MSc Industrial Engineering & Management Thesis

Enhancing Forecasting and Inventory Control for Seasonal Spare Parts at NS

G.J. van Sambeek

UT Supervisors: E. Topan & D.R.J. Prak
NS Supervisor: A. van Houten

July 11, 2024

Industrial Engineering and Management,
Faculty of Behavioural, Management and Social sciences,
University of Twente

Title Page

Intended Audience: For students and employees of the University of Twente
Reference: UT/IEM-G.J.van.Sambeek-11.07.24
Publisher: University of Twente, Industrial Engineering and Management, PO Box 217, 7500 AE Enschede, tel. +31(0)534899111

Title: Enhancing Forecasting and Inventory Control for Seasonal Spare Parts at NS
Author: G.J. van Sambeek
First UT Supervisor: E. Topan
Second UT Supervisor: D.R.J. Prak
Company Supervisor: A. van Houten
Date of Publication: 11 July 2024

Purpose: This thesis is written as part of the Industrial Engineering and Management Master's educational program.

Acknowledgements

Dear reader,

This thesis marks the end of my student career and is the final assignment for my Master's Degree in Industrial Engineering and Management at the University of Twente. The thesis, entitled "Enhancing Forecasting and Inventory Control for Seasonal Spare Parts at NS," has been my biggest project so far. I have overcome a lot of challenges and even the last week, I still learnt something. "Things can go corrupt all the time", is what my Windows software said to me.

First, I would like to thank Joris Emmen for bringing me into contact with the Nederlandse Spoorwegen. Furthermore, I am thankful to Koen Dekkers and Guido Aerts for making this assignment possible. From the start, they gave me the opportunities and confidence to successfully complete my graduation assignment. Within the monthly meetings, they were great discussion partners and gave me helpful feedback. In particular, I want to thank my company supervisor, André van Houten. The openness and willingness of André to help me made me feel very comfortably during the past six months.

Next, I want to thank Engin Topan for his main supervision from the University of Twente. His knowledge about inventory management made me choose the right directions of my research. After meetings, I often had new energy and motivation to dive into my research. The flexibility and frequent meetings were very helpful in the process of this research. Furthermore, Dennis Prak has added great value in the short period of being my second supervisor from the University of Twente. This gave me new insights at the last stage of my research.

Lastly, I want to thank my friends, family, and my fellow students with whom I have studied the past years. Furthermore, I hope you will enjoy reading this thesis.

Gijs van Sambeek
Utrecht, July 2024

Management Summary

This research is conducted at NS which holds a central position in the Dutch transportation sector. The supply chain operations (SCO) department manages the supply and storage of train spare parts and sets up the logistics chains for this purpose. NS wants to improve its inventory process for seasonal spare parts. This group of spare parts has a yearly higher demand in a specific season compared to the demand in the rest of the year. NS believes that these spare parts could have a negative influence on the inventory performance. The current inventory process for these spare parts is outdated and lacks appropriate models. Instead, it heavily relies on interpretations and assumptions which could be be structured by policies. The core problems at SCO are the lack of using specified forecast methods for seasonal spare parts and using inventory policies that do not incorporate changes over time. These lead to the main research question:

"How can the forecasting methods and the inventory policies of seasonal spare parts be changed to improve the performance of the SCO department at NS?"

Currently, SCO identifies 59 spare parts as seasonal. These parts show seasonal variation in either winter or summer. However, this list needs a reclassification because of unclear seasonal demand patterns, the lack of demand, and unknown lead times. The forecasts of spare parts are often manually adjusted to cover seasonal demand. However, the performance of the forecasts including manual adjustments is worse than the forecasts excluding manual adjustments because the supply chain planner often forecasts the demand too high. Furthermore, the network service level performance, also known as fill rate, is above the norm of SCO, while the average train units waiting on materials KPI is below the norm. However, if we let the fill rate get worse, the average train units waiting on materials KPI would also get worse. Another service level measure, like the order line fill rate, would be more appropriate as the shortage of one spare part within an order causes a delay of the complete order. The order line fill rate will be lower than the fill rate and therefore the order line fill rate would represent the average train units waiting on materials KPI better.

For reclassification of seasonal spare parts, we use multiple regression to determine whether the demand pattern of a spare part is correlated to the months of the year. It is a time-efficient method to identify the presence of seasonal variations in comparison to an exploratory data analysis. We find that the demand pattern of 88 spare parts is correlated to the months of the year. Only nine out of the 59 original seasonal spare parts reappear in this list. Furthermore, we forecast demand using eight different forecast methods with the use of growing-window forward validation. We interpret the results for each demand class. The following conclusions are drawn based on the forecast performance measure sMAPE:

- Smooth: The Croston forecast method performs the best,
- Intermittent: The (S)ARIMA forecast method performs the best,
- For erratic demand, the SES forecast method performs the best,
- For lumpy demand, the (S)ARIMA forecast method performs the best.

When we compare forecast performances, we can conclude that our model performs better than the current forecast process. Several forecast methods like the moving average, Croston and (S)ARIMA score better than the current forecast performance. Therefore, we recommend to keep on testing all of our forecast methods and using them according to the corresponding performance.

Also, we control the inventory using the (s, (n)Q)-policy as this is the most used policy within this field of research based on literature research. The current policy uses a variable lot size and a fixed lot size is preferable as spare parts could come in boxes or on pallets. After simulating the inventory model for several years, the performance stabilizes with an average €846,654.80 of inventory costs and a service level of 96.6%. When we compare inventory performances with the current situation, our inventory model performs better on the average stock value and the usage/value ratio and performs closer to the service level KPI target. This reduces costs by 28.3%. Therefore, we conclude that aiming towards the target service level with our model is better than the current process.

We recommend identifying seasonal spare parts twice a year just after the summer and winter seasons. Currently, there is no fixed identification process and the seasonal spare part list is not up to date. Using our seasonal identification tool will help to analyse a big population of spare parts at once and is an addition to the current procedure. Additionally, we recommend implementing the performance measures for the forecast and inventory process of seasonal spare parts. This way supply chain planners can see in the long term which forecast methods perform well for which kind of spare parts and how this effects the inventory service level and costs. A supply chain planner will learn from mistakes as they have insights into their performance.

Contents

Chapter 1

Introduction

The first chapter of this thesis starts with describing the company and motivation of the research. This is followed by the section about the problem statement, addressing the specific challenges the research company faces. Additionally, this chapter presents the research design; describing the research questions, explaining the research methods, and setting the research scope. The last section describes the structure for the remaining of the thesis. This chapter sets the foundation for understanding the company's context and the approach taken in this research to address the main research.

1.1 Company Introduction

The Supply Chain Operations department (SCO) of the Nederlandse Spoorwegen (NS) wants to address a critical challenge in spare part inventory management for train components. NS holds a central position in the Dutch transportation sector, serving as the primary passenger railway operator since 1938 (Nederlandse Spoorwegen, 2024). With an extensive network of railway lines connecting cities and regions, NS provides rail services for travelers. The SCO department manages the supply and storage of train spare parts and sets up the logistics chains for this purpose.

The motivation for this project arises from NS's goal to improve its inventory process for seasonal spare parts. This group of spare parts has a yearly higher demand in a specific season compared to the demand in the rest of the year. The current inventory process for these spare parts is outdated and lacks appropriate models. Instead, it heavily relies on interpretations and assumptions which could be be structured and automated by policies. NS believes that the demand pattern of seasonal spare parts means that these spare parts could have a negative influence on the inventory performance when handled incorrectly. How much it influences the inventory performance is not known as the performance is not measured. Overall, NS wants to ensure enough seasonal spare parts in stock to cover the demand, while keeping its costs low.

1.2 Problem Statement

This section dives into the challenges SCO currently faces. First, this section gives an objective view on the problem context from the researcher's insights. The second and last subsection elaborates the core problem. The core problem is the main problem this research investigates.

1.2.1 Problem Context

It is important to get an objective understanding of the problem context because directly trying to solve the given problem could cause overlooking aspects. By means of interviews with various stakeholders of SCO, we put together what is already known in the organisation. We present these results in a list of issues which connect by cause-and-effect relationships. Linking the known issues leads to a problem cluster. Differently to other models, like for example a fishbone model, a problem cluster shows the connections between problems at a glance providing the core problem instantly (Heerkens and van Winden, 2017).

Figure 1.1 shows the problem cluster with core problems (red boxes), and other problems (orange boxes) given the context of seasonal spare parts at SCO. The selected core problems are the problems that this thesis tries to solve. Note that this figure only identifies the current problems and does not differentiate whether any problems are larger than others. For instance, having a shortage of two spare parts could be a bigger problem than having an excess of two spare parts in inventory. However, the figure does not indicate which situation is a bigger problem.

