risks and enhancing the customer experience.

With more attention paid to forecasting accuracy, Philips successfully improved the accuracy of forecasts. However, the volatility of forecasts has not received as much focus, but it comes back to the spotlight when they find high fluctuations in forecasts. The volatility of forecasts refers to the fluctuations of daily or monthly forecasts monitored by planners in SPS team. Service Part Supply (SPS) team is the demand planning team responsible for meeting the requests for service spare parts across the globe. High volatility will have several inefficiencies, including increased setup times and suboptimal utilization of manufacturing capacities. In addition, the SPS team is confused with the results that came from the system and finds it difficult to provide accurate explanations and maintain stability. At a higher level, such unpredictability can increase inventory holding costs and strain supplier relationships.

Furthermore, the SPS team has raised concerns that the current thresholds for categorizing demands may not be suitable for spare parts that have intermittent patterns. They would like to change the demand threshold ADI (Average Demand Interval) to include more spare parts in erratic and smooth categories. This is due to the recent experiments conducted by them that they applied traditional forecasting methods (i.e., weighted average, single exponential smoothing, etc) and TSB (Teunter, Syntetos and Babai) in all demand categories. But in the future, they will apply TSB only to intermittent demand and apply traditional methods to smooth and erratic demand. Thus, modifying these thresholds can optimize demand classification, which will improve method selection and forecasting results. Therefore, a tailored experiment is proposed and conducted to explore potential threshold adjustments within Philips' data sets.

The main research questions of the research are formulated after discussing with SPS team:

1. *What potential new ADI (Average Demand Interval) can be proposed to accommodate and optimize the demand categorization are Philips?*
2. *Do Neural Networks (NN) and TSB (Teunter, Syntetos and Babai) outperform the existing methods in lumpy pattern under the Philips context?*

Forecasting lumpy demand and intermittent demand is challenging due to the nature of their demand patterns. Therefore, it is vital to adopt suitable forecasting methods and accuracy measurements for these two demand categories. However, within Philips's Supply Chain Management (SPM) system, the lumpy category is not identified, and the methods used for intermittent are used also for lumpy demands. Technically, forecasting methods should be tailored to accommodate the unique features of intermittent and lumpy demand patterns. For instance, Croston's method is more suitable for managing intermittent demand due to its ability to separately forecast the demand size and interval; lumpy demand may benefit from bootstrapping techniques or advanced methods such as machine learning models (Kiefer et al., 2021). Integrating these distinct demand patterns and adopting the same methods for these two categories can introduce significant forecasting volatilities. Thus, in this thesis, I address this issue and seek to determine whether specific forecasting methods designed for lumpy demands are effective for parts categorized as intermittent but exhibiting lumpy demand patterns. Methods from statistics, machine learning and deep learning have been used to predict such demand patterns (Kiefer et al., 2021).

The techniques of forecasting are essential for managing inventory holding and manufacturing plans. With the improvement of artificial intelligence, the application of machine learning models in the supply chain area for Philips can be useful for future demands forecasting. This thesis tests the effect of single-layer Neural Networks (NN). The Teunter, Syntetos, and Babai (TSB) method is another time-series model within the SPM tool, but it has not been considered and utilized so far. However, during the recent experiment conducted by Philips, after allowing the system to choose TSB, 97% of spare parts switched to choosing TSB as best method. In this situation, TSB is another important method to be tested with. This thesis compared the performance of these methods against the existing forecasting techniques used by Philips. The main objective is to determine whether TSB and NN offer superior predictive accuracy in handling Philips' supply chain data. These comparisons are grounded using both the robustness and accuracy of demand forecasting within the organization.

The structure of the thesis is as follows: Chapter 2 introduces backgrounds of the necessary facts for experiments, Chapter 3 reviews the literature, details of research questions and research on an individual spare part are elaborated in Chapter 4. The first research question is explored using traditional models, including TSB on all demand categories. The second research question adds machine learning in, only demands exhibiting lumpy pattern are selected and examined with traditional and NN methods. Steps for processing data sets and the final data sets used for experiments are demonstrated in Chapter 5. Methodologies of the two experiments and the outcomes can be found in Chapter 6 and Chapter 7. Chapter 8 includes the discussion of results, and the suggestions and insights for Philips are in Chapter 9. The supply chain of spare parts can be extremely complex, and in this thesis, the research on forecasting only considers the local demands gathered from locations in each business unit, without parts in reverse logistics or new buy decisions.

1.2 Spare Parts Forecasting and its Relation to Marketing

The SPS team highlights the resilience of its supplier network and the importance of quick service provision. As spare parts and service provision are significant parts of Philips business, delays or insufficiencies in spare parts may lead to a decrease in sales. In addition, the unavailability of spare parts may either result in high emergency costs or increase downtime waiting times for parts. Forecasting plays an important role in quantifying the demands from locations and providing the proper products and services on time. Thus, accurate forecasting has a close relation with the sales and marketing strategies of Philips. This research would assist in promoting the forecasting performance and future business development of Philips.

Chapter 2 Background

In this chapter, the background of the Philips SPS team is briefly introduced. Demand classification and forecasting methods are generally described for better understanding the current situation. Supplementary information such as forecasting procedures in Servigistics and Bestfit is provided. Other metrics Servigistics has but are not considered in this thesis are presented for reference.

2.1 General Background

Philips Service Spare Part department takes on the responsibility of providing the right healthcare equipment's spare parts to patients, clinicians, and engineers. According to the inner files, Global Planning Team Mission: the goal of global planning is to achieve the targeted material availability performance for service parts in support of the customer Fill Rate targets by optimizing inventory levels and minimizing excess and obsolescence costs through the entire lifecycle of a service part. The SPS team (Service Parts Supply Chain team) is dedicated to enhancing customer fill rates and optimizing the service part lifecycle process. The process of planning plays important roles in supplier, warehouse, and customer service (Global Planning Team Mission, 2024). The SPS planning team uses Service Part Management (SPM) tool, named Servigistics, which is a planning tool provided by PTC. PTC (2024) is a third-party computer software and services company that owns the Servigistics product, which is the leading supply chain optimization software. Philips acquired the Servigistics tool and utilized it as the main planning and management tool for service spare parts. The functionalities of Servigistics include autonomous planning, simulating and predicting model uncertainties, increasing service levels through purpose-built, AI-powered optimization capabilities, etc. The SPS Improvement team is part of the SPS team, and its main goals are to support business development and market understanding through adapting new requirements into the Global Planning and Supply processes. During the thesis internship, I worked with the SPS Improvement team to help figure out reasons for forecasting volatility and improvements for future forecasting.

2.2 Demand Classification

Demand classification involves analyzing demand patterns to enhance forecasting accuracy. Prior studies identified demand categorization as a key strategy for improving forecasting effectiveness. There are four categories of demand: erratic, smoothing, intermittent, and lumpy. Erratic demand describes highly unpredictable fluctuations in demand, characterized by irregular surges and drops devoid of discernible trends or patterns. Smoothing demand refers to a relatively stable and constant fluctuation over a given period, and it shows a more predictable and steady flow of demand requests. Intermittent and lumpy, however, exhibit irregularity of customer orders. Intermittent demand often with periods of no demand interspersed with random spikes and lumpy demand refers to a demand pattern with infrequent but significant spikes.

To determine the spare parts' pattern, two coefficients are adopted to define categories: the Squared Coefficient of Variation (CV^2) and the Average Demand Interval (ADI). The ADI-CV demand state space originated from spare parts supply chain research and has been used by the aerospace, steel and retail industries (Nenni et al., 2013; Neu et al., 2024). ADI measures the average number of time periods between two successive requests. CV represents the standard deviation of period requirements divided by the average period. Philips uses squared CV (CV^2) as the threshold measuring the variation of demands. The demands with high variation in interval between two demands but low variation in demand quantity are classified as intermittent demand. The demand with high variation and large quantity is lumpy. Demands with regular quantity and time interval are smooth demands; the erratic demand feature regular occurrences in time with high quantity variations. Philips sets CV^2 equals to 0.49 and ADI equals to 1.32 in Servigistics system for classifying each spare part to categories for further deployment of the model. As mentioned earlier, Philips excludes lumpy demand and integrates it into intermittent, therefore, the demand patterns with ADI higher than 1.32 are treated as intermittent. Figure 2.1 shows the demand category according to Syntetos et al (2005) and the demand category in Philips planning system. After applying the CV^2 and ADI thresholds, 97.6% of spare parts fall into intermittent demand in the system, 1.4% of them are erratic, and 0.9% of them are smooth. Most spare parts show intermittent demand pattern, attributed to the system's policy of excluding lumpy demand. Hence, when the demand category displays "intermittent" in the system, the demand pattern may be intermittent or lumpy.

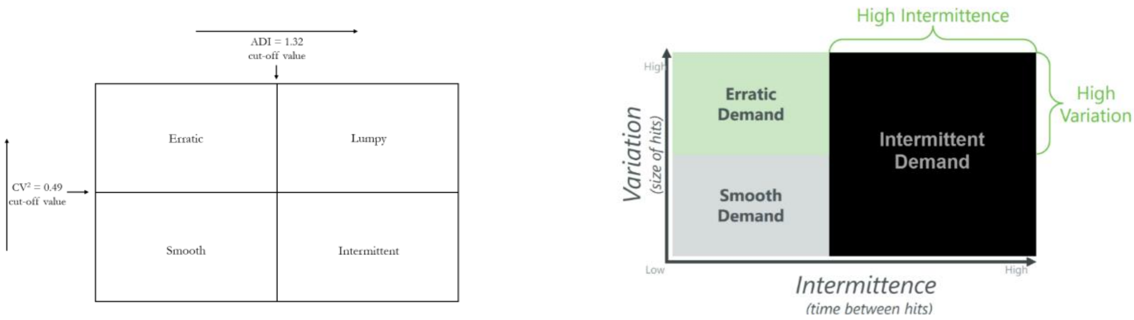

Figure 2.1: Demand category according to Syntetos et al (2005) (left);

Demands into categories in Philips (SPM training, 2023) (right).

2.3 Forecast Methods and Metrics

The main forecast methods in Servigistics are average, weighted average, moving average, single exponential smoothing, double exponential smoothing, intermittence smoothing, and winters multiplicative. In the recent experiment conducted by SPS team, the statistics show that 36.6% of spare parts at location level choose intermittence smoothing, 28.5% of them choose weighted average, 24.9% of them choose single exponential smoothing. These three methods can be considered as primary forecasting methods in the system. The detailed formulas and explanation for the methods are included in the "Methodology" section. Winters multiplicative is applicable to demand patterns that exhibit level, trend, and seasonality. This research excluded this method, as it is under investigation by the SPS team and applies to only a small subset (that is, 0.2%) of spare parts.

Forecast metrics are essential and are derived from the data pulled from the platform. Apart from MAPE, metrics such as tracking signal and bias are also monitored. Tracking signal is used to monitor the accuracy of a forecast by detecting any bias in the forecasting process. It is calculated by comparing the cumulative sum of forecast errors to the mean absolute deviation (MAD) over a given period. Ideally, while ideally the value of Tracking Signal should be zero, a range of 0.5 to 0.5 is used for analytical review (inner file, Philips Forecast version 13, 2024). The primary metric in the system is the tracking signal, which is calculated using the following formula:

$$Tracking\ Signal = \frac{(RSFE/Count)}{MAD}$$

RSFE: running sum of forecast errors.

Count: forecast parameter '# of Slices for Forecast Error Calculation'

MAD: the average of all the forecast errors, disregarding whether the deviations are positive or negative.

Another forecast metric is bias which represents tendency for a forecast to be consistently higher or lower than the demand observations. A forecast bias can be low but with a high error. For example, a forecast which is half of the time 50% higher and 50% lower than the observed demand has no bias; while a forecast which is on average 20% lower than the actual value has 20% error and 20% bias (inner file, Philips Forecast version 13, 2024). For service spare parts, when the demand is intermittent, the bias can be high even though the error may be low. The forecast information provided by the SPS team includes all relevant metrics. But in this thesis, only Mean Absolute Percentage Error (MAPE) is considered for evaluating performance. The formula of MAPE is presented in Methodology chapter.

2.4 Forecasting Procedure with Parameters in Servigistics

The Servigistics's forecasting relies on historical demand data, with method selection primarily determined by calculated MAPE. The number of historical slices represents past actual demand months; and the number of horizon slices which means forward forecasting months. Typically, forecasting is based on 24 or 12 historical slices (including the current forecasting month) and predicts demand for the next 12 months (i.e., the upcoming year). For simplicity, this research uses historical demands from previous 24 time periods (including the current forecasting month) for all spare parts.

Demands are categorized into three types: erratic, smoothing, and intermittent. All built-in methods are evaluated, and the one with the lowest MAPE is considered the final forecasting method. It is important to note that Servigistic TSB has not been implemented in Servigistics, therefore, the system only chooses traditional methods except for Servigistic TSB. Only local demand streams are considered, including all location types (i.e., LDC, DC, FSL). Local Demand Center (LDC) is the geographic area where customer demand for products or services is concentrated. Distribution Center (DC) is a centralized warehouse stores and distributes the products. Forward Stocking Location (FSL) is the warehouse close to the end users. Smoothing parameters like alpha and beta are selected by system algorithm, and the values are around 0.01 to 0.3. Servigistics selects the method with the lowest MAPE, resulting in different methods being chosen for the same spare part across different locations during daily data monitoring. Typically, forecasting for the upcoming year is calculated in each day or month and adjusts when additional demands are incorporated during the forecasting month. The SPS team regards the forecast on the last day of each month as the forecasts of the year in that month.

This thesis mimics the forecasting procedures of Servigistics: I used the past 24 months' historical demands including the current forecasting month. This data was used to make forecasts for the upcoming year (i.e., 12 months), starting from next month. For each spare part in the selected data, the forecasting was conducted each month from January 1, 2023, to December 1, 2023. Finally, I obtained 12 months of forecasting information for each spare part, each month's forecast covering a one-year period.

2.5 Bestfit

Apart from MAPE, Philips set Bestfit when selecting forecasting method and it is auto-approved in Servigistics. Bestfit uses rules to determine whether to replace, eliminate, or keep forecast methods in place (inner file, GPS Forecast Training, 2024). The forecast method that results in the smallest forecast error (MAPE) over the specified period is designated by the application as the Bestfit forecast method (inner file, GPS Forecast Training, 2024). The details of Bestfit rules are listed in Appendix Table 1. In inner GPS Forecast Training file (2024), it introduces Bestfit rules: if Best Fit Analysis encounters a hard rule, the application does not generate a forecast; if Bestfit encounters a soft rule, the application generates a forecast and MAPE, MAD (Mean Absolute Deviation), and RMSE (Root MeanSquare Error); if no rules are encountered, the application generates a normal forecast and forecast errors. This thesis considers only MAPE and disregards the Bestfit effect due to the intransparency of the Servigistics algorithm.

Chapter 3 Literature Review

The literature review of spare parts demand forecasting can be divided into three sections: parametric approaches, non-parametric approaches, and forecast improvement strategies. The framework for this substantive literature review is adapted from Pinçe et al. (2021). A comprehensive review of several time-series methods is in the following sections. Previous studies related to non-parametric techniques are also paid attention to. Past research related to neural networks is also reviewed in this chapter.