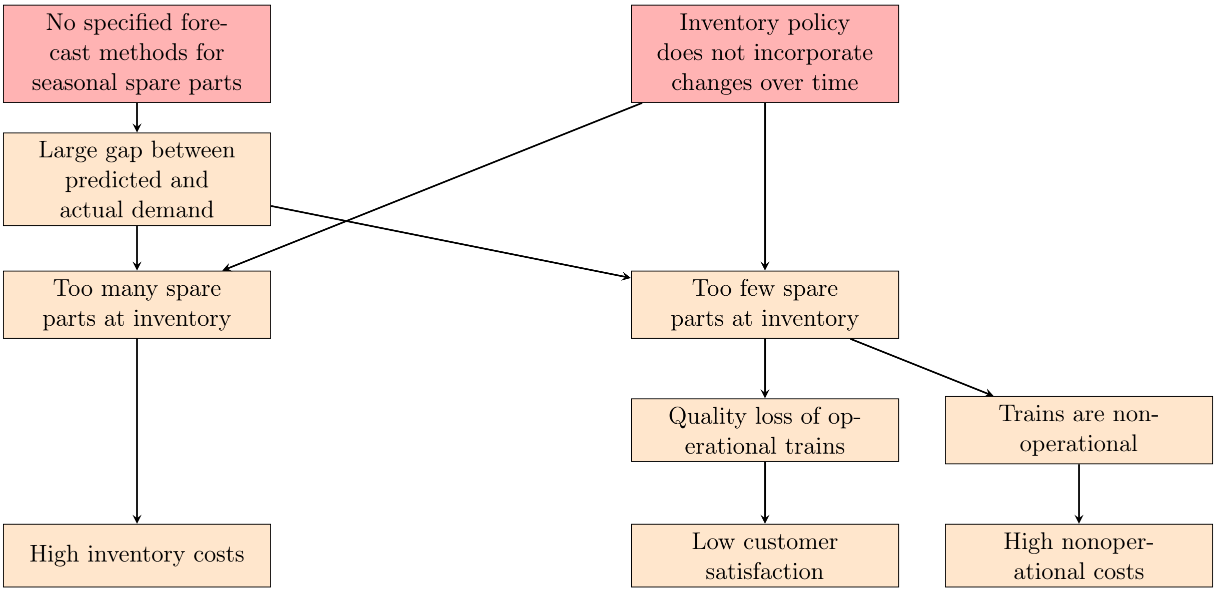

Figure 1.1: Problem Cluster of Spare Parts Forecasting and Inventory Control at NS

There are three problems that indicate the other problems at SCO which are the high inventory costs, the low customer satisfaction, and the high nonoperational costs. The three indicative problems lead back to two core problems. These core problems are the lack of using specified forecast methods for seasonal spare parts and using inventory policies that do not incorporate changes over time. Forecasting accurately is important as this leads to less uncertainty within the inventory process. Less uncertainty can help solving the indicative problems. To give more context about the second core problem: supply chain planners do not frequently review parameters of the policy. Currently, these are reviewed once every few years. Since we assume that solving these two problems within the given time frame of this research is achievable, this thesis will investigate two core problems.

1.2.2 Core Problems

The core problems *"no specified forecast methods for seasonal spare parts"* and *"inventory policy does not incorporate changes over time"* are derived from action problems. The definition of an action problem is a discrepancy between norm and reality as perceived by the problem owner (Heerkens and van Winden, 2017). Here, the problem owner is NS, or more specifically, the problem owner is the department SCO. SCO lacks measurements for the forecast process. Essentially, this is already something that has to change as we can not substantiate the problems with numbers. However, the forecast performance indirectly contributes to the performance of the inventory control which SCO does measure.

SCO quantifies the performance of the inventory control of all spare parts using the key performance indicator (KPI) called BWOM. This is an abbreviation in dutch for *"bakken wachtend op materialen"* which means train units waiting on materials. This KPI measures the average number of train units waiting on materials. When spare parts are out of stock and needed for maintenance, the BWOM KPI increases and scores worse. The longer the spare part is out of stock, the longer the train unit will be waiting on materials. Every month, SCO evaluates the reality and norm in their KPI reports. Table 1.1 presents the reality (the performance in 2023) and norm of the BWOM. The reality and norm do not indicate a big gap. However, this KPI measurement is for all SCO spare parts. SCO lacks direct insight in this KPI performance for seasonal spare parts. NS has the sense that seasonal spare parts have a relatively big share in this number. We will present our forecasting and inventory control performance measures for seasonal spare parts in Section 2.4.

TABLE 1.1: Reality and Norm of Inventory Control KPI of Spare Parts

KPI	**Reality**	**Norm**	**Difference**
BWOM	10.9	10.0	+ 0.9

Solving the core problems will possibly lead to improvements on the BWOM KPI which will contribute to the long term goals of SCO. To achieve the norm, the following main research question holds:

> *"How can we change the forecasting methods and the inventory control policy of seasonal spare parts to improve the performance of the SCO department at NS?"*

The intended deliverables for solving the main research questions are prototype forecast and inventory control tools, but also an implementation plan for using these tools.

1.3 Research Design

To structure solving the core problems, we use the Managerial Problem-Solving Method (MPSM) of Heerkens & van Winden. This method is a systematic problem-solving approach. It is applicable to various problems and takes a problem into account embedded in the context of an organisation (Heerkens and van Winden, 2017). The MPSM is a seven-step approach visualized in Figure 1.2.

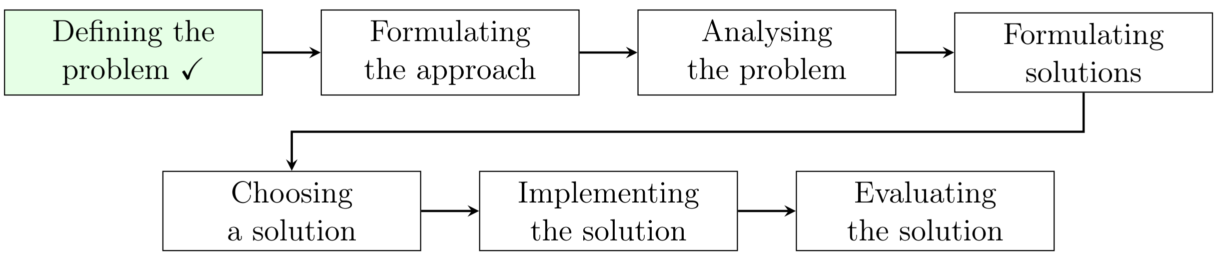

FIGURE 1.2: MPSM Flow Chart - Step 1 Completed

The identification of the core problem completes the first step of the MPSM. The next step of the MPSM is formulating the approach. We present the research approach based on the remaining steps of the MPSM. This section describes the corresponding research questions and data collection methods of each step. Solving these research questions results in the answer to the main research question.

1.3.1 Analysing the Problem

The third step of the MPSM is analysing the problem. This step gives an overview of the current situation at the SCO department which helps understanding the context. To analyse the problem, the first research question and corresponding sub-questions are:

1. What does the current forecast process and inventory control of SCO look like?
   (a) What are the operations playing a role at SCO?
   (b) What are the different features of spare parts managed by SCO?
   (c) What spare part classification does SCO use?
   (d) What methods are currently employed for forecasting and inventory control?
   (e) What is the current forecast and inventory performance?

Information from oral semi-structured interviews and data analysis provide the answer to these research questions. The interviews involve stakeholders of the forecast process and the data analysis uses historical and current data. This step creates a starting position for testing possible solutions.

1.3.2 Formulating Solutions

The next step of the MPSM is formulating solutions. This step involves literature research since various possibilities of improvement should be investigated. The second research question and additional sub-questions are:

2. What are possible forecast methods and inventory policies for the spare parts to apply in the context of SCO?
   (a) What classification methods are present in literature for spare parts management?
   (b) How is seasonality identified in spare part forecasting?
   (c) Which forecasting methods are recommended for forecasting seasonal spare parts?
   (d) What inventory control policies are present in literature?
   (e) In what ways can forecast methods be integrated with inventory control policies to optimize spare parts management?

The outcomes of a literature study contribute to solving the second research question.

1.3.3 Choosing a Solution

After formulating possible solutions, one solution has to be chosen. The possible solutions obtained in the previous step are compared which result in fitting forecast methods and inventory policies. The third research question is:

3. What are the best fitting forecast methods and inventory policies for the seasonal spare parts at SCO?
   (a) What spare parts do we identify as seasonal spare parts according to data analysis?
   (b) How do we classify our seasonal spare parts?
   (c) What forecast methods are most effective for seasonal spare parts?
   (d) What inventory control policy is best suited for seasonal spare parts?
   (e) How can we validate the used methods?

Consulting stakeholders of this research and applying the outcomes of the literature research are steps to take for building a quantitative model which will help choosing the best fitting forecast methods and inventory policies. We split up this step in two chapters as it involves the model design and comparing the outcomes.

1.3.4 Implementing and Evaluating the Solution

The final steps of the MPSM are formulated in two research questions but combined in one chapter. The corresponding research question to the implementation of the solution is:

4. What should SCO do to implement the forecast and inventory tools?

The implementation of a solution into the operations of a department should go carefully. In the end, it impacts a lot of people in their daily operations So, it is of high importance to include the stakeholders throughout the improvement and implementation process. To answer this question, we will write an implementation plan.

The corresponding research question to the evaluation of the solution is:

5. How should SCO evaluate the new forecast methods and inventory control policies?

To prevent a situation where the solution is outdated, SCO should be able to evaluate the solution. When the performance of the chosen solution decreases, SCO should notify it and should have appropriate measures ready. When this is included in the solution, the solution will be robust to changes which will improve the quality of the solution.

1.3.5 Research Scope

The research investigates the current performance of seasonal spare part demand forecasting and inventory control. Chapter 2 analyses information about the current identified seasonal spare parts. Later research in Chapter 4 reclassifies seasonal spare parts to see whether SCO classifies them correctly. The seasonal spare parts that we identify are the input for the rest of the research. Furthermore, the study intentionally excludes supplier management and multi-echelon inventory management from its scope. Historical data from 2019 up to the start of 2024 serves as input for this research. Setting these boundaries makes solving the problems more tangible and achievable within the time constraint of this research.

1.4 Report Structure

The structure of this thesis is as follows: Chapter 2 gives an overview of the existing situation within NS's SCO department answering the first research question. Subsequently, Chapter 3 presents a study of the relevant literature answering the second research question. For answering the third research question, Chapter 4 provides the developed model where Chapter 5 shows a numerical comparison of the solutions. Chapter 6 describes the plan for implementing and evaluating the model, giving answer to the fourth and fifth research question. Finally, the report concludes with Chapter 7, drawing the conclusions, giving recommendations, and explaining the future research and limitations. This last chapter gives an answer to the main research question.

The end of this chapter completes step 2 of the MPSM, formulating the approach. Figure 1.3 shows the progress of the research so far.

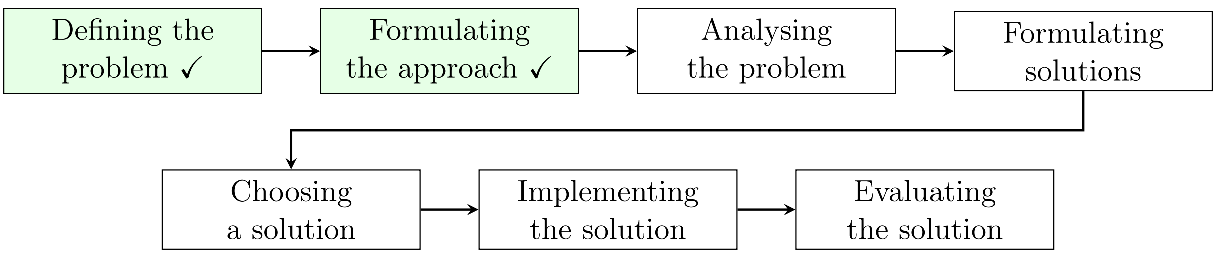

FIGURE 1.3: MPSM Flow Chart - Step 2 Completed

Chapter 2

Description of the Current Situation

This chapter presents the findings to answer the first research question (RQ) of this thesis.

> RQ1: *"What does the current forecast process and inventory control of SCO look like?"*

To solve the research question and corresponding sub-questions, this chapter describes the current operations within the Supply Chain Operations (SCO) department of the Nederlandse Spoorwegen (NS). Furthermore, it gives an elaboration and analysis of the spare parts that SCO deals with. Additionally, the chapter explains the current forecasting and inventory controlling processes and shows the performance of these. The exploration of SCO's current operational procedures and corresponding performances is critical, as it will be used as a benchmark to assess potential improvements. This helps evaluating the efficacy of proposed changes within the department.

2.1 Supply Chain Operations

NS primarily focuses on the transportation of passengers from one destination to another. Trains are the cornerstone of the NS's operations. In simple terms, if the trains aren't working, NS cannot do its job properly. Therefore, it is essential to ensure that these trains are operational for as much time as possible. This means that maintenance is important and especially the part of getting the right amount of spare parts on the right place on the right time. This is where SCO plays the main role.

SCO is responsible for the inventory and storage of spare parts and designs the logistic chains of the maintenance operations. When a train needs maintenance, the goal is to handle this as quick as possible while minimising costs. Within SCO, the organisation is divided by train series. This means that certain teams work on the introduction of new train series, some teams are dedicated to maintaining and guaranteeing train series, and other teams manage the phasing-out train series in a responsible manner.

SCO distinguishes two types of maintenance which are corrective and preventive maintenance. Corrective maintenance is reactive; when a component of a train fails, NS needs to perform maintenance which is unplanned. Preventive maintenance is a proactive measure to prevent failures. It involves replacing functioning components to avoid unexpected breakdowns. To maintain sufficient inventory, SCO must forecast the demand for spare parts resulting from corrective maintenance but also consider the demand for spare parts resulting from preventive maintenance. The demand of these two types of maintenance is segregated, resulting in separate forecasting of corrective and preventive maintenance. Preventive maintenance occurs according cyclical maintenance planning and involves less uncertainty.

2.2 Spare Parts of SCO

Availability of resources is important to make sure trains that need maintenance are up and running as quick as possible. There should be enough spare parts present to repair the train. On the other hand, there should not be too many spare parts since this involves more inventory costs. Three important variables are the product group, seasonality, and criticality of spare parts. The first three subsections of this section describe and explain these variables. The last subsection explains the custom classification that SCO uses for inventory parameters.

2.2.1 Product Group

SCO divides spare parts into two product groups: exchange parts and wear parts. The difference between these product groups is the life cycle. Exchange parts are parts that are not discarded, but these parts are repaired after being used in a train. These parts are repaired because this is technologically possible and economically beneficial. SCO calls unused parts *clean* and used parts *dirty.* Once an exchange part is repaired, the dirty part becomes a clean part which means it can be used again in future replacements. One remark: According to SCO, on average 7% of the exchange parts are beyond repair after replacement and have to be discarded. Wear parts are parts that are always discarded after use. Figure 2.1 shows the flow charts of both exchange and wear parts in case of performing corrective maintenance.

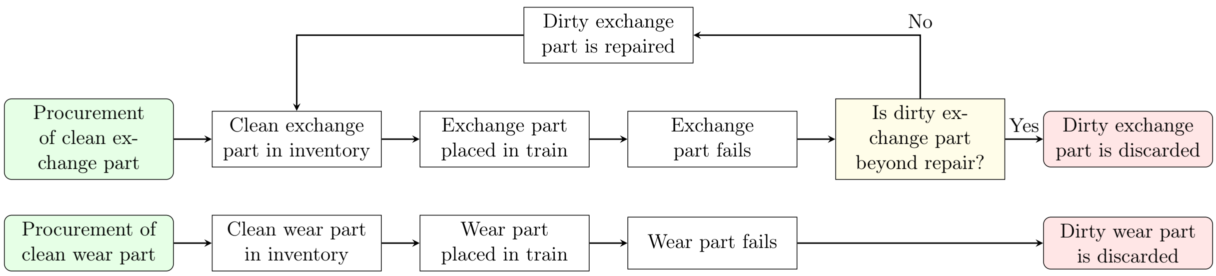

Figure 2.1: Life Cycle of Exchange Part (top) and Wear Part (bottom) in Corrective Maintenance

Just as the treatment of exchange and wear parts differs, so does the inventory management of these parts. The chain size is corresponding to the number of exchange parts within it. The chain size can increase with incoming orders and can decrease when exchange parts are beyond repair. However, the chain size is not seen as the exchange part inventory level as it consists of clean and dirty parts. The dirty parts cannot be used for replacements. Only clean parts can be used for this purpose, thus representing the exchange part inventory level. The exchange part inventory level increases when a new order comes in or when a repaired part comes in, and it decreases when parts are used for replacement. Inventory management for wear parts is a simpler concept as it does not use a chain size. The wear part inventory level increases when a new order comes in and decreases when parts are used for replacement. Overall, the exchange parts require fewer replenishments because dirty parts can be repaired and reused.

To provide some figures, SCO handled 30,958 exchange or wear spare parts at the start of 2024. 15.0% of the spare parts are exchange parts, while 85.0% of these are wear parts. The classification of the product group of a spare part is not fixed. Due to price reductions of exchange parts, it can become economically beneficial to label them as wear parts. Contrarily, if the price of a wear part increases, it may be more beneficial to treat them as exchange parts as long as this is technologically possible.

2.2.2 Seasonality

Every spare part has its own monthly demand. Extreme weather conditions in the summer and winter can increase the chance of a spare part breakdown. Spare parts sensitive to this are labeled as seasonal spare parts. SCO has identified 59 seasonal spare parts. These are defined with use of the technical knowledge of reliability engineers and with use of the historical data provided by the supply chain planners. Section 6.1 explains in more detail how these spare parts are identified.

We look at the difference in demand patterns between a seasonal and nonseasonal spare part. Figure 2.2 visualises the different kind of spare parts and their trendline. Based on the evaluation of demand graphs of all 59 seasonal spare parts, spare part FD332067 is the spare part with the clearest seasonal pattern. Spare part FA500278 is just a arbitrary nonseasonal spare part with a similar size of demand. The figure shows that every year, spare part FD332067 has higher demand during the winter months than during the summer months compared to the trendline. Spare part FA500278 also has variations of demand. However, these do not show that during either the summer or winter months there is a higher demand than in the other season compared to the trendline.

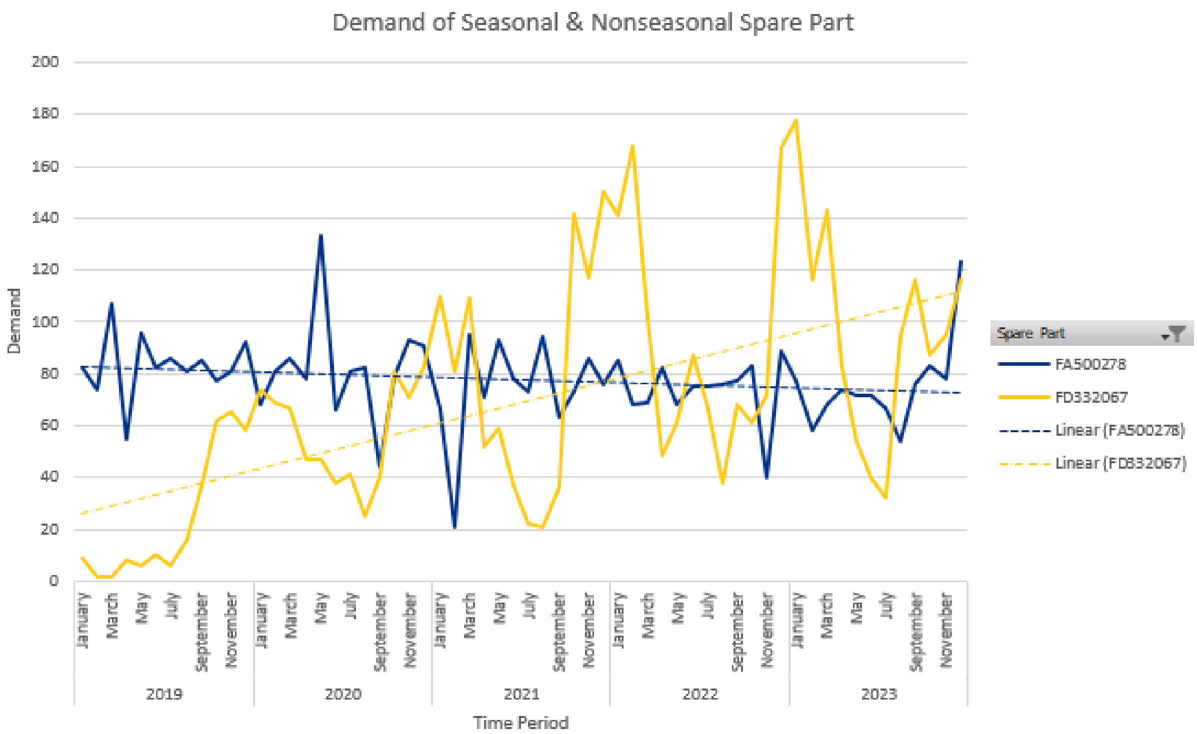

FIGURE 2.2: Demand of a Seasonal (FD332067) and Nonseasonal (FA500278) Spare Part over Time

Appendix A presents two other demand graphs of seasonal spare parts. Here, we will show an example of a seasonal spare part which is prone to failures in the summer months. Also, we show the demand pattern of a seasonal spare part which does not have a clear seasonal demand pattern. By evaluating the demand graphs of seasonal spare parts, we conclude that most spare parts do not have a clear seasonal demand pattern, only six spare parts do.