3.1 Parametric Approaches

3.1.1 Simple exponential smoothing (SES)

Exponential smoothing is a method for smoothing discrete time series to forecast the immediate future. Although SES is widely used to forecast intermittent demand, the method has important limitations (Syntetos et al., 2015). In the spare parts practice, the SES method can produce high bias, since the algorithm weights recent data more heavily. However, in intermittent demand situation, the demands vary in each period, which leads to biased forecasts. Croston (1972) was the first to notice this, and he noted that the exponential smoothing of intermittent demands almost always produces inappropriate stock levels.

3.1.2 Croston's method and modifications

In an attempt to compensate for problems addressed in the SES method, the forecasts two components of time series start from Croston (Syntetos et al., 2015). Intermittent demand appears at random, with many time periods having no demand. Croston's method relies on separate exponentially smoothed estimates of the interval between consecutive demands and the size of the demands (Zied Babai et al., 2014). If there is no demand during one period, the method will increment the counts of time; the mean and corresponding variance of demand per period is calculated for estimating the intermittent future demands. In this way, both time interval and demand size are forecast individually using SES, resulting in more accurate and smoother estimates. The variance of expected value of demand is lower than that of conventional exponential smoothing, but when demand occurs every time interval, the variance from the Croston method will be identical to that of conventional exponential smoothing (Willemain et al., 1994). Croston's method provides a relatively accurate estimation when several assumptions are held successive demand sizes are identical and independently distributed, intervals and sizes are independent of each other, etc. (Willemain et al., 1994). However, the real situation can be complicated, and studies have shown that positive bias lies in the demand per time unit (Syntetos & Boylan, 2001). In addition, obsolescence becomes an issue as the demands of some items decrease over time, but this is not considered in Croston's method. Modifications are made and demonstrate improvements, such as models built by Syntetos and Boylan (Syntetos & Boylan, 2001), Snyder (2002), Teunter-Syntetos-Babai Method (TSB).

Syntetos and Boylan (2001) found Croston's method was biased in 2001 and proposed an improved method in 2005 based on Croston's version. They found the overestimated forecast demand has a positive correlation to the smoothing factor for the demand interval. The Syntetos Boylan approximation (SBA) added a bias correction term $(1 - \alpha/2)$ with a smoothing constant α; and this bias correction means that SBA provides more accurate results for intermittent demand. By using a dataset from the automotive industry, Syntetos and Boylan (2001) showed that SBA gave more accurate results than Croston. But the study carried out by Pinçe et al. (2021) showed that Croston outperformed SBA in terms of service level.

Another adaptation of Croston's Method is the Teunter-Syntetos-Babai Method (TSB), which constructs two separate variables: the demand of the next period (z) and the demand probability (p) of that period. As mentioned before, the obsolescence issue is important but not considered in Croston's method. Teunter et al. noted that Croston method cannot be used to estimate the risk of obsolescence and deal with the removal of excess/dead stock. TSB proposed a new method in which the estimate of the probability of occurrence is updated every time period. In this way, bias and obsolescence issues can be dealt with by providing up-to-date forecasts even after a long period of zero demand. If there is no demand for a period, the forecast will be adjusted downward. A numerical investigation confirmed that TSB is suitable for situations with both stationary and non-stationary demand (Teunter et al., 2011). However, Zied Babai et al. (2014) showed that the performance of TSB was not considerably better than that of SBA and Croston, sometimes Croston outperforms TSB. They proposed that there was a need for more empirical testing of forecasting methods as the two datasets lead to different and sometimes opposite findings. Babai et al. (2019) then proposed a new method that mixes

SBA and TSB. In periods of positive demand, the method updated demand size and interval in the way of SBA; but at any time, if the actual demand interval became higher than the most recent estimated demand interval, the updated technique follows TSB. The results showed the outperformance of the new forecasting method in many cases dealing with obsolescence.

In this master thesis, the performances of SES, TSB, and one forecasting method named intermittence smoothing are examined and compared. Intermittence smoothing is a combination of SES and Croston's method according to the Philips method definition.

3.2 Non-parametric methods

3.2.1 Bootstrapping and empirical method

The parametric methods discussed so far assume the lead time demand follows a certain probability distribution, but when the demand pattern is not accessible, non-parametric methods including bootstrapping and empirical methods can be useful. Nonparametric methods are more flexible and can be used with any kind of demand distribution (Altay & Litteral, 2011). The classic bootstrapping method has been frequently used in the intermittent demand context. Bookbinder and Lordahl (1989) applied bootstrapping to inventory management context which assumed a standard distribution for re-order points. Willemain et al. (2004) modified the classic bootstrapping method that intermittent demands were better modelled with three difficult features: autocorrelation, frequently repeated values, and relatively short time series. After experimenting in six industrial data sets, they showed that their bootstrapping method generates more accurate results than SES or Croston.

The empirical method as one of the distribution-free methods is also used in previous studies. Porras and Dekker (2008) introduced a new nonparametric method where the empirical lead time demand was used without taking sample, and two optimization approaches (i.e., ex-post, ex-ante) were applied. Van Wingerden et al. (2014) extended the empirical method by incorporating randomness into lead times and improved the previous empirical method which they termed as empirical plus. They found this method outperforms the previous methods when the average inter-demand interval was large and the squared coefficient of variation of the demand size was small.

3.2.2 Lumpy demand and Neural Networks

Accurate forecasting is crucial for supply chain efficiency and inventory management of spare parts, especially for lumpy demand. Recent studies utilized advanced machine learning techniques such as Support Vector Regression (SVR), XGBoost, etc. Hua and Zhang (2006) combined the method with a logistic regression approach in which SVR predicted the occurrences of non-zero demand of spare parts. Study by Sapankevych and Sankar (2009) demonstrated that SVR outperformed deep learning techniques such as Multi-Layer Perceptron (MLP). XGBoost, an eXtreme Gradient Boosting framework performs well for electricity consumption prediction in a study by Deng et al. (2017). This thesis focuses more on one basic machine learning method, single-layer neural network, as the extension of traditional existing methods in Servigistics. Single-hidden layer networks trained by back-propagation suggested possible ways for practitioners to improve implementation in real environments (Lolli et al., 2017). In their research, three different single-hidden layer architectures have been adopted such as feed-forward neural network, time-delay neural network. Hoffmann et al. (2022) compared Artificial Neural Networks (ANN) with traditional forecasting methods based on actual demand data from 29 spare parts of a mechanical engineering company. Their results showed that application of neural networks had a high potential for forecasting irregular demands in terms of MAPE mean consideration.

3.3 Forecast improvement strategies

The forecast improvement is important when conducting the time series mentioned above, especially in the Philips Medical service parts planning process. One key problem detected is increasing forecast variability. High forecast variability leads to uncertainty or volatility in forecasts, which makes decision-making and resource allocation difficult. In this paper, demand classification will be mainly discussed.

Demand classification is about matching the demand characteristics with the appropriate estimation methods to improve forecasts and inventory control. Earlier studies on this topic, such as Williams (1984) categorized demands as sporadic, slow-moving, or smooth by decomposing lead-time demand variance into causal elements. Numerical comparisons show that the proposed demand classification scheme leads to a substantial reduction in inventory costs (Williams, 1984). Except for the leading time as one parameter included in the demand classification, Syntetos et al. suggesting an additional parameter (i.e., the squared coefficient of variation of demand). The cut-off values are first found by comparing MSE using different forecasting methods, and then demands are classified into two dimensions based on these cutoffs for finding the best forecasting methods. The categorization of alternative demand patterns facilitates the selection of a forecasting method (Syntetos et al., 2005). But the research also pointed out that the ultimate purpose of inventory management is to reduce stock holdings and improve customer service levels; thus, forecast accuracy is not the only categorization standard. Additionally, Kostenko and Hyndman (2006) criticized this study by claiming that SBA yields smaller MSEs and proposing a more accurate cutoff based on SBA results.

Chapter 4 Problem Statement

This problem statement elaborates research questions presented in Introduction part in detail. The first question explores the potential modifications for thresholds of demand category to better suit the specific context of Philips. The second question examines the performance of NN (neural networks) and TSB (Teunter, Syntetos and Babai) in Philips's spare parts business environment. Both research questions examined forecasting methods

4.1 Research Question 1: What potential new ADI can be proposed to accommodate and optimize the demand categorization in Philips?

As stated in the Background section, the current thresholds are an ADI of 1.32 and a squared CV of 0.49. Recent experiments revealed that most SKUs switch to Servigistics TSB method upon its adoption in all demand categories. This finding has prompted the SPS improvement team to consider several next steps. They plan to restrict Servigistics algorithm to select Servigistics TSB only for intermittent demand categories and allowing erratic and smooth demand to choose other forecasting methods.

Concerns have arisen from including too many spare parts in the intermittent demand category, which can introduce complexity and increase the time required to implement the Servigistics TSB method. Additionally, for SKUs with erratic and smooth demand, methods such as average, weighted average, and single exponential smoothing are more suitable according to the previous studies. Implementing TSB for these SKUs may not be necessary for these demand categories. While TSB can enhance forecast accuracy, it may also result in increased volatility which is undesirable in forecasts. This highlights the need to re-evaluate and potentially adjust the demand thresholds to ensure appropriate application of forecasting methods, thereby optimizing efficiency and reducing unnecessary complexity. Since the goal of this research question is to recategorize demands, and lumpy pattern is excluded from the system, ADI is the key threshold proposed to be modified.

4.2 Research Question 2: Do NN (Neural Networks) and TSB (Teunter, Syntetos and Babai) outperform the existing methods in the Philips context?

Within Healthcare, Spare Parts Management is directed and managed by the SPS team. For the past work, SPS team carried routine work and forecasting tasks on Servigistics, with the methods and parameters automatically selected by the system algorithm: the method has the lowest MAPE and suitable for the demand classification is chosen as the final forecasting methods and return the corresponding values. BestFit is automatically approved by the system which contains several other criteria for selecting the best method. However, the foundation of BestFit remains MAPE, it uses rules to determine whether to replace, eliminate, or keep forecast methods in place (see details in Bestfit in Chapter 2). Thus, in this thesis, I only consider the MAPE for picking the methods. During recent work, problems are observed by planner that when setting MAPE as the critical standard for selecting method which is supposed to lead to high accuracy, the high volatility also happened in the forecasting. As requested by the managers of the SPS team, they are figuring out ways of reducing the volatility when maintaining the accuracy of forecasting results at the same time.

As outlined in the Background section, a significant 97.6% of demands are categorized as intermittent, yet in practice, they may include lumpy demand integrated into intermittent patterns. Currently, no tailored method exists for intermittent or lumpy demand, as method selection prioritizes error minimization. For instance, in the forecasts for part 459801352572 from each warehouse location, the algorithm primarily selects weighted average and single exponential smoothing. However, this part falls under lumpy demand, for which these methods are generally not well-suited. Though yielding low error rates, these methods exhibit high fluctuation. This example was thoroughly examined in Research on Individual Part in Appendix 1. The SPS improvement team collaborated with the PTC team to enhance forecasting performance, proposing Servigistics TSB as a new method to be adopted in the system. Recent experiments conducted by the PTC Servigistics team in the SBOX environment (Philips' supply chain experimentation laboratory) revealed that transitioning from existing methods to TSB accounted for 97.7% of the total spare parts sample in the experiments. This is a significant change and Philips is closely monitoring the impacts of implementing the TSB method on both forecasting accuracy and volatility.

Based on the above scenarios, this research focuses on lumpy demands categorized within intermittent demand to investigate forecasting methods and their potential to mitigate fluctuations. The development of tailored forecasting methods for intermittent and lumpy demands is important, with a need to enrich and enhance the current methods in the system. Moreover, exploring Servigistics TSB's performance relative to existing methods is also important: assessing its comparative efficacy and its potential to enhance performance for lumpy demand. The effect of NN is explored in Philips's data set, but since the focus of Philips Servigistics still lies in the traditional methods, the results of NN serve as an extension and reference for future experiments.

Chapter 5 Data Description

This chapter introduces the procedure of processing raw data sets and main characteristics of the data sets used for each forecasting method. Since the raw data sets provide limited information about SKUs, the process of preparing them for experiments is crucial and is outlined in Processing Data Sets. Figure 5.1 illustrates the data preparation process. Section Data Sets for Experiments present the final data sets used in the two experiments. All original data sets were provided by the Philips SPS team or extracted directly from Servigistics. Some of the processes of preparing data sets were conducted under the requirements of SPS team. For example, SPS team defines high-value spare parts as parts' unit cost larger than 100 euros.

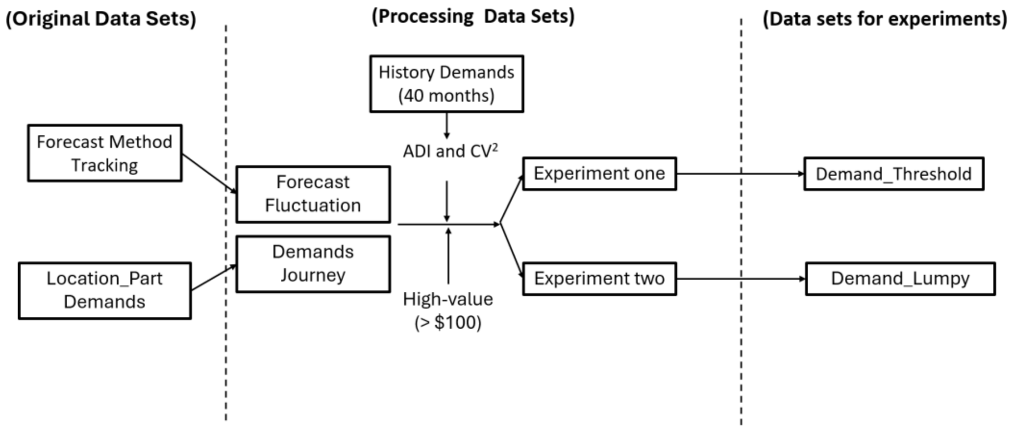

Figure 5.1: Data sets and processing procedures

5.1 Original data sets

In this sub-section, the original data sets provided by Philips Team are elaborated. There are two original data sets: Location_Part Demands data set (hereafter "Location_Part Demands") and Forecast Method Tracking (hereafter "Forecast Method Tracking").

"Location_Part Demands" represents the historical monthly demands of each part of each location. The demands from January 1, 2021, to April 1, 2024, are extracted. The reason why I focus on this period is that all spare parts use either 12 or 24 months, and for forecasting metrics such as standard deviation of demands, Servigistics uses either 24 or 36 months for calculation. Therefore, I included as much information about historical demands as possible (i.e., 3-years demands) so that forecasting bases can be easily adjusted. The demands are rolled up to month buckets, which means demands for that month are aggregated and shown on the first day of that month. In total, there are 2,519,759 observations, 34,749 unique spare parts, and 11 variables in this data set. The research only considered the local demands stream, which is from external demand or customer demands. Additionally, Philips employs a multi-echelon inventory optimization approach, which integrates all distribution levels to optimize inventory and achieve a balance that meets customer demands. As shown in the "Location.Type", historical demand records for LDC (Local Demand Center), DC (Distribution Center), and FSL (Forward Stocking Location) are all included and taken into consideration when conducting forecasting.

The "Forecast Method Tracking" utilizes daily forecasts extracted from Servigistics, and SPS team extracted 32 weekdays' forecasts from February 28, 2024, to April 12, 2024, which serves as a sample period to provide insight into the fluctuation of daily forecasts. The forecasting process uses historical demand data from the preceding 12 months or 24 months, including the current forecasting month. For instance, forecasts in March are derived from historical demands from March and the previous 11 months, with forecasts for a year starting from April (i.e., 12 months). Each day, approximately 191,440 observations are recorded, with each observation representing the forecast from a specific location. Some variables overlap with those found in "Location_Part Demand", such as "Part Number" and "Location Name". The "Forecast.method" is included which indicates the method used for that part in a specific location on the day of extracting the data. It is important to highlight that different locations may have different forecasting methods on different days. This variation depends on the demand patterns of the parts in each location and the system choice of the forecasting. The Servigistics adopts MAPE as the primary standard for choosing the forecasting methods.

5.2 Processing Data Sets

This sub-section details the processing and transformation of the original data. For each research question, the specific data set containing the actual demands is used. As illustrated in Figure 5.1, processed data set Demand_Threshold (hereafter "Demand_Threshold") is for the first research question exploring the potential threshold change, and processed data set Demand_Lumpy (hereafter "Demand_Lumpy") is for the second research question. This section explains four steps of processing data.

Step 1: calculating forecast fluctuation for each Stock Keeping Unit (SKU)

The initial step is identifying the TopX volatile spare parts prior to partitioning them into experimental datasets. "Forecast Method Tracking" is transformed to provide insights into daily forecast fluctuations of each SKU, which is data set Forecast_fluctuation (hereafter "Forecast_fluctuation") in Figure 5.1.

The locations of forecasts are ignored since the SPS team's emphasis is on aggregated spare parts demand. The average values of daily forecasting fluctuations are calculated for further selecting the most fluctuating parts. Daily fluctuation is calculated by taking the difference between the forecasts at time t and the previous period t-1, then dividing it by the forecast at the previous period t-1 (formula 5.1 as below). The "mean.volatility.day" shown in formula 5.2 calulates the average daily forecast fluctuation, where $\bar{DF}$ represents the mean value of daily fluctuations. The "mean.volatility.day" is used for selecting top fluctuating parts. The daily forecast fluctuations for each part on the first day, February 28th, are set as 0. The formulas for calculating fluctuation are as follows. The maximum and minimum of daily forecasting fluctuations are also presented for reference.

$$Daily\ Fluctuation = \frac{Forecast_t - Forecast_{t-1}}{Forecast_{t-1}} \quad (5.1)$$

$$mean.volatility.day = \bar{DF} \quad (5.2)$$

Step 2: Create a historical demand journey for each spare part.

The "Location_Part Demand" contains historical demand records. This thesis ignores the location under instructions from SPS team and aggregates the historical demands from all locations for each SKU. The historical demands are shown on the first day of each month in an aggregated form. To build the forecasting models, I created a demand journey for each SKU which contains spare part number, demand date, and demand quantity on that date. The complete demand journey of each SKU is created by filling in the months without demands with 0, and demand date is from January 1, 2021, to April 1, 2024. The new data set is shown as "Demand Journey" in Figure 5.1.

Step 3: Calculate the squared CV and ADI for each spare part based on the historical demand journey.

For determining the demand categories, squared CVs and ADI were manually calculated in step 3. As mentioned in section 2.2 in Chapter 2 , Philips removed lumpy demand and classified demands with ADI larger than 1.32 to intermittent demand category. In the parameter information provided by Servigistics, I found that around 20% of SKUs are classified into demand categories that differ from those they should be placed in if the Servigistics algorithm adhered to the previously mentioned thresholds (i.e., ADI = 1.32, CV^2 = 0.49). For example, the part in one specific location should be classified into "Erratic Demand" but shows as "Intermittent Demand" in the data which Servigistics provided. I questioned the incorrect categorization of demand to the PTC team, but as of the submission of this thesis, I have not received a response from them. Therefore, I didn't use squared CVs and ADI data provided by Servigistics in this thesis. Instead, these two parameters are calculated using the actual demands of the past 40 months (more than 3 years) from January 1, 2021, to April 1, 2024.

The squared of CV is calculated by squaring the ratio of the standard deviation (SD) to the mean of the demands (formula 5.3 as below). Zero demand is included in the calculation; however, if the demand has been zero for the past 40 months, the spare part is removed from consideration. ADI is calculated as the ratio of the total number of months (i.e., 40 months) to the number of periods with actual demand (i.e., demand quantity larger than 0). The calculated squared CVs and ADI are matched and combined in the "Forecast_fluctuation" using corresponding part numbers of each SKU.

$$CV^2 = \left(\frac{SD\ of\ demands}{mean\ \ of\ demands}\right)^2 \quad (5.3) \qquad ADI = \frac{Total\ number\ of\ months}{Number\ of\ non-zero\ demands\ occasions} \quad (5.4)$$

Step 4: Select demand samples for two experiments: "Demand_Threshold" for first experiment, "Demand_Lumpy" for second experiment.

Philips team suggested focusing on high-value spare parts that have unit cost of more than 100 euros. Thus, for both research questions, only spare parts with values exceeding 100 euros are selected for model building. In total, there are 9,738 SKUs that have a value of more than 100 euros.

For the first experiment, SKUs are sampled from all three categories. Prior to sampling, the dataset consists of 850 erratic spare parts, 6,993 intermittent spare parts, and 1,895 smoothing spare parts. A random sample of SKUs is selected according to the proportions of these three categories, resulting in 300 erratic, 1,500 intermittent, and 500 smoothing spare parts being chosen for analysis. The “Demand_Threshold” is then formed by extracting historical demands journey using part number of the selected SKUs from “Demand Journey”. The “Demand_Threshold” has 92,000 observations and three variables: spare part number, demand date, and demand quantity.

For the second experiment, I used “mean.volatility.day”, the average daily fluctuation of each spare part, to focus on the top 1,000 most volatile spare parts. Threshold parameters are calculated in Step 3. For each SKU, I classified them without removing lumpy demand pattern. The results indicate that among the top 1,000 most volatile spare parts, 995 are lumpy demand, 2 are erratic demand, 1 is smoothing demand, and 2 were removed because demand is zero in most months that ADI and CV^2 cannot be calculated. The “Demand_Lumpy” is formed by extracting demands journey using part number of the selected 995 lumpy SKUs from “Demand Journey”. The “Demand_Lumpy” has 40,000 observations and three variables: spare part number, demand date, and demand quantity.

5.3 Data Sets for Experiments

As indicated in Figure 5.2, two main data sets are used for applying forecasting models: Demand_Threshold and Demand_Lumpy; red line represents ADI=1.32, blue line is CV^2=0.49. “Demand_Threshold” and “Demand_Lumpy” only contain demand journey of each SKU, but the characteristics of these spare parts, such as volatility, ADI, are fetched from the “Forecast_fluctuation” using corresponding spare part number and presented in this section.

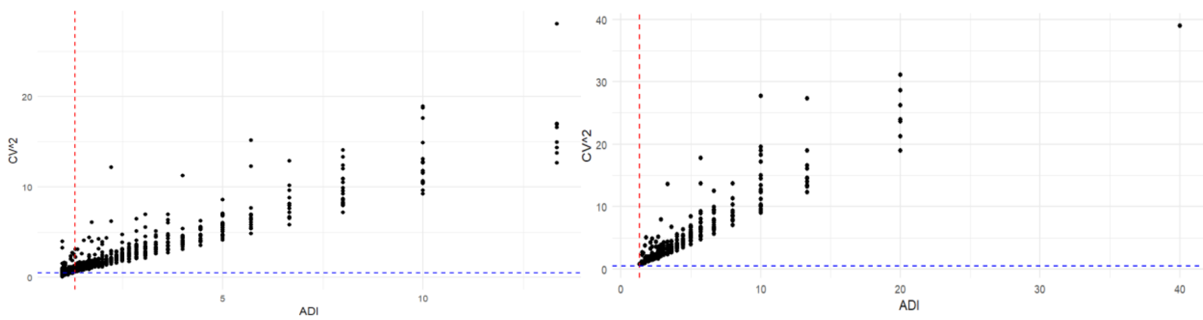

Figure 5.2: distribution of selected SKUs, “Demand_Threshold” (Left), “Demand_Lumpy” (Right).

Figure 5.2 shows the distribution of the selected SKUs in the two data sets according to squared CV and ADI. There are 2,500 SKUs in the "Demand_Threshold" and 995 SKUs in the "Demand_Lumpy". The steps of choosing spare parts are given in section 5.2. A statistical overview of characteristics of the spare parts selected in the two data sets can be found in Table 5.1 and Table 5.2.

Table 5.1 exhibits spare parts randomly selected for the first experiment, which includes three demand categories. In Table 5.1, the average daily fluctuation of spare part demand has a median value of 0.012 (1.2%), and mean value of 0.031 (3.1%). This means that, on average, the daily forecasts would change 3.1% among all selected spare parts, and more than 50% of spare parts have 1.2% change in average volatility of daily forecasts. Among all spare parts, more than half of them have a demand mean of more than 0.6, and 25% of them have a demand mean of more than 2.35. Table 5.2 shows spare parts with lumpy demand patterns for the second experiment, the maximum change of average volatility of daily forecasts is 1.244 (124.4%), and more than 50% of them have at least 0.106 (10.6%) of changes. Since all of them belong to lumpy demand, the minimum ADI is 1.33, which is larger than traditional threshold (i.e., ADI =1.32). The demand mean has a maximum value of 5.675, and more than half of spare parts have a demand mean smaller than 0.2.

Variables	Min	Q1	Median	Mean	Q3	Max	SD
Unit.Cost	100.09	199.282	442.38	1642.044	1268.358	274629.52	5490.535
Mean.volatility.day	0	0.007	0.012	0.031	0.023	4.626	0.086
Min.volatility.day	0	0.009	0.04	0.088	0.147	6.745	0.167
Max.volatility.day	0	0.105	0.198	0.724	0.403	70.562	2.247
ADI	1	1.212	2.5	6.748	6.667	40	9.753
demandmean	0.025	0.175	0.6	3.892	2.35	443.75	13.997
CV^2	0.014	0.692	2.314	6.512	7.179	40	9.987

Table 5.1: statistical overview of SKUs in data set Demand_Threshold.

Variables	Min	Q1	Median	Mean	Q3	Max	SD
Unit.Cost	100.100	180.075	387.390	1342.075	1003.640	57719.480	3630.862
Mean.volatility.day	0.048	0.066	0.106	0.18	0.256	1.244	0.156
Min.volatility.day	0.001	0.089	0.162	0.216	0.204	6.745	0.419
Max.volatility.day	0.446	1.427	2.583	4.9	6.988	37.944	4.852
ADI	1.333	4	6.667	10.079	13.333	40	10.115
demandmean	0.025	0.1	0.2	0.302	0.375	5.675	0.359
CV^2	0.747	3.497	6.347	9.924	12.333	39	10.041

Table 5.2: statistical overview of SKUs in data set Demand_Lumpy.

Data sets Demand_Threshold and Demand_Lumpy contain part number of spare parts, demand date, and demand quantity on that date. Table 5.3 shows the descriptive information of spare parts in three demand categories in Demand_Threshold data set. Variable “Demand” is the average demand over 40 periods, “Mean.V.” is the mean of volatility according to 32 days tracking provided by SPS team. From the table, it can be found that intermittent demand category has more fluctuating SKUs since the max and average of “Mean.V” are much higher than other two categories. In addition, the demands in intermittent category exhibit sporadic patterns: the minimum demand is close to 0 and maximum demand is around 35, but the average demand among all intermittent SKUs are much smaller than other two categories.

Variables	**Erratic**			**Smooth**			**Intermittent**		
	min	max	mean	min	max	mean	min	max	mean
CV^2	0.49	3.98	0.75	0.02	0.49	0.23	0.60	28.03	4.70
ADI	1.00	1.29	1.17	1.00	1.25	1.02	0.60	28.03	4.70
Unit.Cost	100.16	31748.19	1391.15	100.28	68060.70	2473.01	100.16	54922.07	1416.28
Demand	1.23	34.52	3.60	1.43	328.50	15.90	0.07	35.35	0.69
Mean.V.	0.00	0.07	0.01	0.00	0.03	0.01	0.00	4.63	0.04

Table 5.3: descriptive statistics of SKUs in “Demand_Threshold”.

Figure 5.3 to 5.6 shows three examples of distribution of demands over 40 months with ADI = 1.33, ADI = 3.33, ADI = 10, and ADI = 40 in “Demand_Lumpy”. The example SKU is randomly selected to illustrate how demands are distributed in lumpy demand category. The demands follow a sporadic pattern with some periods when there is zero demand.

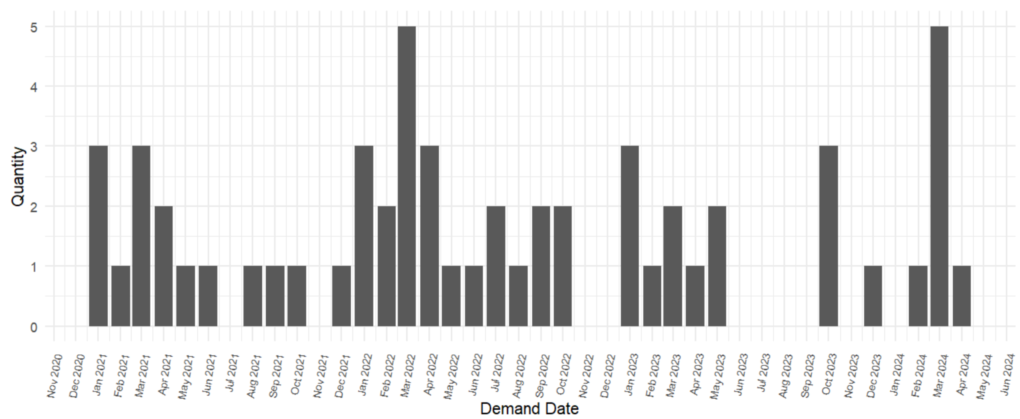

Figure 5.3: demands bar chart of part 459801658621 (ADI =1.33)

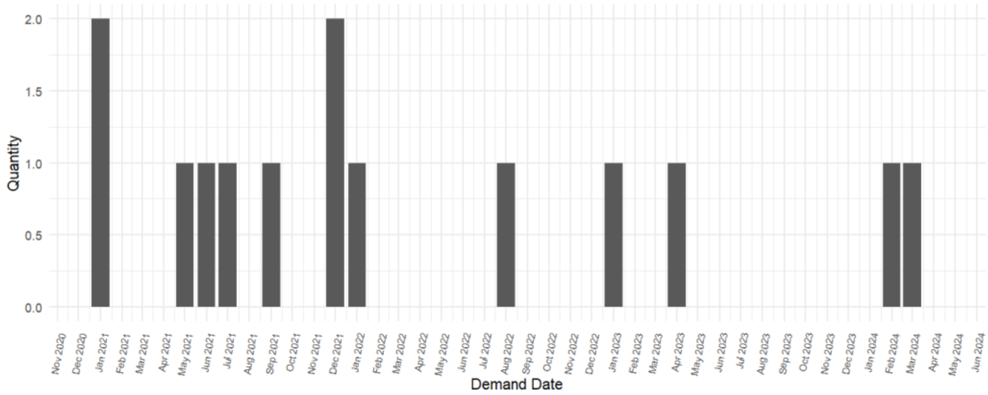

Figure 5.4: demands bar chart of part 451213376211 (ADI =3.33)