Figure 2.3 and Figure 2.4 illustrate the distribution of demand and planned lead times (respectively) of the seasonal spare parts using box plots. The yellow bar represents the interquartile range, which accounts for the middle 50% of the data points. The whiskers indicate the last data point within 1.5 times the interquartile range below the 25th percentile or above the 75th percentile of the data points. Any data points outside the whiskers are considered as outliers.

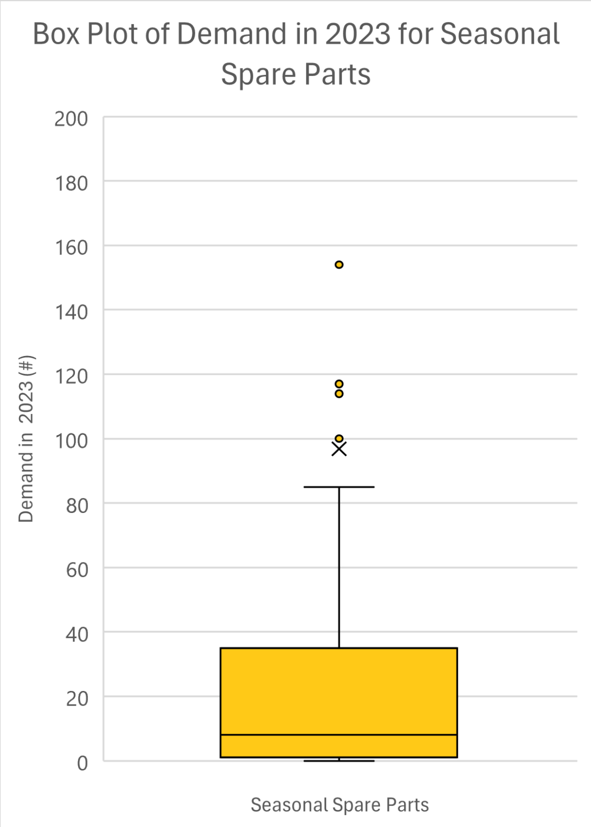

FIGURE 2.3: Demand Distribution of Seasonal Spare Parts

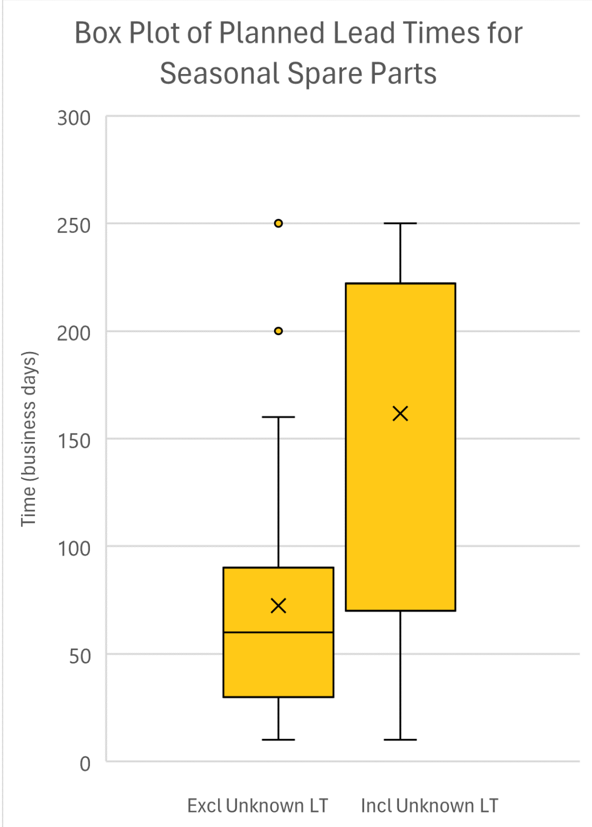

FIGURE 2.4: Planned Lead Times Distribution of Seasonal Spare Parts

The box plot of Figure 2.3 indicates a positive skewed distribution of seasonal spare part demand. This reveals that most of the seasonal spare parts have low demands as the median is below 10 spare parts in 2023. It is contradictory that there are spare parts selected as seasonal when they have an annual demand close to zero. These are spare parts that should be removed from the seasonal spare part list as it does not make sense to monitor these spare parts closely in advance of a season when they have this little demand. A few spare parts have such a high demand rate in 2023 that the average demand (the X in the figure) surpasses the maximum whisker. This shows the variation of seasonal spare parts and indicates to need to monitor individually as they are different. Note that not all outliers from the demand of the seasonal spare parts are visual in Figure 2.3 as they exceed 200. Including these outliers in the box plot visualisation would make the box plot unreadable.

The box plot of Figure 2.4 shows two box plots. A lot of seasonal spare parts do not have a known planned lead time. Only 27 spare parts have a known planned lead time which means that there are agreements with the supplier about time until delivery. In reality, these lead times can differ but SCO assumes these as deterministic. When the lead time is unknown, SCO assumes a lead time of 222 business days, hence the second box plot in the figure. This is a safe assumption as it exceeds the upper whisker from the known lead time box plot. The fact that more than half of the seasonal spare part population do not have a known lead time indicates that the process is outdated. Assuming a lead time of 222 business days leads to high safety stocks and a lot of inventory costs. Having known lead times for all spare parts could decrease these safety stocks and corresponding costs. Note that it looks like there is no median value in the box plot including unknown lead times. However, the median is the same value as the 75th percentile, which is at 222 business days making the median disappear. Furthermore, there are no outliers.

Managing the seasonal demand variations is important for optimizing inventory levels and reducing costs. By accurately identifying the seasonal demand patterns, we can explain the seasonal variation and decrease the safety stocks. This approach will save costs while ensuring enough inventory. The fact that a lot of current seasonal spare parts do not have a clear seasonal demand pattern, do not have demand, or do not have a known lead time, indicates the need to update this classification.

2.2.3 Criticality

The last variable to consider is criticality, which is based on safety and logistics. Criticality affects the current classification and inventory parameters. Some parts are critical for the safety of a train. When these parts fail, the train is not allowed to travel any further with passengers due to safety regulations. Failures in parts that are not critical to safety might affect passengers' comfort and the customer satisfaction, but do not pose a danger. Additionally, several spare parts could be critical from a logistic standpoint. When such a part fails, the train cannot operate logistically, meaning it cannot drive at all. These two variables are independent from each other and important because these say something about the consequences of a breakdown of a spare part. Subsection 2.2.4 uses this variable to assign inventory parameters.

Figure 2.5 shows two bar charts visualising the critical seasonal spare parts. For 17 seasonal spare parts, the criticality based on safety is undetermined. SCO assumes these as safety critical unless proven differently. Therefore, we consider 28 seasonal spare parts to be safety critical. For four seasonal spare parts, it is unknown whether they are logistic critical. SCO assumes these spare parts to be logistic critical to be on the safe side as logistic critical spare parts are better monitored. This means in total 52 of the 59 seasonal spare parts are logistic critical which is the vast majority. When comparing the safety and logistic aspect, there are only two seasonal spare parts which are neither safety nor logistic critical. These are spare parts FA702049 and MD805059.

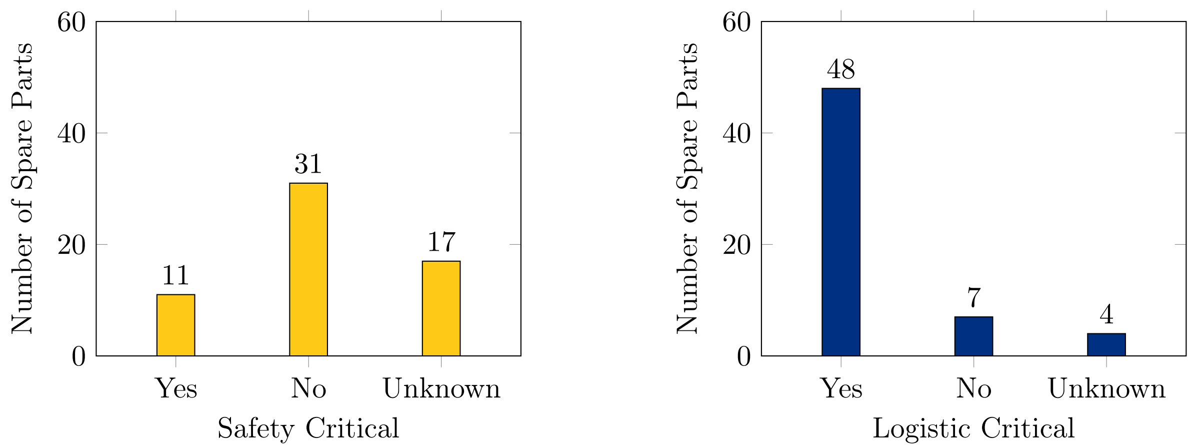

(A) Classification based on Safety Critical

(B) Classification based on Logistic Critical

FIGURE 2.5: Critical Classification Bar Charts

2.2.4 Custom Classification

SCO uses a custom classification for their spare parts which they update quarterly. The classification of a spare part determines some of the inventory parameters like target service level (TSL) or lot size. This classification is based on whether the spare part is an exchange or wear part, the usage (displayed as letter), the value (displayed as number), and whether it is logistic critical or not. So, the custom classification is four-dimensional.

The letters of the classification range from A until D based on the spare part's usage of the past 18 months with the following thresholds:

- Class A: if $demand > 17$
- Class B: if $2 < demand \leq 17$
- Class C: if $0 < demand \leq 2$
- Class D: if $demand = 0$

The numbers of the classification range from one until three based on the spare part's value. This classification differs for exchange and wear parts as they have other thresholds. Table 2.1 shows the difference between the exchange and wear part classification.

TABLE 2.1: Spare Part Value Classification Thresholds in Euros

Spare Part	**Class 1**	**Class 2**	**Class 3**
Exchange Part	$value \leq 999.99$	$999.99 < value \leq 3,499.99$	$value > 3,499.99$
Wear Part	$value \leq 49.99$	$49.99 < value \leq 499.99$	$value > 499.99$

When only looking at the letter and number classification, it leads to twelve different classes for spare parts. Table 2.2 shows the number of seasonal spare parts per class. Most seasonal spare parts are fast movers with a high value as most spare parts are in classes A2, A3, B2, and B3.