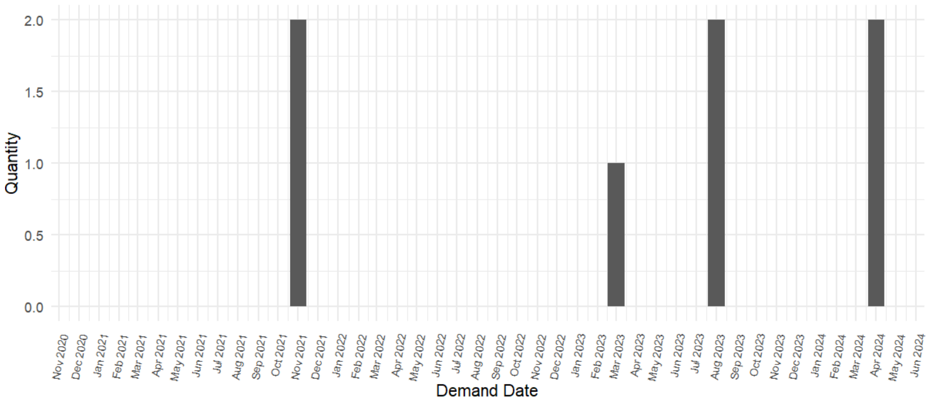

Figure 5.5: demands bar chart of part 459800422471 (ADI =10)

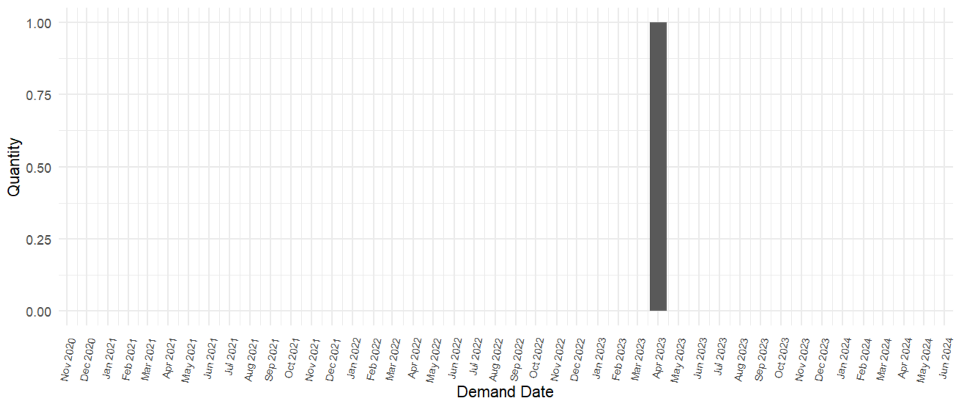

Figure 5.6: demands bar chart of part 453567451632 (ADI = 40)

Chapter 6 Methodology

In this section, the experimental designs for two research questions are illustrated with Figure 6.1 and Figure 6.2. The first set of experiments focuses on ADI threshold modifications corresponding to the first research question. The second experiment examines the performance of the TSB method and NN in relation to the second research question. Additionally, the selected forecasting methods, as mentioned in the Background section, are introduced and described. The metrics and criteria used to measure performance are also elaborated.

6.1 Research Design

6.1.1 Experiment one: Threshold Modifications

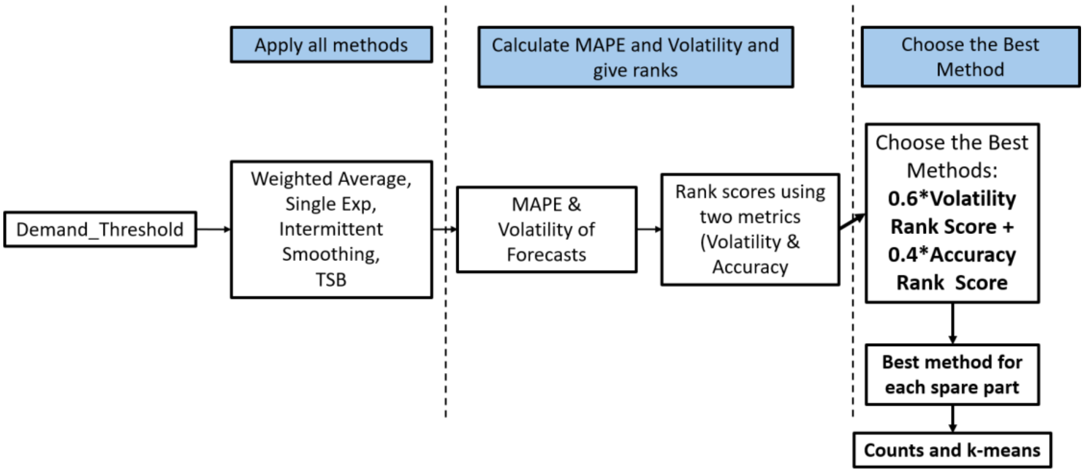

Figure 6.1: the first experiment design.

In the first experiment, SKUs are sampled from erratic, smooth, and intermittent demand categories under definition of Philips for further application of forecasting methods. The methods adopted include Weighted Average (WA), Single Exponential Smoothing (SES), Intermittent Smoothing (IS), and the TSB method. The methodologies of forecasting methods are in section 6.2, the details of implementation of these methods are explained in this section.

To evaluate both accuracy and volatility in selecting the best method for each spare part, I assigned rank scores to methods based on their performance in two dimensions: MAPE and volatility of forecasting. The methods were first implemented in all demand categories, then the MAPE and volatility of forecasting of each spare part were calculated. The details of calculations of accuracy and volatility are in section 6.3. Subsequently, all forecasting methods were ranked separately using MAPE and volatility. The method leading to the lowest MAPE was assigned to a score of 1, the second lowest received a score of 2, and the third received a score of 3. Same as MAPE, the methods were ranked again using volatility: the lowest volatility was assigned to a score of 1, the second lowest received a score of 2, and the third one received a score of 3. I suggested a metric to balance the weights of the two dimensions: a weight number 0.6 was assigned to the volatility rank score and 0.4 was assigned to MAPE rank score. A weighted rank score was then calculated for each forecasting method of each spare part using the following formula 6.1:

$$weighted\ rank\ score = 0.4 \times MAPE\ score + 0.6 \times volatility\ score \qquad (6.1)$$

For example, for one specific spare part, WA gets score of 3 according to accuracy and score of 1 according to volatility, the weighted rank score of WA of that spare part is 1.8. The calculation itinerated until all spare parts got weighted rank score for four forecasting methods. Finally, the method with the lowest weighted rank score was chosen as the best method for that spare part.

To identify the potential thresholds, statistical counts and a simple machine learning algorithm k-means method were used. First, the methods of spare parts which have ADI larger than a certain number were counted. For example, set ADI larger than 2, then the spare parts that have ADI larger than 2 are considered, and final best methods these spare parts choose are counted. The potential ADI change is the switch point when counts from traditional methods (i.e., WA, SES, IS) exceed TSB. As previously mentioned in Literature, Weighted Average (WA), Single Exponential Smoothing (SES), and Intermittent Smoothing (IS) are generally more effective for erratic and smooth demand patterns, whereas the Teunter, Syntetos, and Babai (TSB) method is proven to be more effective for intermittent and lumpy demand patterns. Thus, it is assumed that SKUs selecting WA, SES, and IS are more suitable for classification into erratic demand, while those selecting TSB are more appropriate for intermittent or lumpy demand. The Servigistics system allows SKUs to select all forecasting methods, including TSB. Thus, by analyzing the statistical counts, when the counts of spare parts choosing WA, SES, and IS exceed those of TSB under a certain ADI, it is preferable to classify spare parts with ADIs smaller than this threshold as having erratic or smoothing demand. Conversely, when the counts for the spare parts choosing TSB method exceed those for the other three methods, this indicates that more parts are choosing TSB as the final method. Corresponding ADI can be identified as the switch point and considered as a potential ADI threshold change.

The K-means method is an unsupervised machine learning clustering method that identifies similar groups of data points. The algorithm is easy to implement and only requires that a kd-tree be built once for the data points (Kanungo et al., 2000). The algorithm aims to minimize the distance between points in a cluster with their centroid. K-means has been widely used in a great deal of research from both optimization and data perspectives, and it was examined on data stream and high-dimensional data. The k-means method can give brief overviews of characteristics of each cluster (Wu, 2012). For clustering the spare parts and highlighting the characteristics of them, I made clusters using squared CV, ADI, and final chosen forecasting method of each SKU as the criterion. Silhouette shows which objects lie well within their cluster, and which ones are merely somewhere in between clusters (Rousseeuw, 1987). In this thesis, R package "factoextra" was used for determining the optimal clusters in silhouette plot. Since the k-means clustering is adopted for reference, it is not the focus of the first experiment.

6.1.2 Experiment two: the performance of TSB and NN

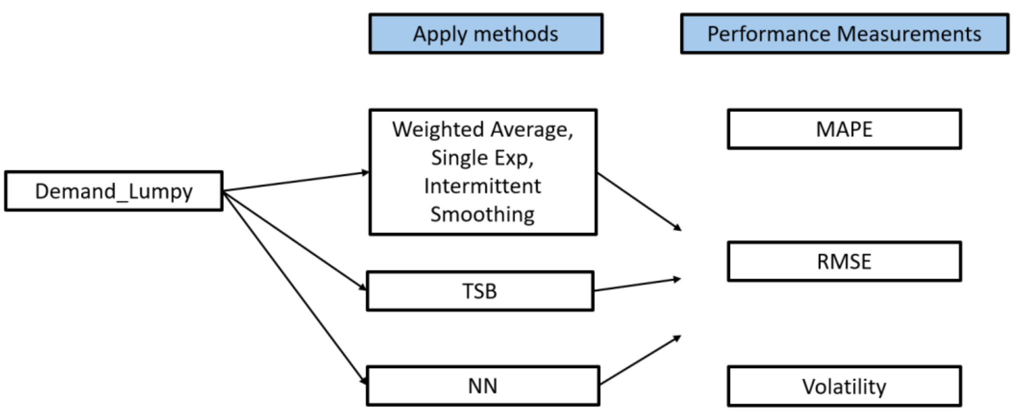

Figure 6.2: the second experiment design.

The second experiment design targets a demand category identified by Philips as intermittent but exhibiting lumpy patterns. In addition to the existing forecasting methods, a machine learning approach using neural networks (NN) is adopted to compare performance with the four traditional methods. The details of forecasting techniques and performance measurements are in section 6.2 and 6.3 of this chapter.

6.2 Forecasting Methods

In this section, five forecasting models applied in this thesis are introduced, along with their detailed implementation parameters. It is important to note that all models are implemented using R packages through their respective function commands. Additionally, smoothing parameters such as alpha are optimized for each method. The forecasting procedures are mentioned in Forecasting Procedure in Chapter 2.

Weighted Average (WA) and Single Exponential Smoothing (SES) models were used directly according to the documentations in SPS team. Intermittence smoothing (IS) is described in PTC website (2024) as "this method generates a modified Croston forecast, and when demand is not intermittent, this method behaves much like Single Exponential Smoothing". Due to the intransparency of Servigistics's models of IS, I assume that Croston's method is used when encountering intermittent demand pattern, and for erratic and smooth pattern, SES is used. Servigistics TSB is interpreted as the Teunter, Syntetos & Babai (TSB) method because of invisible Servigistics's algorithm.

Weighted Average

The Weighted Average (WA) method refers to the average values scaled by their importance. The formula is provided below (formula 6.2). Since forecasts are made for the entire year, monthly forecasts are multiplied by 12 to represent the annual forecast. The weights assigned to each period correspond to the period number, meaning the demand in the first month is multiplied by 1, the demand in the second month by 2, and so on according to Philips training slides (SPS training, 2024). This approach ensures that more recent demand data has a greater influence on the forecast, reflecting its increased relevance.

$$WA_t = \frac{(D_{t-24} \times 1) + (D_{t-23} \times 2) + \cdots + (D_t \times 24)}{(1 + 2 + 3 + \cdots + 24)} \tag{6.2}$$

D_t: historical demand in that forecasting month t.

Croston's method and SES

Croston's method is elaborated in detail in Croston (1972) and is built on the Simple Exponential Smoothing method (SES). In this thesis, the "tsintermittent" R-package by Kourentzes (2014) was used. Croston (1972) focuses on two separate components: the non-zero demand size z_t and inter-demand interval p_t. The prediction from Croston's method is given by: $\hat{y_t} = \frac{\hat{z}_t}{\hat{p}_t}$. The initial value for predictions and both z_t and p_t are using SES, the final output from Croston is the average estimated demand for each period in the forecasting horizon Kourentzes (2014). z_t has to be non-zero in at least two periods because the predictions are updated only when demand occurs. For optimizing the parameters, α values equal to 0.05, 0.1, 0.15, and 0.2 are tried, and the parameter is for both demand and intervals. The sum of forecasts is compared to actual demand in past 12 months and the α which results in least forecasting error MSE is selected as best α. Sum of 12 months' forecasts (that is, the forecasting for the upcoming year) from best α are considered the final forecasts in that forecasting month.

The SES method uses a single smoothing parameter alpha (α) to control the effect of past observations. This smoothing parameter is usually set somewhere between 0.1 and 0.3 in a setting with intermittent demand (Syntetos and Boylan, 2005). The theoretical equation shown below is a simple formular to calculate the forecast for the current period using previous values (both the actual and forecasting values). Same as Croston's method, α values equal to 0.05, 0.1, 0.15, and 0.2 are tried and forecast horizon is set to 12; the results with best performing α will be used as final forecast in that month.

$$\hat{y_t} = \alpha y_t + (1 - \alpha)\hat{y}_{t-1} \tag{6.3}$$

$\hat{y_t}$: forecast for current forecasting month t.

$\hat{y}_{t-1}$: forecast for previous forecasting month t-1.

TSB method

The TSB model comprises two main components: the probability of demand occurrence and the interval between positive demands. Unlike Croston's method, TSB does not alter the level estimation and uses d_t as the probability of the demand occurrence. d_t is 1 when demand does occur and otherwise it is 0. For optimizing the parameters, α values equal to 0.05, 0.1, 0.15, and 0.2 are tried, and the parameter is for both probability and demand. The "tsintermittent" R package is utilized, with the cost function set to mean squared rate (MSR), which is most suitable for TSB, yielding more reasonable forecasts. The forecast horizon is set to 12 periods, and the results are summed to represent the annual forecast. The predictions by TSB are given by the calculation formula:

$$\widehat{y_t} = \hat{d}_t \hat{z}_t \tag{6.4}$$

Neural Network (NN)

NN was found to generally perform better than the traditional methods (Gutierrez et al., 2008). To align with the research goal and offer a potential approach for Philips future experiments, this thesis adopted a single-hidden layer neural network to give a glimpse of the effect of NN. The functions in "nnfor" R package are used. Feedforward neural networks are applied for time series forecasts. The algorithm follows a single-layered network with N hidden neurons and activation function G as follows, where input weight vector is $w_i = (w_{i1}, w_{i2}, \dots, w_{in})^T$:

$$G_N(x_t) = \sum_{i=1}^{N} \beta_i g(w_i \, x_t + b_i) \; with \; t = 1, \dots., T \tag{6.5}$$

Given a set of T samples, {(x_t,d_t)t = 1, …, T}, x_t represents input vector and d_t is the target vector for the supervised learning. The weight vector connects the hidden nodes to the output neurons and b_i is the bias. Figure 6.3 illustrates the algorithm of feedforward neural networks.

In this thesis, the 18 months of historical demands were used to train NN model, and the rest of 6 months were considered as test set. For measuring the accuracy of model, Mean Absolute Error (MAE) within "accuracy" function is used. The mean value of MAE is used to assess the accuracy of the NN model for all spare parts. The hidden nodes are set to number of 3 as it is a reasonably low number to approximate any complex function; maximum number of iterations for training is set to 150. No seasonal lags are considered, and regularization parameter “decay” is 0.8. In this thesis, high “decay” is chosen for preventing overfiting, other values of parameters can be tried in future experiments.

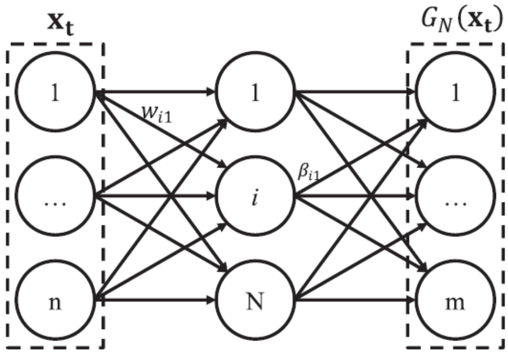

Figure 6.3: the feedforward neural network.

6.3 Selected Performance Measurements

To assess the forecasting and model performance, accuracy and robustness of forecasts are measured. The accuracy metrics give insights into the deviation between forecasting amounts and the actual values of demands; robustness in this thesis refers to the volatility of forecasts. For accuracy metrics, MAPE, MAD, and RMSE are tested; for volatility, the standard deviation (SD) of forecasts is utilized.

The actual demands are calculated under Formula 6.6. I balanced the past first 12 months and last 12 month of past 24 months and assigned a weight number 0.3 to the first 12 months, 0.7 to the last 12 month. The “last 12 months” refers to the 12 months preceding the current forecasting month, including the current month itself. The “first 12 months” refers to the 12 months preceding the “last 12 months”. y_t is the actual demands at time period t.

$$weighted\ actual\ demand = \sum_{i=12}^{23} 0.3 \times y_{t-i} + \sum_{j=0}^{11} 0.7 \times y_{t-j} \qquad (6.6)$$

In Philips’s Servigistics, they mainly use the mean absolute percentage error (MAPE) as the accuracy measurement. MAPE is the average of the absolute differences between actual and forecasted values y_t, expressed as a percentage format. Formula 6.7 is the traditional MAPE (i.e., $MAPE_1$). Given that lumpy demand includes many periods with zero demand, the traditional form of MAPE may fail. Therefore, I adopted an adjustment for MAPE which is given by Formula 6.8 below (Gilliland, 2002). $\hat{y}_t$ is the forecasting demands at time period t.

$$MAPE_1 = \frac{1}{n}\sum\nolimits_{t=1}^{n} |\frac{y_t - \hat{y}_t}{y_t}| \quad (6.7) \qquad MAPE_2 = \frac{\sum_{t=1}^{n} |y_t - \hat{y}_t|}{\sum_{t=1}^{n} y_t} \quad (6.8)$$

The Mean Absolute Deviation (MAD) takes the average of the absolute differences between the forecasted and actual values. MAPE is built under the MAD and actual demands y_t, this metric is calculated as a reference.

$$MAD = \frac{1}{n}\sum_{t=1}^{n} |y_t - \hat{y}_t| \quad (6.9)$$

Pinçe et al. (2021) also find that the most common forecasting accuracy measures used in recent spare parts demand literature are the absolute accuracy measures. The Root Mean Square Error (RMSE) measures the average magnitude of the errors between forecasts and actual values. The metric formula is presented in Formula 6.10 as below:

$$RMSE = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(y_t - \hat{y}_t)^2} \quad (6.10)$$

For measuring the robustness of model performance, volatility of forecasts is used. Standard deviation of forecasts serves as the main measure of volatility (Formula 6.11). $\bar{\hat{y}}$ is the mean of forecasts, and the calculation of SD takes mean of deviation between the mean of forecasts and the forecast at time t. n is the number of periods, and in this thesis, the total forecasting periods are 12 months.

$$\sigma_i = \sqrt{\frac{1}{n-1}\sum_{t=1}^{n}\left(\hat{y}_t - \bar{\hat{y}}\right)^2} \quad (6.11)$$

Chapter 7 Results

This chapter presents results from the experiments for two research questions. The potential ADI thresholds are included in in section 7.1. Forecasting performances of each method are demonstrated in section 7.2. Forecastings are based on the best performance parameters, details of parameters chosing are in section 6.2.

7.1 Experiment one: Threshold Experiment

The first experiment proposed potential ADI thresholds to change in Servigistics. Table 7.1 presents several examples for explaining the outputs of weighted rank score of each forecasting methods for each spare part. The column “Category” refers to demand category, “WA.W” indicates the weighted rank score of the weighted average method, “SES.W” is the weighted rank score of single exponential smoothing, “IS.W” is the weighted rank score of intermittence smoothing, “TSB.W” is the weighted rank score of TSB. “Final.Method” refers to the best method the spare part chose based on weighted rank. The calculation formula of weighted rank score is in Formular 6.1, the method that results in the lowest rank score is chosen as the final method. Examples are randomly selected from three demand categories.

Part.Number	ADI	Category	WA.W	SES.W	IS.W	TSB.W	Final.Method
459800931362	1.026	E	1	2	3	5	WA
453566489171	1.111	E	1.6	2.6	4.2	2.6	WA
453567552651	1.111	E	1	3	5	2	WA
452215021961	1.081	S	1	2.6	4.2	3.2	WA
453567306162	1.081	S	1	2.4	3.8	3.8	WA
459800336413	1.143	S	3.4	2	3	2.6	SES
459801685051	3.636	I	1	4.2	2	3.8	WA
452215031292	4.444	I	2.4	5	2.2	1.4	TSB
453560068371	10.000	I	2.4	3.4	3.8	1.4	TSB

Table 7.1: examples for weighted rank scores of spare parts.

Table 7.2 indicates the first evidence of new ADI: counts of spare parts under different forecasting methods when changing ADI thresholds. As shown in Table 7.1, each spare part has the final method based on the calculated rank scores. For deciding a potential new threshold, ADI values were tried, the spare parts which have ADI larger or equal to that ADI were considered for statistical counting. For example, when ADI is larger than or equal to 1.5, 733 spare parts choose WA, SES, and IS as final method, while 631 spare parts choose TSB as final method.

As explained in the Methodology section 6.1.1, the switch point is when the count of spare parts using TSB exceeds counts of the other methods chosen by the rest of spare parts. This may indicate that TSB is becoming more dominant among SKUs, as an increasing number of spare parts are being selected for this final method. “ADI” refers to the ADI values tried for finding potential new threshold. In the “ExceptTSB” column, it is the sum of spare parts adopting weighted average, single exponential smoothing, and intermittence smoothing; the “TSB” column represents the number of parts that pick TSB as the final method.

ADI	WA	SES	IS	TSB	ExceptTSB	TSB
>=1.3	746	34	52	663	832	663
>=1.5	651	34	48	631	733	631
>=1.9	542	30	42	567	614	567
>=2	519	28	41	549	588	549
>=2.5	432	26	35	486	493	486
>=3	355	25	31	426	411	426
>=3.5	308	20	25	381	353	381
>=4	283	17	24	363	324	363

Table 7.2: statistical counts of spare parts under certain forecasting methods.

From Table 7.2, there is a switch point when ADI equals or is larger than 3. When ADI is smaller than 3, the counts of parts using TSB are fewer than counts of parts using WA, SES, and IS. This means that if ADI is set under 3, if the system allows to choose TSB, most of the SKUs will still choose weighted average, single exponential smoothing, and intermittence smoothing as the best method. When ADI is larger than 3, more spare parts were included in erratic and smooth demand categories; and from Table 7.2, I found that most of the spare parts chose TSB as the best method. There are 426 spare parts choose TSB and 411 spare parts choose WA, SES, and IS. As ADI increases, TSB becomes preferable among spare parts, with an increasing number of them being selected as the final best method.

Therefore, the potential threshold of ADI is proposed to be 3 or more than 3 using statistical counting as evidence. The key behind logic is to classify more SKUs earlier marked as intermittent demand to erratic or smooth demand. Within erratic and smooth categories, TSB is not necessary to be implemented in Servigistics since traditional methods like WA are more effective than TSB according to results from the experiment.

Second evidence for choosing new ADI is through k-means clustering. Figure 7.1 indicates the optimal clusters result from silhouette plot. The plot suggests that three clusters are suitable for clustering, but to include more information, I used four clusters. In Figure 7.2, each SKU in squared CV and ADI dimensions is labelled with different colors according to the assigned clusters. The "+" sign marks center of that cluster, red line is "ADI=1.32", blue line is "CV^2= 0.49". Table 7.3 gives a detailed description of characteristics of each cluster. From four clusters segmented by algorithm, generally the possibility of choosing TSB as best method increases when ADI increases. When ADI is equals to 4.58, the percentage of TSB selected as final best method exceed that weighted average. Thus, the ADI modified to around 4.5 can be a potential threshold for future Philips's experiments.

Additionally, since I set the algorithm balancing volatility and accuracy of the forecasts when choosing the methods, the final methods involved a fluctuation factor (i.e., volatility). Thus, the results can be more convincing according to the requirements from Philips. However, I observed that half of the SKUs chose WA as the final method when ADI is smaller than 3. It can be addressed that WA is a crucial method in the experiment. The further performance of WA is investigated in experiment two in the next section. Moreover, TSB takes a large part of all chosen methods, the effectiveness and performance of TSB is elaborated in experiment two.

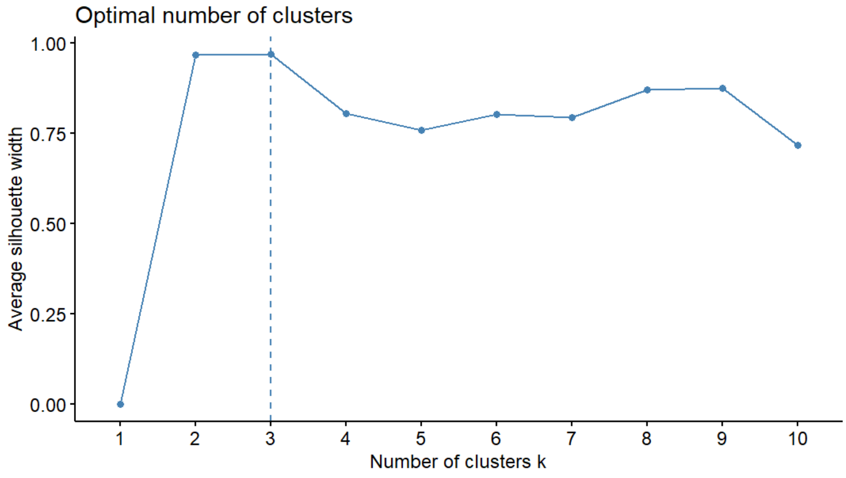

Figure 7.1: distribution of spare parts with assigned clusters.

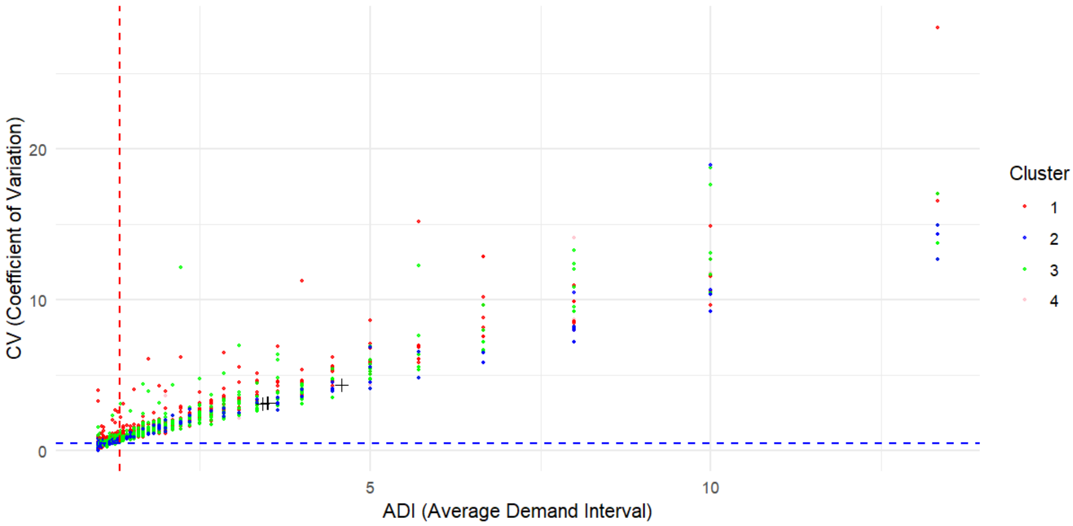

Figure 7.2: distribution of spare parts with assigned clusters.

Center	ADI	CV²	IS	SES	TSB	WA
1	3.415	3.119	0.033	0.028	0.348	0.590
2	3.504	3.179	0.026	0.068	0.350	0.556
3	3.482	3.160	0.043	0.025	0.413	0.519
4	4.583	4.353	0.040	0.051	0.485	0.424

Table 7.3: clusters information from k-means clustering.

7.2 Experiment two: the performances of TSB and NN

The second experiment examines the performance of TSB and NN in Philips's lumpy demand data sets. The results from WA, SES, IS, TSB, and NN are presented and compared. Results are presented in two ways: at an aggregated level and each part level. "At an aggregated level" means summing up the forecasts from all spare parts for each forecasting month. "At a part level" means forecasts are evaluated for each individual spare part, with key metrics like MAPE and RMSE calculated for each one of them.

7.2.1 Results at an aggregated level

Table 7.4 shows the aggregated amount of forecasted demand from all SKUs of each forecasting month. The aggregated amounts from different forecasting methods are compared and analyzed. Absolute errors of each method, shown as "Error" in Table 7.4, are the difference between actual historical demands and the forecasts. The actual historical demands were calculated by assigning a weight value of 0.4 to the first year's actual demands and weight value of 0.6 to the second year's actual demands. The calculation formula is given in section 6.3. The average of each month of forecasts are presented as a reference. Table 7.5 shows various key metrics on the aggregated demands.

Neural networks perform well considering volatility and accuracy: it has the lowest MAPE and standard deviation, which means that the results from NN are closer to the real demands and the fluctuations are relatively stable. The WA method is important as many spare parts choose it as the final method in Servigistics, and the performances using the top fluctuating parts show that WA is important in forecasting. It is ranked as second best according to MAPE, but WA exhibits volatile forecasting demands. TSB is another focus of this research, but the evidence shows that the accuracy it achieves is relatively low (i.e., 7.3%, 6% higher than MAPE of NN). However, TSB exhibits less volatile forecast results, with a standard deviation of 219.65, it is the second smallest among the traditional methods.