TABLE 2.2: Number of Seasonal Spare Parts per Class (Letter+Number Classification)

Usage \Value	**1**	**2**	**3**
A	4	13	9
B	1	8	10
C	2	3	2
D	3	1	3

When looking at the custom four-dimensional classification, it leads to 48 different classes. Figure 2.6 shows the number of seasonal spare parts per class in the custom classification. Currently, there are no exchange parts that are not logistic critical, so this is excluded from the figure. Only 23 of the 48 custom classes are used in the selection of seasonal spare parts.

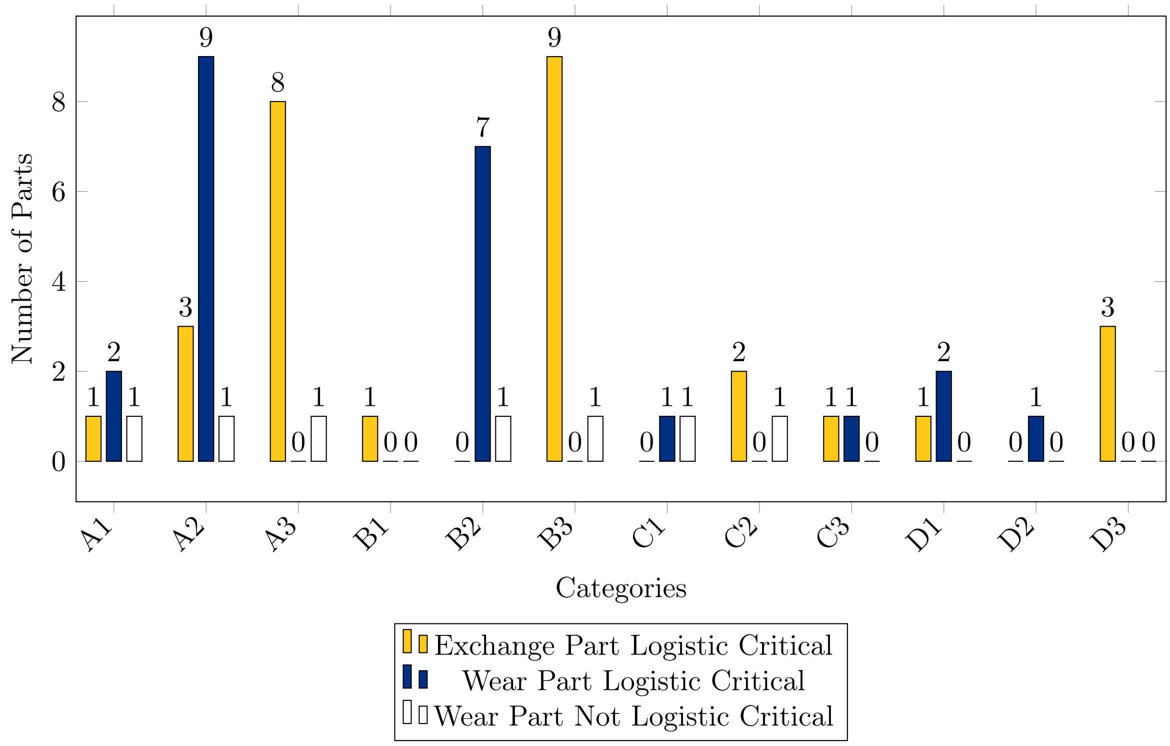

FIGURE 2.6: Number of Seasonal Spare Parts per Class (SCO Custom Classification)

Figure 2.7 shows the influence of the classification on an inventory parameter, the TSL. Here, the TSLs of C1 & D1, C2 & D2, and C3 & D3 are the same. This means that there is no differentiation between a slow- and non-moving (C and D respectively) spare part regarding the TSL. Also, the custom classification gives spare parts with higher values a lower TSL. It is cheaper to reach a higher service level for cheap spare parts as it is cheaper to keep more inventory for these parts.

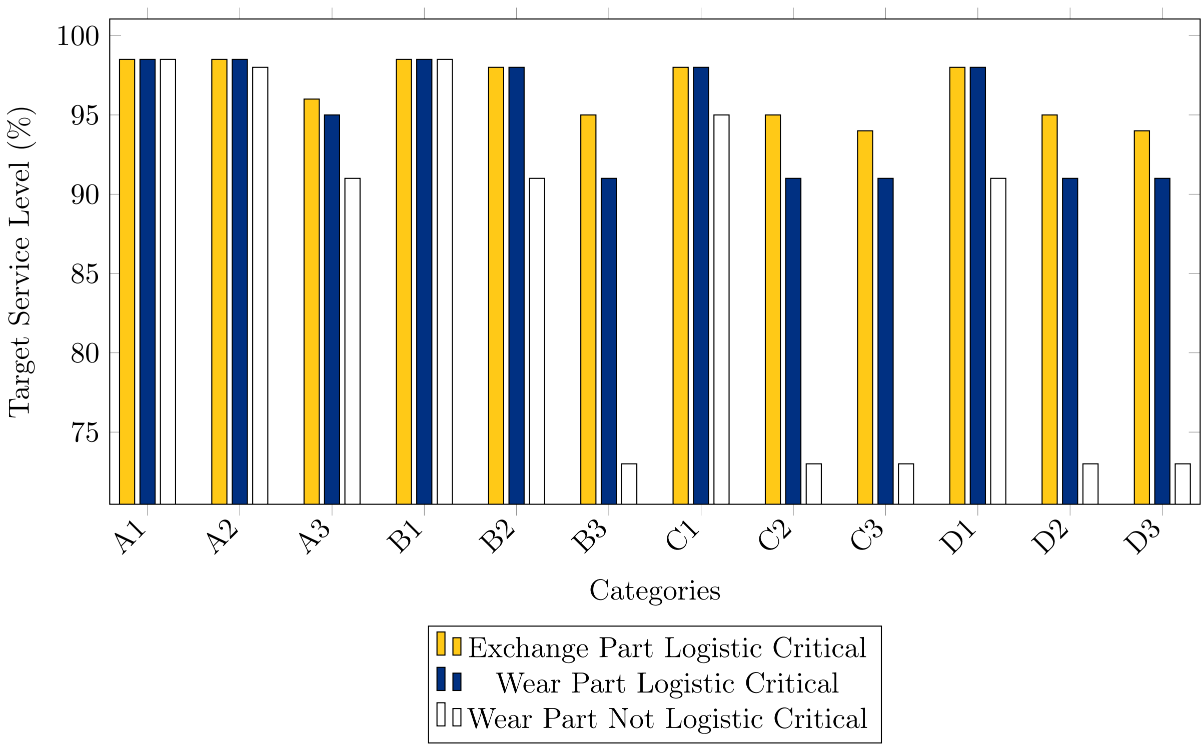

FIGURE 2.7: Target Service Level per Class (SCO Custom Classification)

There are other factors that influence some of the inventory parameters. However, these are not influencing the TSL and therefore not relevant.

2.3 Forecasting and Inventory Control

The forecasting and inventory control processes are very important for SCO. If these are of high quality, this means that there is enough inventory for maintenance while the costs are reduced as much as possible. Figure 2.8 shows the echelon structure of the inventory locations at NS. Logistics Center Tilburg (LCT) is the highest echelon level. The second level consists of the *"onderhoudsbedrijven"* (OB), or the maintenance companies, located in four different locations. The *"servicebedrijven"* (SB), or the service companies, represent the lowest echelon level. In total there are 30 service companies spread throughout The Netherlands, each associated with a specific region and one OB. This thesis focuses on managing the demand and inventory of the entire system. The inventory and demand of specific locations are not considered in this research as explained by Subsection 1.3.5.

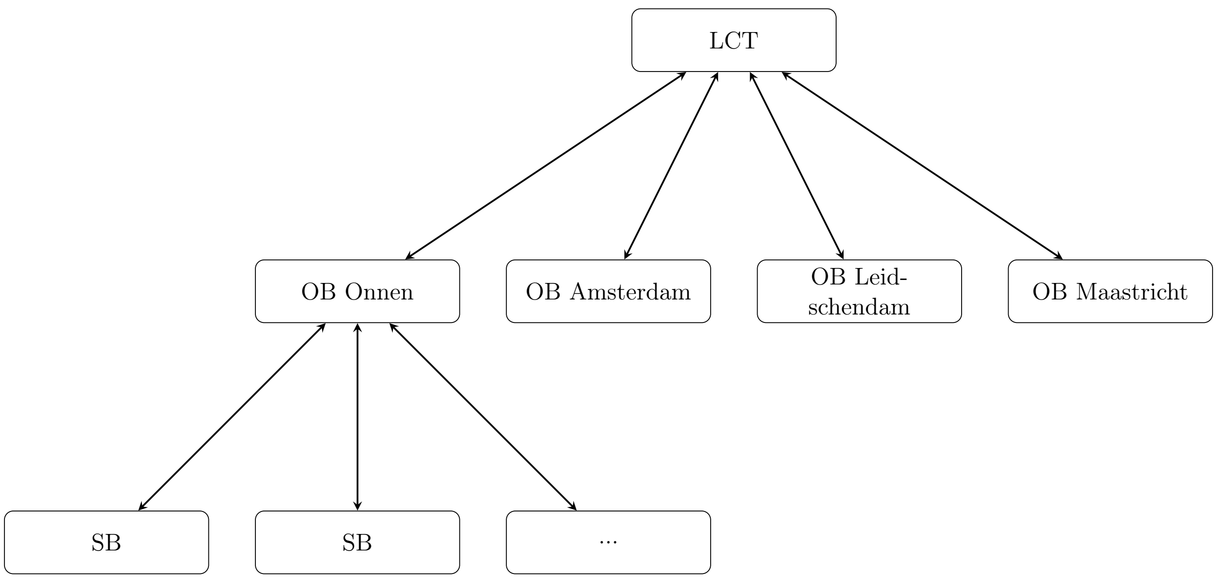

FIGURE 2.8: Echelon Structure and Material Flow of Maintenance Locations

As Chapter 1 identifies the core problems in the forecast and inventory control, the forecast and order processes are where improvements can be made. NS uses the software Servigistics XelusParts for their inventory control. This program generates monthly forecasts based on historical data and implements inventory policies. As Servigistics XelusParts will be replaced by another program in the near future, this research is not bounded to the limitations of the program. To give an idea how the process of forecasting and ordering looks like, Figure 2.9 visualizes these. Here, SCP is referring to the supply chain planner and ME to the maintenance engineer.