SES is one of the traditional methods in Servigistics, and experiment outcomes support that SES leads to 4.8% of MAPE, which is relatively low. But when it comes to volatility, SES has 328.55 of standard deviation in forecasting, which is the most fluctuating among other methods. The Philips experiment shows that 24.9% of spare parts choose SES as their method. From the above results, it suggests that the application of SES can be one of the reasons causing high fluctuating forecasts.

As stated in Table 7.4, the intermittence smoothing method results in the highest MAPE at 9.8% and RMSE at 359.1. But when considering the volatility of forecasting, the standard deviation of intermittence smoothing method is the smallest, only 131.26. The deficiencies of the IS model may be attributed to the algorithm's setups of my implementation. But due to the invisible operation of Servigistics, this stands as the constraints of the thesis.

Figure 7.3 illustrates trends of aggregated amounts of forecasts, WA, SES, TSB, and NN display a drop in February's forecast and an increase in March's forecast; however, only IS shows an opposite trend (i.e., a rise in February and a drop in March). Line "histdemand" represents the actual historical demands, the calculation can be found in section 6.3. Lines from NN, WA, and TSB are close to line of historical demands, this reveals that forecasts from these methods are more accurate. The forecasts from SES, however, display a volatile trend and are not as accurate as the other three methods during February to August.

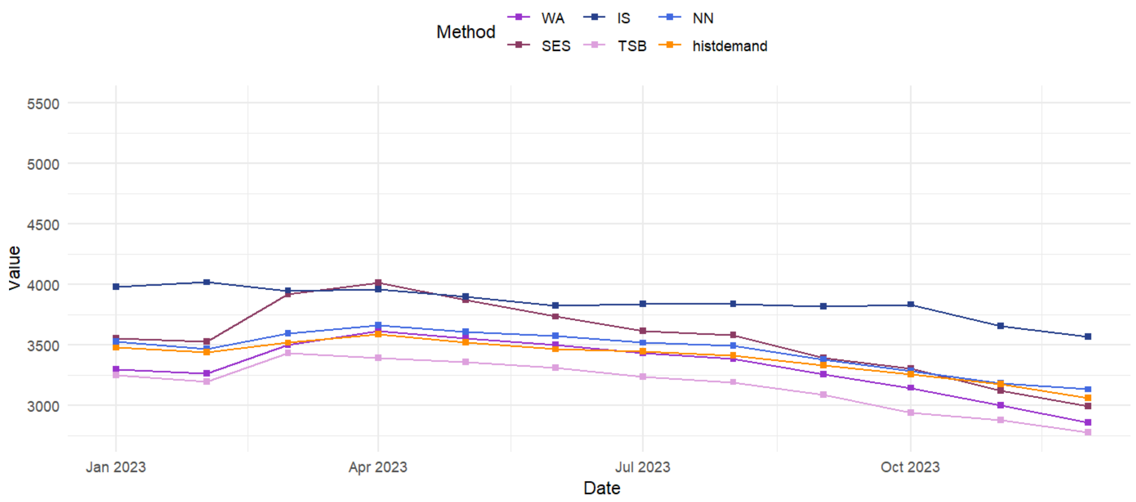

Figure 7.3: line plot of aggregated amounts of forecasts in each month from five methods.

Date (Mm/Dd/YY)	WA	SES	IS	TSB	NN	Actual Demand	Error (WA)	Error (SES)	Error (IS)	Error (TSB)	Error (NN)
01/01/2023	3388.9	3627.9	3978.1	3319.8	3604.1	3487.4	98.5	140.5	490.7	167.6	116.7
02/01/2023	3338.2	3585.8	4016.1	3257.8	3538.7	3456.6	118.4	129.2	559.5	198.8	82.1
03/01/2023	3622.6	4009.4	3947.4	3522.3	3680.8	3603.5	19.1	405.9	343.9	81.2	77.3
04/01/2023	3786.8	4143.0	3959.9	3502.7	3769.5	3761.3	25.5	381.7	198.6	258.6	8.2
05/01/2023	3719.6	3997.7	3898.5	3465.9	3713.0	3677.6	42.0	320.1	220.9	211.7	35.4
06/01/2023	3649.7	3848.8	3822.8	3409.1	3671.7	3625.3	24.4	223.5	197.5	216.2	46.4
07/01/2023	3557.2	3710.3	3835.2	3315.9	3602.5	3585.6	28.4	124.7	249.6	269.7	16.9
08/01/2023	3501.4	3671.2	3838.9	3265.0	3571.1	3555.5	54.1	115.7	283.4	290.5	15.6
09/01/2023	3367.6	3484.2	3817.8	3158.6	3456.8	3470.5	102.9	13.7	347.3	311.9	13.7
10/01/2023	3256.6	3390.0	3832.7	3014.7	3368.5	3379.2	122.6	10.8	453.5	364.5	10.7
11/01/2023	3114.2	3199.6	3654.0	2952.2	3269.7	3294.3	180.1	94.7	359.7	342.1	24.6
12/01/2023	2970.6	3059.6	3565.6	2851.3	3223.2	3178.6	208.0	119.0	387.0	327.3	44.6
Average	3439.5	3644.0	3847.3	3252.9	3539.1	3506.3	85.3	173.3	341.0	253.3	41.0

Table 7.4: aggregated amounts of forecasts.

Metrics	WA	SES	IS	TSB	NN
MAD	85.3	173.3	341.0	253.3	41.0
RMSE	104.9	214.4	359.1	265.2	52.6
MAPE (%)	0.025	0.048	0.098	0.073	0.012
SD	247.08	328.55	131.26	219.65	175.42

Table 7.5: measurements result of each forecasting method.

7.2.2 Results at a part level

The measurements are also compared at a part level where each SKU's forecasting accuracy and volatility are calculated, and the statistical information is indicated in Table 7.6 to Table 7.8. Table 7.6 indicates the MAPE results: WA has the highest accuracy if comparing the mean of MAPE. TSB performs well as it has the third lowest average MAPE.

NN model was tested with the model accuracy using testing data, and the mean MAE result is 0.42. This indicates that the forecasts deviate from the actual demand by approximately 0.412 units. NN has an average MAPE of 22.18% and median MAPE of 18.71%, which means 50% of spare parts have a mean absolute error less than 18.71%. Compared to previous outcomes at an aggregated level, the MAPE results show that TSB, WA, or SES outperform NN at a part level. RMSE measures the deviation of actual and forecasts and Table 7.7 shows the statistical overview of RMSE in all SKUs.

IS has the lowest accuracy and the highest fluctuation across all three metrics. WA is the most accurate method, and NN performs better when comparing RMSE: 0.74 of RMSE, which is 0.15 higher than RMSE of WA. Overall, the performance of TSB is satisfactory that it has the relatively high accuracy and low fluctuation. TSB has a MAPE of 21.28, which is 3% higher than MAPE of WA; its volatility is 0.9, which is 0.2 higher than NN. From the results of two experiments, the accuracy performances of TSB are comparable to those of WA and SES. However, TSB does reduce volatility to some extent at both the aggregated and part levels.

Volatility is another focus of this thesis which measures the fluctuation of forecasts. Under this situation, NN becomes the least fluctuating method, WA performs well in minimizing the volatility of forecasts, and TSB has standard deviation of 0.9 which is also a small fluctuation. SES demonstrates high fluctuation which is consistent with the previous findings at an aggregated level.

Method	Min	Q1	Median	Mean	Q3	MAX
WA	2.792	10.524	14.338	16.01	17.428	93.333
SES	1.043	12.878	17.67	19.97	24.077	65.501
IS	0	15.25	27.492	41.553	46.572	697.232
TSB	5.287	14.867	19.471	21.286	26.782	204.29
NN	4.321	13.027	18.712	22.183	25.436	511.835

Table 7.6: MAPE (%) result of each forecasting method at part level.

Method	Min	Q1	Median	Mean	Q3	MAX
WA	0.086	0.258	0.401	0.546	0.642	9.117
SES	0.015	0.348	0.586	0.777	0.957	11.689
IS	0.154	0.851	1.208	1.631	1.805	16.8
TSB	0	0.316	0.526	0.806	0.921	11.374
NN	0.001	0.312	0.507	0.691	0.822	12.927

Table 7.7: RMSE result of each forecasting method at part level.

Method	Min	Q1	Median	Mean	Q3	MAX
WA	0	0.363	0.592	0.821	0.944	9.177
SES	0	0.487	0.769	1.045	1.233	13.713
IS	0.193	0.75	1.055	1.388	1.577	12.061
TSB	0	0.322	0.576	0.901	1.029	10.181
NN	0	0.268	0.482	0.743	0.836	10.246

Table 7.8: Volatility result of each forecasting method at part level.

Chapter 8 Discussion and Limitations

This section discusses the results from the previous chapter and provides insights into the research questions. Besides, another potential factor named "Archived Forecast" is discussed. Thorough research was conducted on one spare part for explaining "Archived Forecast" and this research is included in Appendix 1 as reference. The thresholds suggested for Philips are reviewed, and the performances of five methods in the lumpy pattern are discussed further under Philips situation. The limitations of this research and problems that occur during the thesis are also pointed out.

8.1 The effects of "Archived Forecast" and the research on an individual spare part.

During the exploration of Servigistics, I found that a variable named "Archived Forecast" may have an important effect on the forecasts. SPS team had provided several spare parts as examples when stating problems and one of the spare parts, part 459801352572, was used by me for examining the problems of adopting "Archived Forecast".

"Archived Forecast" is the stored forecast in Servigistics based on the previous forecast or demands. I questioned the function and calculation logic of the "Archived Forecast" with PTC team, but by the time this thesis was completed, no response had been received. Thus, the discussion was based on my assumptions and results from implementation of assumed forecasting logics.

I observed that the daily forecasts in April 2024 in an aggregated amount (that is, sum of forecasts from all locations) were absurdly volatile (Appendix Figure 1.1). After closely examining the forecasts and facts in Servigistics, I assumed: due to Servigistics's algorithm, forecasts are using "Archived Forecast" for forecasting instead of historical actual demands and forecasting models.

The detailed investigation on part 459801352572 is included in Appendix 1. The results confirm that the assumption is valid. The Servigistics algorithm can cause significant forecast fluctuations, partly due to the use of "Archived Forecast". Some calculations of forecasts do not follow the Servigistics 's forecasting methodology but instead rely on archived forecasts. For example, in February, the forecasts of the year are calculated by "Archived Forecast" times 12 months.

8.2 Thresholds for categorization.

Philips sets a threshold of ADI to 1.32 to distinguish smooth from intermittent demand, which includes more SKUs as intermittent demand items. The research goal is to explore the possibility of increasing ADI. For the scope of this thesis, the effects of adopting a new ADI cannot be easily researched after discussing with SPS team. Instead, SPS can implement the recommended ADI in SBOX to accurately determine its effects.

This thesis created a simple experiment to propose potential changes in thresholds. The key idea of this experiment is finding the switch point of SKUs choosing forecasting methods. Finally, ADI equals 3 is a possible number when using statistical counting. K-means is adopted for reference and suggests an ADI equalling 4.5 to be a viable option. But according to the outcomes listed in Table 7.3, there is no obvious change in ADI after comparing the methods used in different clusters. Therefore, I would suggest ADI equals to 3 as the most likely new ADI. When increasing ADI to 3 or 4.5, 35% more SKUs categorized as intermittent demand are dealt with by methods for erratic and smooth demand. TSB is planned to implement only in intermittent demand. By increasing the ADI threshold, the number of SKUs in intermittent demand are reduced by 35% which streamlines the process.

8.3 Methods discussion: the performances of TSB and NN implementation in lumpy pattern.

In Servigistics, WA, SES, and IS are the top three methods in forecasting; TSB is another important focus that the performance on the lumpy or intermittent demand pattern has a significant impact on the further Philips's experiments. NN is suggested by previous studies and has also been investigated in Philips data set as well.

The second experiment examined the main five forecasting models on the lumpy demand pattern, especially on those most fluctuating spare parts. Three measurements were used to evaluate the forecasts, MAPE and SD are the key metrics. From the results at a part level, WA achieves the highest accuracy, with a MAPE of 16.1% and an RMSE of 0.54, which are less than half of the values obtained from other methods. During the recent experiment conducted by Philips, 28.5% of the spare parts chose WA as the main method before the simulation. The results from my experiment echo the previous findings in Philips that WA is a crucial technique in forecasting. This also resonates with some studies on intermittent demand patterns. Ghobbar (2004) conducted a study on aircraft operators in their components that have intermittent pattern. The results show continued superiority of weighted average, Holt, and Croston methods (Ghobbar, 2004).

TSB and NN are explored and performance of these two methods are compared with the traditional methods in the system. TSB has not been implemented in the system, and Philips pays attention to the effectiveness of this method. From the results shown above, TSB performed worse compared to WA and SES when considering MAPE and volatility metric. The mean of MAPE for TSB in all spare parts is 21.28% which is 7% higher than WA and 1% higher than SES. The standard deviation is 0.901 across all spare parts, which is 0.1 higher than WA and indicates that the fluctuation of forecasts from this method is relatively low. Therefore, though TSB is comparable to WA or SES in terms of accuracy, TSB can still be considered as one option for lumpy or intermittent demand. Doszyn (2020) conducted research on SBA, Croston, TSB, and SES, he found that TSB method outperforms other methods for all products. In case of erratic and smooth items, TSB method yielded the poorest results (Doszyn, 2020). Another research modified TSB, and TSB method achieves the best results on MASE and RMASE among all comparison methods (Yang et al., 2021).

NN is a machine learning method and has been proved to perform well in some previous studies. Results from Shahwan & Said (2012) assure that when the demand data is more sporadic, i.e. have more zero values, then neural network becomes a better forecasting tool. This research carried out NN in the most fluctuating SKUs exhibit lumpy pattern. The MAPE value of NN is 22.18% which did not show the better performance of NN compared to TSB, WA, or SES. But when considering RMSE and volatility, NN outperformed WA, SES, and TSB. NN has mean values of 0.69 in RMSE and mean volatility value of 0.743 across all spare parts. The volatility of NN is around 0.2 to 0.3 lower than other methods. In addition, when focusing on the aggregated amounts, it is noteworthy that the results of NN exhibit the highest accuracy and the lowest volatility at an aggregated level (i.e., MAPE equals to 1.8%).

After balancing the results at an aggregated level and part level, NN did lead to the best forecasting results; and for TSB, it has the lowest fluctuation which indicates that adopting TSB method may improve volatility to some extent. WA has the highest accuracy at an aggregated level, but the volatility is not desirable compared to TSB. Since in Servigistics, it is not possible to adopt NN, thus TSB is a potential method for reducing the volatility of forecasting. The main aspects Philips focuses on regarding the forecasting are accuracy and volatility. As stated before, they would like to know methods that improve accuracy while decreasing the fluctuations to some extent. TSB has results that indicate that TSB ranks third in accuracy and second in volatility. SES has relatively high accuracy, which is only 0.8% lower MAPE than TSB, but it has 0.1 higher volatility compared to TSB. After balancing the effects, TSB still stands out and it is meaningful to apply in intermittent demand in Philips system.