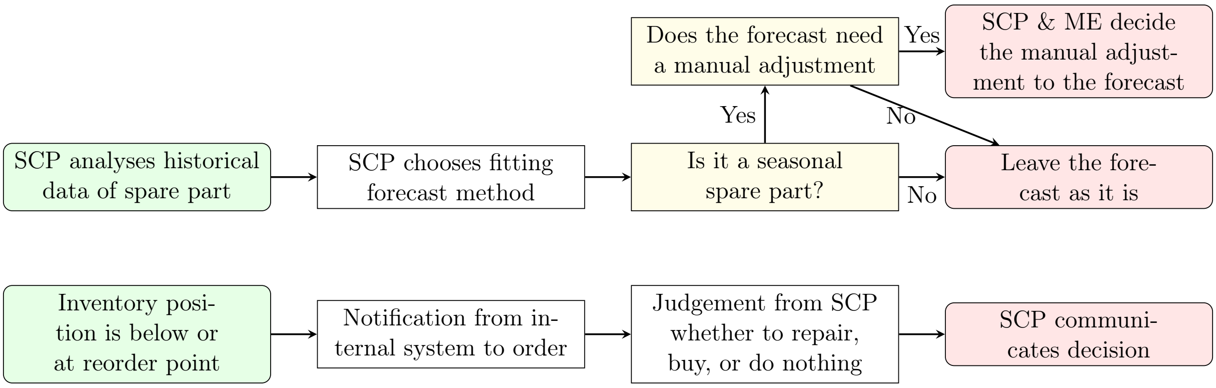

FIGURE 2.9: Flow Charts of Forecast Process (top) and Order Process (bottom)

Figure 2.9 shows that the SCP analyses historical data and chooses a fitting forecast method for a specific spare part, but there is no specified policy to follow. This has lead to the situation that the SCPs look at only four methods out of the fourteen possibilities of Servigistics Xelus-Parts and do not look at any parameters as these are unknown for them. The process of choosing a suitable method involves altering between forecast methods and assessing which would fit the best based on the forecast visualisations of these methods. When a spare part is labelled as seasonal, the process can take a different route. Due to a lack of knowledge about using forecast methods that account for seasonal variations, there is an evaluation step in the forecast process where the planner can manually adjust the forecast to meet the seasonal demand. Figure 2.10 illustrates the number of seasonal spare parts per forecast method.

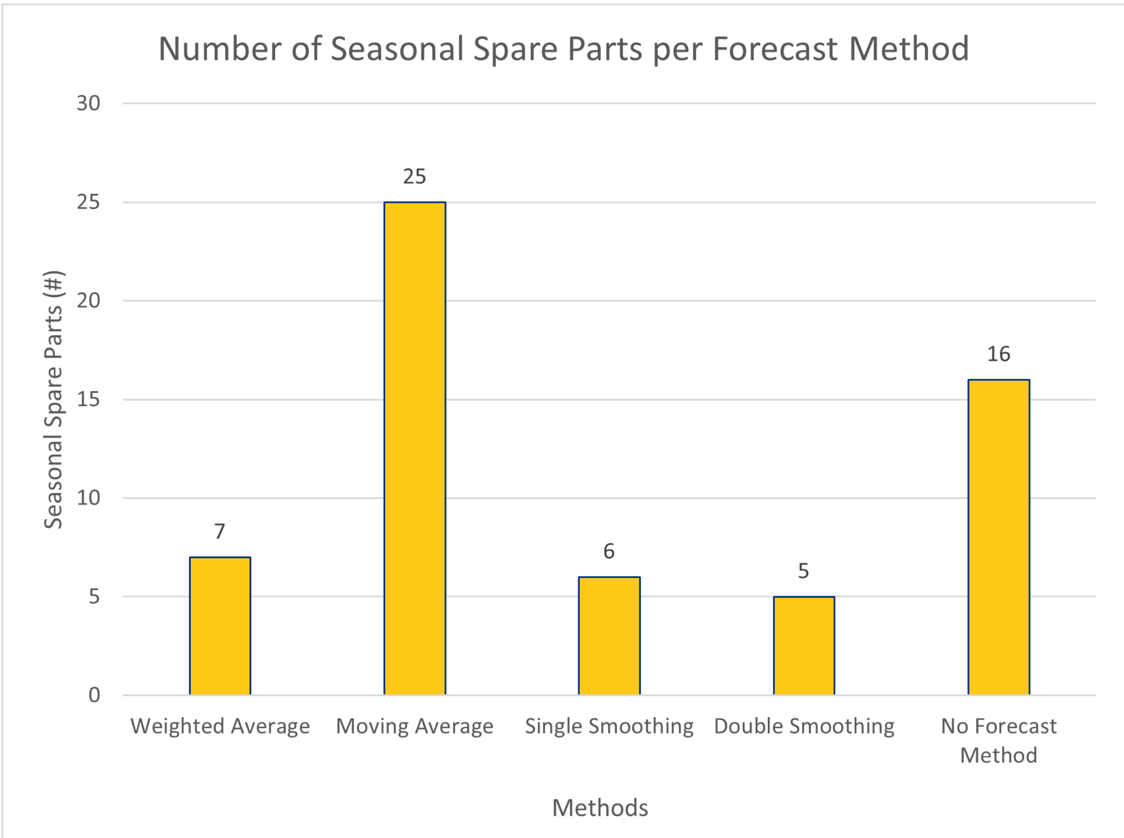

FIGURE 2.10: Number of Seasonal Spare Parts per Forecast Method

SCPs use four forecast methods that only include the level and trend for forecasting the demand of the selected spare parts. Chapter 3 explains the methods from the figure. None of these forecast methods include a seasonality factor which is contradictory when forecasting seasonal spare parts. Figure 2.10 also shows that sixteen seasonal spare parts do not have a selected forecast method. This means their forecast is zero every month. This overview of current forecast methods could indicate that the selected seasonal spare parts do not follow a seasonal demand pattern or the forecasts are manually adjusted to cover the future demand. But, it could also be the case that some of these methods cover the demand well enough.

SCO considers safety stock for determining the reorder point. The safety stock is calculated with use of TSL and lot size. This safety stock is added on the demand over lead time to find the reorder point per spare part. SCO uses inventory control in Servigistics XelusParts with continuous review. When the inventory position is below or at the reorder point the internal system suggests to order for the total of a certain amount of months demand. Ordering a quantity based on a certain amount of months demand is called a periodic order quantity (POQ). The value corresponding to the POQ is dependent on the custom classification but only based on the letter + number classification, shown in Table 2.3. There is no difference between the C and D classification just like in Figure 2.7. The POQ is low for expensive spare parts which makes sense as a low POQ results in a low holding costs and expensive parts have higher holding costs than cheap parts. The POQ first decreases, then increases when going from class A, to B, to C. There is no logic behind this, indicating the need for an update. Furthermore, the POQ is often lower than the lead time as most spare parts have a lead time of 222 business days, which translates to 11 months. This could be a problem when allowing one order at a time. Overall, changing the inventory parameters does not happen proactively, only when there is a reason to change them. For example during the Corona pandemic in 2020, these parameters were manually adjusted.

TABLE 2.3: POQ in months of Seasonal Spare Parts per Class (Letter+Number Classification)

Usage \ Value	**1**	**2**	**3**
A	6	5	4
B	4	2.5	2
C	6	3.5	2
D	6	3.5	2

The SCP can decide to take over the decision to order or to do something else as described in Figure 2.9. For example, the SCP looks at the minimum order quantity (MOQ) of the part. If the internal systems advises to order less parts than the MOQ, the supply chain planner possibly chooses to increase the order quantity. Or, the SCP could choose to wait to order that part and order it only when the advice to order exceeds the MOQ.

If a mechanic wants to perform maintenance and cannot replace a part because there are no spare parts in inventory, it is called a stockout. When there is a stockout, it could be that the order is already underway or that no order has been placed yet. In either case, SCO communicates with the supplier to deliver their goods as quickly as possible.

2.4 Current Performance

As Subsection 1.2.2 describes, SCO does not use performance measures in the current forecast operations. Therefore, we will analyze the performance in this section ourselves. The literature in Subsection 3.2.8 explains various measures for forecast accuracy. We compare the historical demand and forecast with and without manual adjustment based on the previous four years and calculate the performance. Table 2.4 and 2.5 present the forecast performance which is a starting point for the further analysis of possible interventions. Here, we calculate the performance as an average performance of all seasonal spare parts. It is the higher the measure, the worse the performance for the symmetric mean absolute percentage error (sMAPE). In Subsection 3.2.8, we also show the calculation for this measure.

TABLE 2.4: Average Forecast Performance of Seasonal Spare Parts without Manual Adjustment

Measures	2020	2021	2022	2023
sMAPE	53.74%	61.68%	63.24%	57.27%

TABLE 2.5: Average Forecast Performance of Seasonal Spare Parts with Manual Adjustment

Measures	2020	2021	2022	2023
sMAPE	67.57%	65.68%	64.55%	66.44%

Table 2.4 shows a significant better average forecast performance than Table 2.5. At first glance, this does not seem logical. When adjusting the forecast based on expert experience the expectation is that a forecast improves. However, the performance tables prove different. An explanation of these results could be that the SCP often forecasts the demand too high just to be on the safe side of availability of spare parts. The bias performance confirms this explanation as the manual adjusted performance is constantly positive indicating that on average the forecast is too high while the other performance is fluctuating around zero. Figures 2.11 and 2.12 also show this behavior. The SCP recognises seasonal demand peaks and acts according these by manually adjusting the forecast. This is common behavior according to Hewage et al. (2022). Many organizations use human judgement when forecasting for contextual information like sales promotions, weather, and seasonal conditions. In our case, this behavior results in an overestimation of demand in almost every time period.