8.3 Limitations

The primary and most significant limitation of this thesis is the model building for the demand forecasting. PTC is the third-party IT service company for Philips; thus, the operation of Servigistics is unseeable which cause obstacles to the research. The models of forecasting and calculation formulas are constructed under documented instructions from SPS team. The Servigistics algorithms created by PTC especially for Philips's demand forecasting are not included, this may contribute to the inaccuracies in results compared to forecasts number observed in planning platform.

In addition, this research applied NN in the simplest form, which is a single layer neural network. MAPE and volatility are the main measurements to echo with the settings in the system. The future research area can adopt multi-layered networks and select more metrics for evaluating the performance of models. Furthermore, due to the limited capacity of software, only a small subset of parameter values was tested. Expanding the range of values can provide more insights into model performance.

The intransparency of Servigistics operation hampered the research. For example, some of the spare parts are classified into different categories when relying on the documented definition of demand pattern. Archived forecast mentioned in part research is a potential reason for fluctuating forecasts. The requests for the definition and implementation scenarios of archived forecast are not answered by PTC team. This serves as another limitation of this thesis.

Chapter 9 Suggestions for Philips

This chapter presents the suggestions after experimenting with Philips's data sets. The suggestions are in two aspects: Servigistics algorithm and forecasting methods selection.

Regarding the Servigistics algorithm, I observed on Servigistics online platform that, variables like "Archived Forecast" are used for getting yearly forecasts. This can simplify the forecasting process since only the "Archived Forecast" are calculated based on historical demands and yearly forecasts are equal to "Archived Forecast" times 12 periods. However, the usage of "Archived Forecast" may lead to difficulties explaining the forecasts in each month. Moreover, forecasting results from "Archived Forecast" may cause fluctuations in forecasts. From my observations, for some spare parts in certain locations, the forecasts may shift from using embedded Servigistics forecasting methods to "Archived Forecast" calculations. This contributes to some of the volatility in forecasting. The future work can be conducted to reduce the fluctuations on such reasons.

For forecasting methods, Philips is experimenting with TSB in intermittent demand SKUs. From results of this research, TSB exhibits relatively high accuracy and low volatility compared to SES and IS. Thus, it is worth noting that TSB implementing in intermittent demand may lead to better performance of forecasts. WA is another aspect that Philips should pay attention to since WA has more stable and accurate forecasts. The calculation of WA is easier to implement in Servigistics. Therefore, I suggest that WA can be the second option to consider except for TSB when figuring out methods used for intermittent demand.

Appendix 1

Research on Individual Part

This section is on research for an individual part, some facts are included in Appendix. To better understand Philips' inquiries and the demand forecasting issue, I researched a specific spare part with abnormal fluctuations. The part number is 459801352572, and this spare part is also indicated in the problematic parts reported by Philips.

Characteristics of Part 459801352572

Part 459801352572 is the chosen part for investigation, with product price of €15,490. It represents the service spare part Detector PX3040, and the warehouse locations are in China (see Appendix Table 1). The demand data from January 1, 2021, to April 1, 2024, was extracted, with demands rolling up to the first day of each month. A zero in the demand record indicates that there may have been demand, but the product was out of stock, preventing any sales. The records of demand without zero values are listed in Table 3.2, the months with no records in the system means there is no demand during that period. From the demand record, this spare part has a max demand of 2 and minimum demand of 1, with most of the time being no demand (that is, zero demand).

Parameters related to CV^2 and ADI are presented in Table 3.1. These are calculated and queried from database of PTC or copied from the online platform. CV^2 and ADI are not calculated for certain locations due to insufficient demand history information. According to theory, part 459801352572 should be categorized as lumpy demand, but it is shown as intermittent since Philips excludes lumpy demand. In addition, this section uses existing demand thresholds, CV^2 and ADI, for exploring potential reasons and research directions. For the experiments of the thesis, CV^2 and ADI are manually calculated using the past 40 periods historical demands.

Philips SPS team provided daily forecasts information pulled from PTC system from February 28, 2024, and April 12, 2024, to illustrate the volatility of forecasting. Daily forecasts information pulled from PTC planning platform is in Figure 3.1, which shows an unstable forecasting line. Regarding the forecasting methods of part 459801352572, the algorithms select four methods: average, weighted average, single exponential smoothing, and intermittent smoothing. The forecasting numbers of Figure 3.1 for this are shown Appendix, which is the aggregated number of forecasts of each location.

After closely examining the forecasts from each location over the period, I have identified two potential reasons for the fluctuations: Firstly, it could be attributed to the system algorithm within PTC's planning platform, which may not be accessible to the SPS team. Secondly, significant fluctuations in forecasts may be due to changes in forecasting methods. The subsequent sections provide a detailed explanation of these two reasons. A preliminary experiment is conducted, and results are demonstrated for elaborating the assumptions and research directions of the thesis.

Location	Part Number	SD	Mean	CV^2	ADI
CN6SU690U	459801352572	0.61	0.28	4.82	4.67
CN6SU687U	459801352572	0.20	0.04	25	14
CN6SU688U	459801352572	0.28	0.08	11.30	2
CN93U632U	459801352572	0.41	0.17	5.82	4

Appendix Table 1.1: parameters information of part 459801352572.

Location	Part Number	Unit Cost	Location Type	Demand Type	Quantity	Demand Date
CN6SU690U	459801352572	15490	LDC	Local Demand	2	2021-06-01
CN6SU690U	459801352572	15490	LDC	Local Demand	2	2021-09-01
CN6SU690U	459801352572	15490	LDC	Local Demand	1	2021-10-01
CN6SU690U	459801352572	15490	LDC	Local Demand	1	2022-05-01
CN6SU690U	459801352572	15490	LDC	Local Demand	1	2022-10-01
CN6SU690U	459801352572	15490	LDC	Local Demand	1	2023-01-01
CN6SU687U	459801352572	15490	DC	Local Demand	1	2023-02-01
CN6SU690U	459801352572	15490	LDC	Local Demand	2	2023-05-01
CN6SU688U	459801352572	15490	LDC	Local Demand	1	2023-10-01
CN6SU688U	459801352572	15490	LDC	Local Demand	1	2023-11-01
CN93U632U	459801352572	15490	FSL	Local Demand	1	2023-12-01

Appendix Table 1.2: historical demands of part 459801352572.

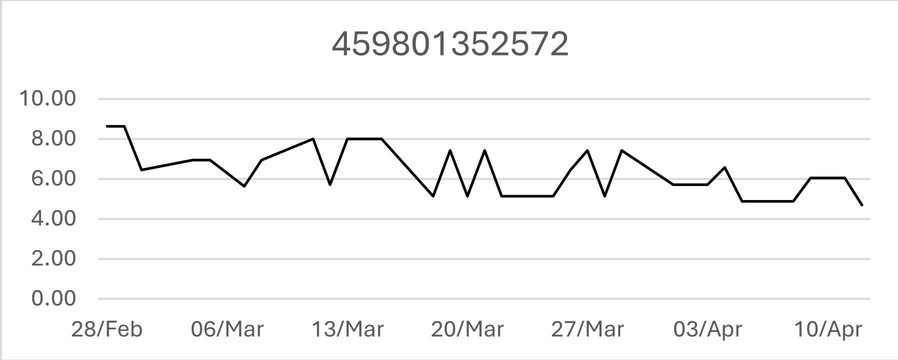

Appendix Figure 1.1: daily forecasts of part 459801352572 (aggregated).

Assumption: Servigistics algorithm may cause the fluctuation of forecasting.

Upon careful examination of forecasts from each location, it was observed that in location "CN93U632U", the forecasted values were significantly higher than the actual demands recorded over the past three years. Specifically, the actual demand from "Location_Part Demand" (see Data section) in the past three years has consistently been only 1, which is further supported by information obtained from online platforms. However, the forecasts for this location are consistently around 2. The variability in forecasting, fluctuating from 2.86 to 2.29 to 1.71, can partially account for the fluctuations observed.

After consulting with the SPS team, it has been suggested that this discrepancy may arise from the presence of an "archived forecast" setting in the system. In this particular location, each month's archived forecast is saved and presumed to be the forecast for the subsequent 12 periods. Consequently, the forecasting for a given year is calculated by multiplying the archived forecast by 12 months. The proposed calculations are outlined in Table 3.3, where "assumed forecasts" represent the calculated values under my assumption, while "actual forecasts" are the forecasts provided directly by the SPS team. The assumed forecasts under assumption one match with the actual forecasts extracted from the system.

Month	Archived Forecast	Assumed Forecast	Actual Forecasts
February	0.24	2.88	2.88
March	0.19	2.28	2.28
April	0.14	1.71	1.71

Appendix Table 1.3: test assumption.

Summary of Individual Part Research

Part 459801352572 is investigated, and two assumptions are examined. The results confirm that both assumptions are valid. The Servigistics algorithm can cause significant forecast fluctuations, partly due to the use of "archived forecasts". Some calculations do not follow Servigistics forecasting methodology and instead rely on archived forecasts. Appendix Table 1.3 supports the first assumption. Due to the intransparency of Servigistics problem, further research can be conducted to looks into the effectiveness of "archived forecasts".

Appendix 2

Rule Type	Status	Eliminate forecast Methods
Soft	Trend detected	Trend detected. Average eliminated, Weighted Average Eliminated, Moving Average eliminated. Single Exponential eliminated.
Soft	LastYearAvg < 3	LastYearAverage less than 3. Moving Average eliminated.
Hard	NumHistDemands < 2	Total number of historical demands less than 2. Intermittence Smoothing eliminated.
Hard	#StreamHistSlices < 12	Stream config history slices less than 12. Intermittence Smoothing eliminated.
Hard	Insufficient History	# of history slices less than 12. Intermittence Smoothing eliminated. Winters Multiplicative eliminated.
Soft	#HistSlices < x	Number of history slices less than value. Intermittence Smoothing eliminated.
Soft	Not intermittent	Intermittency test failed. Intermittence Smoothing eliminated.
Hard	#StreamHistSlices < x	Stream config history slices less than global setting INIT_EXPSMOOTHING_MONTHS+3. Double Exponential eliminated.
Hard	#HistSlices < y	# of history slices less than global setting INIT_EXPSMOOTHING_MONTHS+3. Double Exponential eliminated.
Soft	#HistSlices < z	Number of history slices less than val(minDoubleExpHistory). Double Exponential eliminated.
Soft	TotalDemandSlices < 5	Number of history slices less than 5. Double Exponential eliminated.
Hard	LastYearAverage < 2	LastYearAverage less than 2. Winters Multiplicative eliminated.
Hard	#StreamHistSlices < 12	Stream config history slices less than 12. Winters Multiplicative eliminated.
Hard	#HistSlices < x	Number of history slices less than val(minWintersHist). Winters Multiplicative eliminated.
Hard	Not seasonal	Auto correlation test failed. Winters Multiplicative eliminated.
Hard	#Slides < 1	Number of slides less than 1. All forecast methods eliminated. The number of times that MAPE, MAD, and RMSE are calculated using the holdout window positioned within the forecast window. The holdout window is initially aligned with the first time slice in the forecast window, and the error values are calculated. Best Fit will then "slide" the holdout window one time slice later and recalculate the error values until the number of slides is reached. The resulting error values are averaged over the number of slides.

Appendix 2 Table 1: Bestfit Rules

Part Number	Date	Total
459801352572	28/02/2024	8.63
459801352572	29/02/2024	8.63
459801352572	01/03/2024	6.45
459801352572	04/03/2024	6.95
459801352572	05/03/2024	6.95
459801352572	07/03/2024	5.63
459801352572	08/03/2024	6.95
459801352572	11/03/2024	8
459801352572	12/03/2024	5.71
459801352572	13/03/2024	8
459801352572	14/03/2024	8
459801352572	15/03/2024	8
459801352572	18/03/2024	5.13
459801352572	19/03/2024	7.42
459801352572	20/03/2024	5.13
459801352572	22/03/2024	5.13
459801352572	25/03/2024	5.13
459801352572	26/03/2024	6.45
459801352572	27/03/2024	7.42
459801352572	28/03/2024	5.13
459801352572	29/03/2024	7.42
459801352572	01/04/2024	5.71
459801352572	02/04/2024	5.71
459801352572	03/04/2024	5.71
459801352572	04/04/2024	6.58
459801352572	05/04/2024	4.87
459801352572	08/04/2024	4.69
459801352572	09/04/2024	6.05
459801352572	10/04/2024	6.05
459801352572	11/04/2024	6.05
459801352572	12/04/2024	4.69

Appendix 2 Table 2: Aggregated forecasts from system during February 28 and April 12 on part level.

Location	Part.Number	Unit.Cost	Quantity	Demand.Exte	Demand.Date	Quantity.Fille
CN6SU687U	459801352572	15490	0	0	2021-06-01 00:00:00.0	
CN6SU688U	459801352572	15490	0	0	2021-06-01 00:00:00.0	
CN6SU690U	459801352572	15490	2	0	2021-06-01 00:00:00.0	
CN6SU687U	459801352572	15490	0	0	2021-09-01 00:00:00.0	
CN6SU688U	459801352572	15490	0	0	2021-09-01 00:00:00.0	
CN6SU690U	459801352572	15490	2	0	2021-09-01 00:00:00.0	
CN6SU688U	459801352572	15490	0	0	2021-10-01 00:00:00.0	
CN6SU690U	459801352572	15490	1	0	2021-10-01 00:00:00.0	
CN6SU688U	459801352572	15490	0	0	2022-05-01 00:00:00.0	
CN6SU690U	459801352572	15490	1	0	2022-05-01 00:00:00.0	
CN6SU690U	459801352572	15490	1	0	2022-10-01 00:00:00.0	
CN6SU687U	459801352572	15490	0	0	2023-01-01 00:00:00.0	
CN6SU690U	459801352572	15490	1	0	2023-01-01 00:00:00.0	
CN6SU687U	459801352572	15490	1	0	2023-02-01 00:00:00.0	
CN6SU690U	459801352572	15490	0	0	2023-02-01 00:00:00.0	
CN6SU687U	459801352572	15490	0	0	2023-05-01 00:00:00.0	
CN6SU690U	459801352572	15490	2	0	2023-05-01 00:00:00.0	
CN6SU688U	459801352572	15490	1	15490	2023-10-01 00:00:00.0	
CN6SU688U	459801352572	15490	1	15490	2023-11-01 00:00:00.0	
CN93U632U	459801352572	15490	1	15490	2023-12-01 00:00:00.0	
CN6SU688U	459801352572	15490	0	0	2023-12-01 00:00:00.0	

Appendix 2 Table 3: actual demand for past 3 years.

Part Number	Location Nam	Forecast Method	Total	Date
459801352572	CN6SU687U	Single Exp Smoothing	0.34	2024-02-28 00:00:00
459801352572	CN6SU687U	Single Exp Smoothing	0.34	2024-02-29 00:00:00
459801352572	CN6SU687U	Weighted Average	1.08	2024-03-11 00:00:00
459801352572	CN6SU687U	Weighted Average	1.08	2024-03-12 00:00:00
459801352572	CN6SU687U	Weighted Average	1.08	2024-03-13 00:00:00
459801352572	CN6SU687U	Weighted Average	1.08	2024-03-14 00:00:00
459801352572	CN6SU687U	Weighted Average	1.08	2024-03-15 00:00:00
459801352572	CN6SU687U	Weighted Average	1	2024-04-01 00:00:00
459801352572	CN6SU687U	Weighted Average	1	2024-04-02 00:00:00
459801352572	CN6SU687U	Weighted Average	1	2024-04-03 00:00:00
459801352572	CN6SU687U	Weighted Average	1	2024-04-04 00:00:00
459801352572	CN6SU687U	Weighted Average	1	2024-04-05 00:00:00

Appendix 2 Table 4: Daily forecasts in location CN6SU687U.