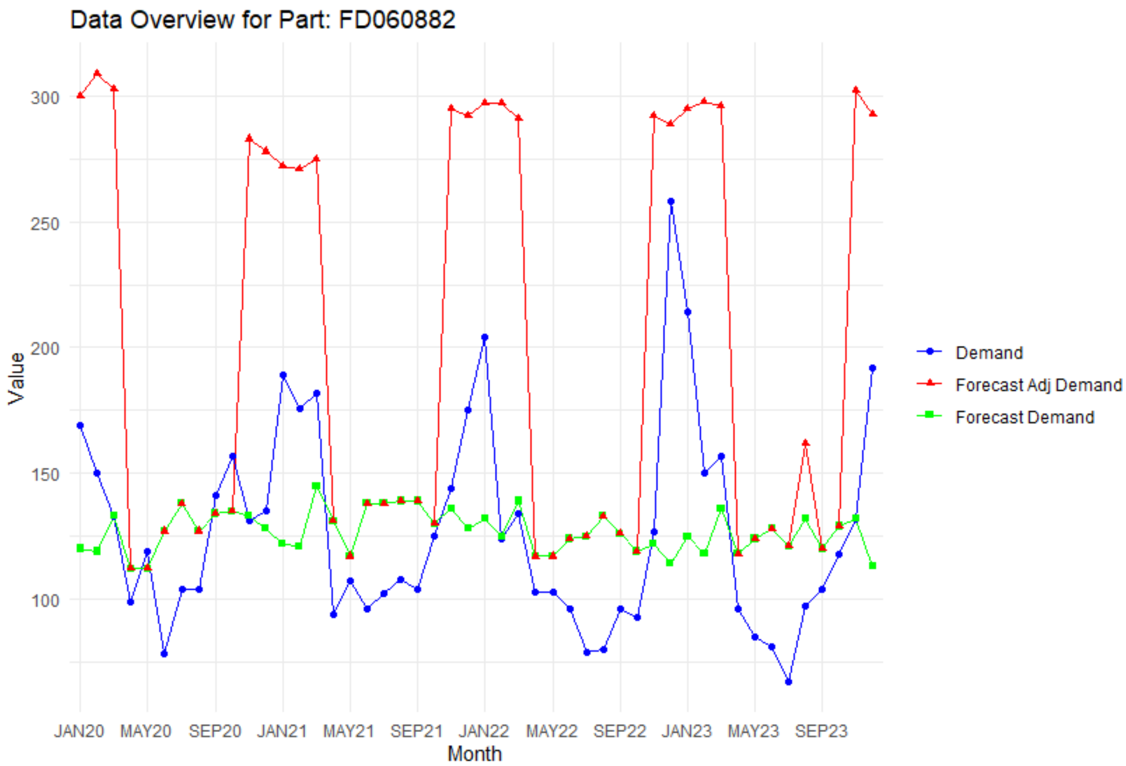

FIGURE 2.11: Demand and Forecast Pattern of Part FD060882

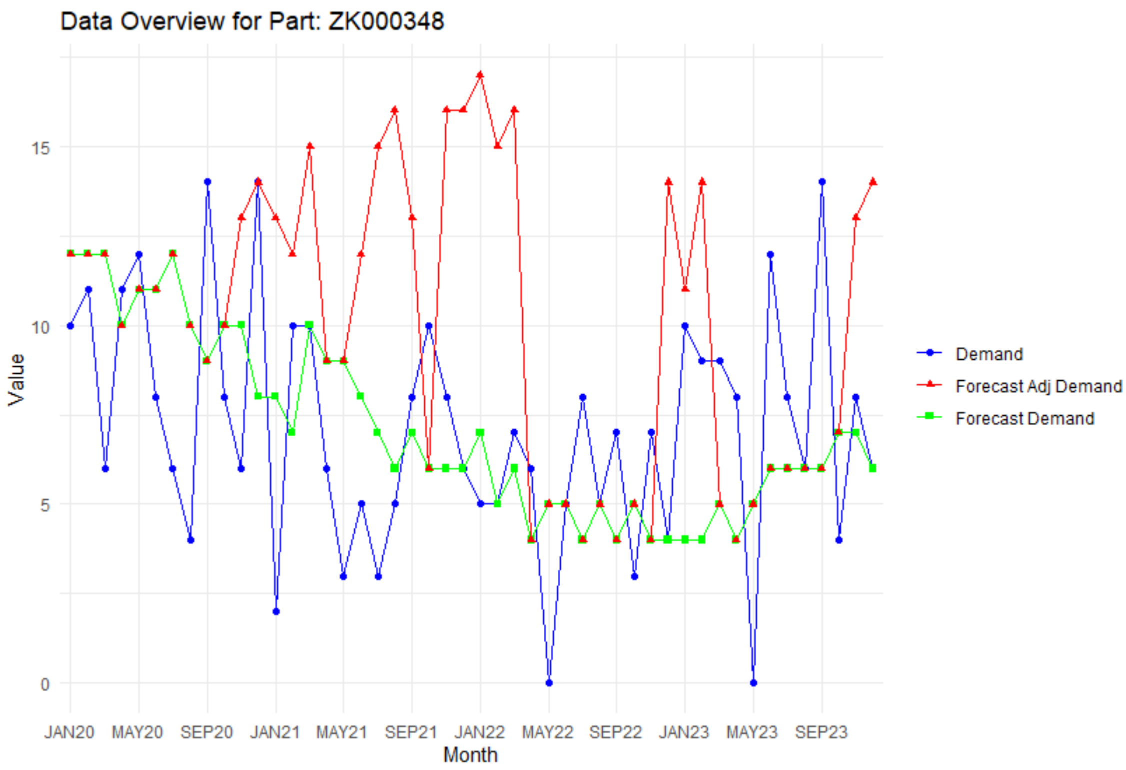

FIGURE 2.12: Demand and Forecast Pattern of Part ZK000348

We evaluate the performance of the inventory control with the KPIs network service level, usage and the stock value of spare parts. The network service level, also known as fill rate, is the service level of the whole inventory network of NS. Every time a spare part is needed and there is no inventory on any location of NS, this KPI decreases. Table 2.6 and 2.7 show the inventory KPIs of 2021 until 2023. There is no data from 2020 or earlier.We expected that the network service level of seasonal spare parts was lower than the network service level of all spare parts as this is a difficult group of spare parts according to SCO. The network service level for seasonal spare parts is 99.9% in 2023 as there have been four stockouts with a usage of 2,775 spare parts. This service level is higher than the network service level of all spare parts with a usage of 120,185 spare parts. The usage represents both corrective and preventive maintenance. As the target service level is 97.0%, these service levels indicate that the stock levels of the seasonal spare parts are often too high because the supply chain planners predict their seasonal spare parts rather safely than risky. This does also mean that the service level can become worse, while we observed that the BWOM should improve according the norm in Table 1.1. As both these KPIs are related to each other, this is contradictory. Another service level measure, like the order line fill rate, would be more appropriate as the shortage of one spare part within an order causes a delay of the complete order. Furthermore, the usage to stock value ratio indicates that inventory for all spare parts is twice as cost-efficient as for seasonal spare parts. This indicates the need for improved inventory management for seasonal spare parts.

TABLE 2.6: Inventory Performance of Seasonal Spare Parts

Inventory Control KPIs	**2021**	**2022**	**2023**
Network Service Level	99.8%	99.8%	99.9%
Total Usage	2,758	3,053	2,775
Average Stock Value	€3.6M	€3.2M	€3.5M
Usage/Value	766.1	848.1	792.9

TABLE 2.7: Inventory Performance of All Spare Parts

Inventory Control KPIs	**2021**	**2022**	**2023**
Network Service Level	99.6%	99.6%	99.7%
Total Usage	104,961	112,422	120,185
Average Stock Value	€82.0M	€83.4M	€82.5M
Usage/Value	1,280.0	1,347.9	1,456.8

2.5 Conclusion of the Current Situation

This chapter answers the first research question of this thesis.

> RQ1: *"What does the current forecast process and inventory control of SCO look like?"*

First, we identified the current operations within the SCO department of NS. These operations are the supply of spare parts and execution of corrective and preventive maintenance. These operations are divided by train series.

Furthermore, we analysed the spare parts that SCO deals with. There are exchange and wear spare parts that differ in inventory management. Currently, SCO labels 59 spare parts that show seasonal variation in either winter or summer. These spare parts require more attention in advance of the seasons for forecasting and inventory control. Besides, spare parts are classified critical when a breakdown means that trains are not allowed to drive any further. Based on these variables, SCO uses a custom classification for its spare parts based on the demand, value, product group, and criticality. This custom classification determines several inventory parameters. The seasonal spare part analysis showed that the current selection is not up to date and needs a reclassification. The most evident arguments are the unclear seasonal demand patterns, the lack of demand, and the unknown lead times.

Additionally, the chapter explained the current process of forecasting and the inventory control of spare parts. The forecasting process consists five different forecast methods which is often selected automatically with use of the forecast program and do not include seasonality. So, when SCO considers a spare part as seasonal, the SCP often manually adjusts the forecasts to incorporate the seasonal variations. The inventory control uses a safety stock, reorder point, continuous review period, POQ, and MOQ which are not updated proactively.

Lastly, we calculated the performance of these processes. The performance of the forecasts including manual adjustments is worse than the forecasts excluding manual adjustments. This is probably because the supply chain planner often overpredicts demand to be on the safe side of availability of spare parts. The network service level, also known as fill rate, performance of the inventory control for the seasonal spare parts is similar to the performance for all spare parts. Probably, the average stock value can be decreased while keeping a high network service level. Furthermore, the network service level is above the norm of SCO and the BWOM is below the norm. This is contradictory as these two KPIs are related to each other. Another service level measure, like the order line fill rate, would be a more appropriate service level as the shortage of one spare part within an order causes a delay of the complete order.

The findings of this chapter give direction to the literature study. We identify the important aspects classification, identification of seasonality, forecasting, and inventory control. Performing a literature review on these aspects creates an overview of the possibilities for our model.

The end of this chapter completes step 3 of the MPSM, analysing the problem. Figure 2.13 shows the progress of the research so far.