Part Number	Location Nam	Forecast Method	Total	Date
459801352572	CN6SU688U	Single Exp Smoothing	2.93	2024-02-28 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.93	2024-02-29 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	1.66	2024-03-01 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	1.66	2024-03-04 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	1.66	2024-03-05 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-07 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	1.66	2024-03-08 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-11 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-12 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-13 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-14 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-15 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-18 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-19 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-20 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-22 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-25 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	1.66	2024-03-26 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-27 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-28 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.63	2024-03-29 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	1.5	2024-04-01 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	1.5	2024-04-02 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	1.5	2024-04-03 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.37	2024-04-04 00:00:00
459801352572	CN6SU688U	Single Exp Smoothing	2.37	2024-04-05 00:00:00
459801352572	CN6SU688U	Intermittence Smooth	1.5	2024-04-08 00:00:00
459801352572	CN6SU688U	Intermittence Smooth	3.43	2024-04-09 00:00:00
459801352572	CN6SU688U	Intermittence Smooth	3.43	2024-04-10 00:00:00
459801352572	CN6SU688U	Intermittence Smooth	3.43	2024-04-11 00:00:00
459801352572	CN6SU688U	Intermittence Smooth	2.07	2024-04-12 00:00:00

Appendix 2 Table 5: Daily forecasts in location CN6SU688U.

Part Number	Location Nam	Forecast Method	Total	Date
459801352572	CN6SU690U	Average	2.5	2024-02-28 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-02-29 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-03-01 00:00:00
459801352572	CN6SU690U	Average	3	2024-03-04 00:00:00
459801352572	CN6SU690U	Average	3	2024-03-05 00:00:00
459801352572	CN6SU690U	Average	3	2024-03-07 00:00:00
459801352572	CN6SU690U	Average	3	2024-03-08 00:00:00
459801352572	CN6SU690U	Average	2	2024-03-11 00:00:00
459801352572	CN6SU690U	Average	2	2024-03-12 00:00:00
459801352572	CN6SU690U	Average	2	2024-03-13 00:00:00
459801352572	CN6SU690U	Average	2	2024-03-14 00:00:00
459801352572	CN6SU690U	Average	2	2024-03-15 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-03-18 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-03-19 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-03-20 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-03-22 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-03-25 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-03-26 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-03-27 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-03-28 00:00:00
459801352572	CN6SU690U	Average	2.5	2024-03-29 00:00:00
459801352572	CN6SU690U	Average	1.5	2024-04-01 00:00:00
459801352572	CN6SU690U	Average	1.5	2024-04-02 00:00:00
459801352572	CN6SU690U	Average	1.5	2024-04-03 00:00:00
459801352572	CN6SU690U	Average	1.5	2024-04-04 00:00:00
459801352572	CN6SU690U	Average	1.5	2024-04-05 00:00:00
459801352572	CN6SU690U	Single Exp Smoothing	0.91	2024-04-08 00:00:00
459801352572	CN6SU690U	Single Exp Smoothing	0.91	2024-04-09 00:00:00
459801352572	CN6SU690U	Single Exp Smoothing	0.91	2024-04-10 00:00:00
459801352572	CN6SU690U	Single Exp Smoothing	0.91	2024-04-11 00:00:00
459801352572	CN6SU690U	Single Exp Smoothing	0.91	2024-04-12 00:00:00

Appendix 2 Table 6: Daily forecasts in location CN6SU690U.

Part Number	Location Nam	Forecast Method	Total	Date
459801352572	CN93U632U	Weighted Average	2.86	2024-02-28 00:00:00
459801352572	CN93U632U	Weighted Average	2.86	2024-02-29 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-01 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-04 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-05 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-08 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-11 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-13 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-14 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-15 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-19 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-26 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-27 00:00:00
459801352572	CN93U632U	Weighted Average	2.29	2024-03-29 00:00:00
459801352572	CN93U632U	Weighted Average	1.71	2024-04-01 00:00:00
459801352572	CN93U632U	Weighted Average	1.71	2024-04-02 00:00:00
459801352572	CN93U632U	Weighted Average	1.71	2024-04-03 00:00:00
459801352572	CN93U632U	Weighted Average	1.71	2024-04-04 00:00:00
459801352572	CN93U632U	Weighted Average	1.71	2024-04-08 00:00:00
459801352572	CN93U632U	Weighted Average	1.71	2024-04-09 00:00:00
459801352572	CN93U632U	Weighted Average	1.71	2024-04-10 00:00:00
459801352572	CN93U632U	Weighted Average	1.71	2024-04-11 00:00:00
459801352572	CN93U632U	Weighted Average	1.71	2024-04-12 00:00:00

Appendix 2 Table 7: Daily forecasts in location CN93U632U.

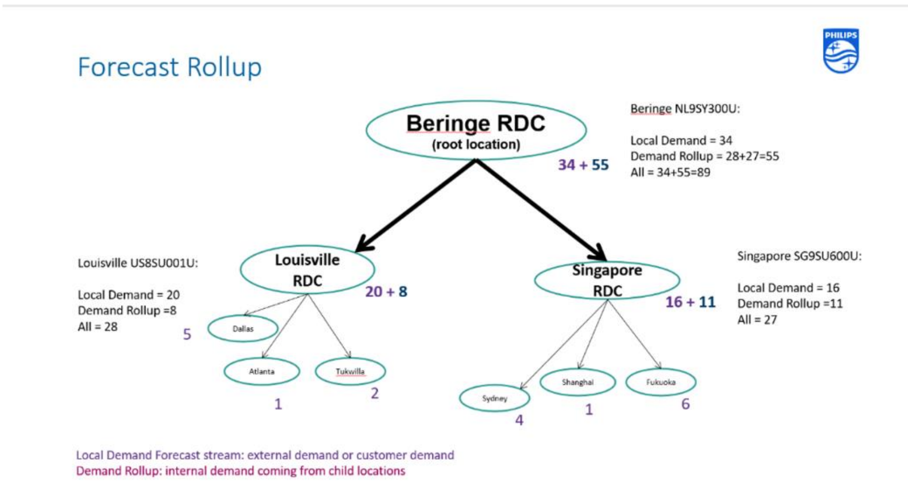

Appendix 2 Table 8: Demand Roll-up, from Philips SPM trainings by PTC.

References

Altay, N., & Litteral, L. A. (2011). *Service Parts Management Demand Forecasting and inventory control*. Springer.

Babai, M. Z., Dallery, Y., Boubaker, S., & Kalai, R. (2019). A new method to forecast intermittent demand in the presence of inventory obsolescence. *International Journal of Production Economics*, *209*, 30–41. https://doi.org/10.1016/j.ijpe.2018.01.026

Bookbinder, J. H., & Lordahl, A. E. (1989). Estimation of inventory re-order levels using the bootstrap statistical procedure. *IIE Transactions*, *21*(4), 302–312. https://doi.org/10.1080/07408178908966236

Cheng, C.-Y., Chiang, K.-L., & Chen, M.-Y. (2016). Intermittent demand forecasting in a tertiary pediatric intensive care unit. Journal of Medical Systems, 40(10), 217. https://doi.org/10.1007/s10916-016-0571-9

Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Journal of the Operational Research Society*, *23*(3), 289–303. https://doi.org/10.1057/jors.1972.50

Deng, W., Lyu, G., Shi, Y., & Wang, W. (2017, September). Electricity consumption prediction using XGBoost based on discrete wavelet transform. In *2nd International Conference on Artificial Intelligence and Engineering Applications* (pp. 716–729). Guilin, China.

Doszyn, M. (2020). Accuracy of intermittent demand forecasting systems in the enterprise. *EUROPEAN RESEARCH STUDIES JOURNAL*, *XXIII*(Issue 4), 912–930. https://doi.org/10.35808/ersj/1723

Durlinger, P., & Paul, I. (2012). *Inventory and holding costs*. Durlinger Consultant.

Gilliland, M., 2002. Is forecasting a waste of time? *Supply Chain Management Review 6* (4), 16–23.

Global leader in product lifecycle management software. PTC. (2024, April 24). https://www.ptc.com/en

Ghobbar, A. A. (2004). Forecasting intermittent demand for aircraft spare parts: A comparative evaluation of methods. *Journal of Aircraft*, *41*(3), 665–673. https://doi.org/10.2514/1.851

Gutierrez, R. S., Solis, A. O., & Mukhopadhyay, S. (2008). Lumpy demand forecasting using neural networks. *International Journal of Production Economics*, *111*(2), 409–420. https://doi.org/10.1016/j.ijpe.2007.01.007

Hoffmann, M. A., Lasch, R., & Meinig, J. (2022). *Forecasting irregular demand using single hidden layer neural networks* (6th ed.). Bundesvereinigung Logistik (BVL) e.V. https://doi.org/10.23773/2022_6

Hua, Z., & Zhang, B. (2006). A hybrid support vector machines and logistic regression approach for forecasting intermittent demand of spare parts. *Applied Mathematics and Computation, 181*(2), 1035–1048. https://doi.org/10.1016/j.amc.2005.11.157

Kanungo, T., Mount, D. M., Netanyahu, N. S., Piatko, C., Silverman, R., & Wu, A. Y. (2000). The analysis of a simple *k* -means clustering algorithm. *Proceedings of the Sixteenth Annual Symposium on Computational Geometry*, 100–109. https://doi.org/10.1145/336154.336189

Kiefer, D., Grimm, F., Bauer, M., & Van, D. (2021). Demand forecasting intermittent and lumpy time series: Comparing statistical, machine learning and deep learning methods. http://hdl.handle.net/10125/70784

Kostenko, A. V., & Hyndman, R. J. (2006). A note on the categorization of demand patterns. *Journal of the Operational Research Society*, *57*(10), 1256–1257. https://doi.org/10.1057/palgrave.jors.2602211

Kourentzes, N. (2014). On intermittent demand model optimisation and selection. *International Journal of Production Economics*, *156*, 180–190. https://doi.org/10.1016/j.ijpe.2014.06.007

Lengu, D., Syntetos, A. A., & Babai, M. Z. (2014). Spare parts management: Linking distributional assumptions to demand classification. *European Journal of Operational Research*, *235*(3), 624–635. https://doi.org/10.1016/j.ejor.2013.12.043

Lolli, F., Gamberini, R., Regattieri, A., Balugani, E., Gatos, T., & Gucci, S. (2017). Single-hidden layer neural networks for forecasting intermittent demand. *International Journal of Production Economics*, *183*, 116–128. https://doi.org/10.1016/j.ijpe.2016.10.021

Martin, D., Spitzer, P., & Kühl, N. (2020). A new metric for lumpy and intermittent demand forecasts: Stock-keeping-oriented prediction error costs. Hawaii International Conference on System Sciences. https://doi.org/10.24251/HICSS.2020.121

Nenni, M. E., Giustiniano, L., & Pirolo, L. (2013). Demand forecasting in the fashion industry: A review. *International Journal of Engineering Business Management*, *5*, 37. https://doi.org/10.5772/56840

Neu, Z., Hicks, B., & Gopsill, J. (2024). Operating minimally intelligent agent-based manufacturing systems across the Average demand Interval – coefficient of variation (Adi-cv) demand state space. *Production & Manufacturing Research*, *12*(1), 2323479. https://doi.org/10.1080/21693277.2024.2323479

Pinçe, Ç., Turrini, L., & Meissner, J. (2021). Intermittent demand forecasting for spare parts: A Critical review. *Omega*, *105*, 102513. https://doi.org/10.1016/j.omega.2021.102513

Porras, E., & Dekker, R. (2008). An inventory control system for spare parts at a refinery: An empirical comparison of different re-order point methods. *European Journal of Operational Research*, *184*(1), 101–132. https://doi.org/10.1016/j.ejor.2006.11.008

PTC help center. (n.d.). https://trne-prod.ptcmanaged.com/servigistics_help/en/index.html#page/PTC_Servigistics_Help_Center/glossary/intermittence_smoothing_forecast_method.html

Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, *20*, 53–65. https://doi.org/10.1016/0377-0427(87)90125-7

Sapankevych, N., & Sankar, R. (2009). Time series prediction using support vector machines: A survey. *IEEE Computational Intelligence Magazine, 4*(2), 24–38. https://doi.org/10.1109/MCI.2009.932254

Shahwan, T., & Said, R. (2012). A comparison of bayesian methods and artificial neural networks for forecasting chaotic financial time series. *Journal of Statistics Applications & Probability*, *1*(2), 89–100. https://doi.org/10.12785/jsap/010202

Snyder, R. (2002). Forecasting sales of slow and fast moving inventories. *European Journal of Operational Research*, *140*(3), 684–699. https://doi.org/10.1016/S0377-2217(01)00231-4

Syntetos, A. A., & Boylan, J. E. (2001). On the bias of intermittent demand estimates. *International Journal of Production Economics*, *71*(1–3), 457–466. https://doi.org/10.1016/S0925-5273(00)00143-2

Syntetos, A. A., Boylan, J. E., & Croston, J. D. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society*, *56*(5), 495–503. https://doi.org/10.1057/palgrave.jors.2601841

Syntetos, A. A., Zied Babai, M., & Gardner, E. S. (2015). Forecasting intermittent inventory demands: Simple parametric methods vs. bootstrapping. *Journal of Business Research*, *68*(8), 1746–1752. https://doi.org/10.1016/j.jbusres.2015.03.034

Teunter, R. H., Syntetos, A. A., & Zied Babai, M. (2011). Intermittent demand: Linking forecasting to inventory obsolescence. *European Journal of Operational Research*, *214*(3), 606–615. https://doi.org/10.1016/j.ejor.2011.05.018

Van Wingerden, E., Basten, R. J. I., Dekker, R., & Rustenburg, W. D. (2014). More grip on inventory control through improved forecasting: A comparative study at three companies. *International Journal of Production Economics*, *157*, 220–237. https://doi.org/10.1016/j.ijpe.2014.08.018

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. https://doi.org/10.1016/S0169-2070(03)00013-X

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), 529–538. https://doi.org/10.1016/0169-2070(94)90021-3

Williams, T. M. (1984). Stock control with sporadic and slow-moving demand. *Journal of the Operational Research Society*, *35*(10), 939–948. https://doi.org/10.1057/jors.1984.185

Wu, J. (2012). *Advances in K-means clustering: A data mining thinking*. Springer.

Yang, Y., Ding, C., Lee, S., Yu, L., & Ma, F. (2021). A modified Teunter-Syntetos-Babai method for intermittent demand forecasting. *Journal of Management Science and Engineering*, *6*(1), 53–63. https://doi.org/10.1016/j.jmse.2021.02.008

Zied Babai, M., Syntetos, A., & Teunter, R. (2014). Intermittent demand forecasting: An empirical study on accuracy and the risk of obsolescence. *International Journal of Production Economics*, *157*, 212–219. https://doi.org/10.1016/j.ijpe.2014.08.019