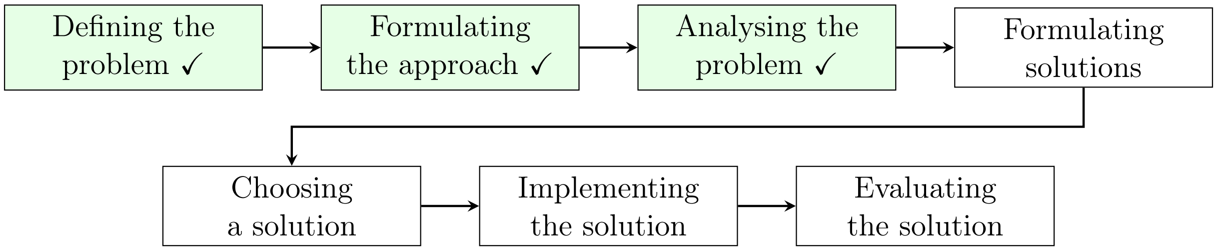

FIGURE 2.13: MPSM Flow Chart - Step 3 Completed

Chapter 3

Literature Study

The third chapter presents a literature study to answer the second research question (RQ) of this thesis.

> RQ2: *"What are possible forecast methods and inventory policies for the spare parts to apply in the context of SCO?"*

To solve the research question, this chapter describes the findings of the relevant literature. First, it will elaborate classification of spare parts. Next, it will describe several forecast methods and the way to identify seasonality. Lastly, it continues with an explanation of inventory control and how to measure the performance. This chapter will be a tutorial based literature study to understand how a problem in the context of this research can be solved.

3.1 Spare Part Classification

The aim of inventory management is to determine the best fitting forecast method and inventory policy which result in the most accurate predictions for stock keeping units (SKUs) (Heinecke et al., 2013). In the context of this thesis, SKUs represent spare parts. The choice of these methods or policies is partially determined based on the spare parts' characteristics (Heinecke et al., 2013; Millstein et al., 2014). This section describes classification by means of different methods.

3.1.1 ABC-Classification

Silver et al. (2016). suggest using a distribution by value curve for an ABC-classification of spare parts. Figure 3.1 shows this curve representing the cumulative percentage of yearly usage value and the cumulative percentage of the total number of spare parts in inventory.

The ABC-classification of spare parts divides spare parts into three categories. Class A represents the most important spare parts. These spare parts, often 20% of the spare parts with the highest yearly usage value, account for 80% or more of the total yearly usage value. These parts need the most management attention. This involves reviewing decision parameters frequently and determining precise values of control parameters. Class B are the spare parts with a medium importance, usually 30% of the total spare parts, accounting for 15% of yearly usage value. These parts often have a computerized inventory control with management-by-exception (Van der Heijden, 2022). Class C is the group of spare parts that are the least important. This group is relatively large (50%) for the proportion of yearly usage value it represents (5%). Typically, most companies keep a large inventory for class C spare parts to minimize the possible difficulties caused by stockouts of these spare parts (Silver et al., 2016).

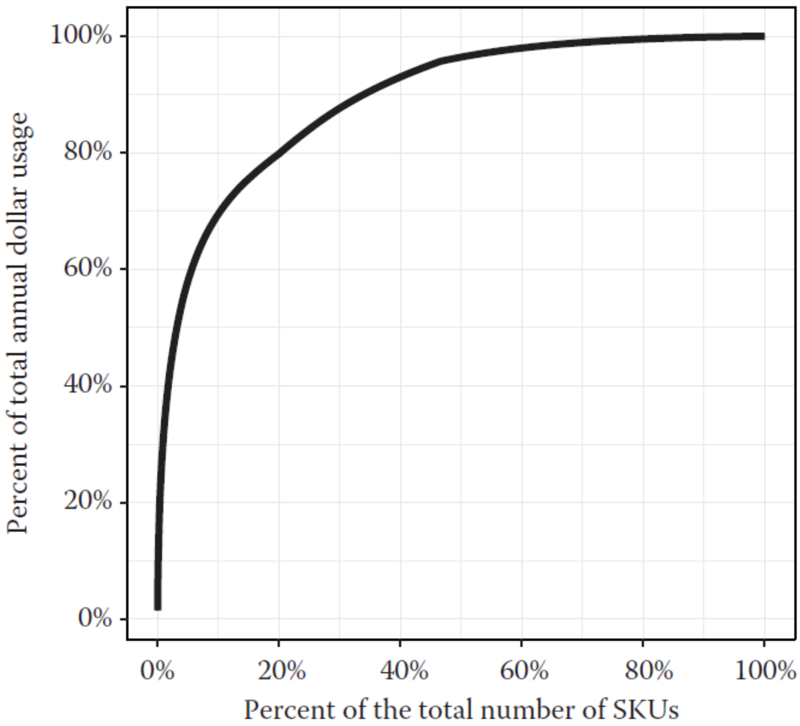

FIGURE 3.1: Distribution by Value Curve of a Fictional Situation (Silver et al., 2016)

Despite the benefits of using ABC-classification, Dhoka (2013) states that the classification method has a limitation. Critical spare parts that are low in consumption value may be overlooked as they could be classified as C-item. Silver et al. (2016) propose a solution by shifting spare parts that are crucial for some operations manually to another class.

3.1.2 XYZ-Classification

Another way of classifying spare parts in inventory is based on the demand uncertainty of spare parts, also known as the demand predictability (Bhalla et al., 2021). The XYZ-analysis consists of calculating the coefficient of variation (Dhoka, 2013). Similar to the ABC-classification, XYZ-classification defines three groups of spare parts according an increasing coefficient of variation. Spare parts with a uniform demand and the lowest coefficient of variation are classified as class X. Class Y represents spare parts with varying demand and moderating fluctuations. Spare parts with an irregular demand are classified as class Z. This means that spare parts with the highest coefficient of variation are in class Z (Nowotyńska, 2013). An example of dividing the three different classes can be for class X a coefficient of variation of 0-10%, for class Y a coefficient of variation of 10-25%, and for class C a coefficient of variation larger than 25%. Of course, these boundaries are not strict as each industry differs, so adjusting the boundaries is acceptable. Figure 3.2 shows a fictional example of how the demand pattern of spare parts relates to the possible classes of the XYZ-classification.

Dhoka (2013) identifies a drawback for the XYZ-classification. New spare parts are often classified as class Z because for these spare parts, it could take some time to establish a stable demand. Another relevant drawback is that the XYZ-classification overlooks seasonality. Seasonal spare parts could be removed from this analysis.

Combining classification methods can address some of the limitations as the limitation of a classification method fades when combining it with another method. Combining the ABC- and

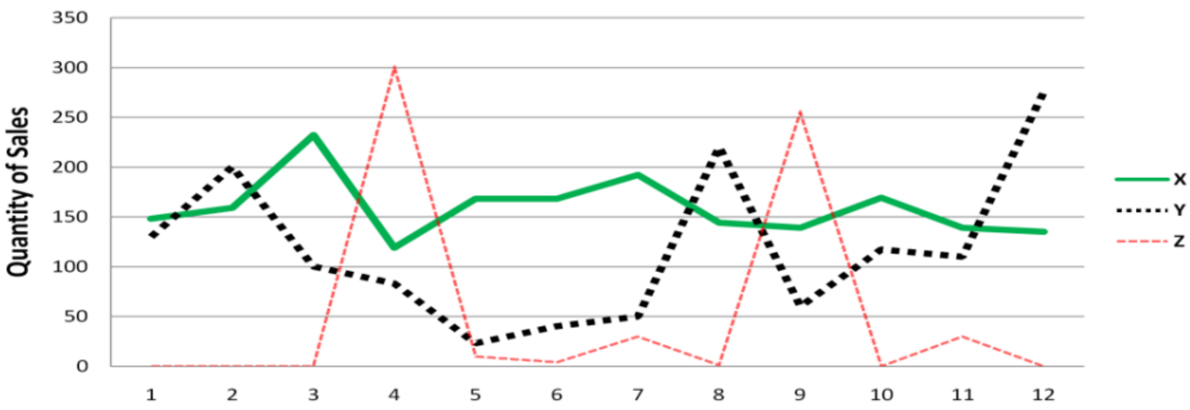

FIGURE 3.2: XYZ - Demand over Time and Classification based on Coefficient of Variation (Dhoka, 2013)

XYZ-classifications creates a more effective 2-dimensional approach for classification (Dhoka, 2013). With the combination of ABC- and XYZ-classification, it becomes possible to establish service level targets for each class. Class AX, being the most crucial and stable, should have the highest service level target since the service level raises quickly when increasing the service stock. On the other hand, class CZ, being the least important and stable, should have the lowest service level target (Van der Heijden, 2022).

3.1.3 Other Classification Methods

Mobarakeh et al. (2017) suggest that the most widely accepted spare parts demand classification method is based on a two-dimensional matrix using the average inter-demand interval (ADI) and coefficient of demand variation (CV). The classification regions are smooth, intermittent, erratic and lumpy with cutoff values of $ADI = 1.32$ and $CV = 0.49$.

- Smooth: Both low ADI and demand variation,
- Intermittent: High ADI and low demand variation,
- Erratic: Low ADI and high demand variation,
- Lumpy: Both high ADI and demand variation.

Smooth demand has the most constant demand. It has the most potential to be forecasted accurately. Many forecast methods are suitable for this demand class. Based on a study for intermittent demand of aircraft's spare parts, exponential smoothing and Croston forecasting methods outperformed other forecasting methods like the weighted moving average and Winters' forecast method (Mobarakeh et al., 2017). Not many articles suggest forecast methods for erratic demand. However, one paper states that linear forecast models like exponential smoothing and Croston are not suited to capture nonlinear dynamics and uncertainties while machine learning models are more suitable methods (Jiang et al., 2017). Forecasting lumpy demand is hard and not very accurate. However, the best approaches are the weighted moving average and Croston forecast methods (Regattieri et al., 2005).

Bhalla et al. (2021) write about the FSN-classification which is an abbreviation for fast-moving, slow-moving, and non-moving and is a classification method based on the demand pattern. Fast-moving means that a spare part has a high demand per time unit, while non-moving means it does not have any demand. The FSN-classification looks like the ABC-classification as it is also based on the criteria demand but has other classification thresholds. Furthermore, it is not different and therefore not relevant for an in-depth review.

All mentioned classification methods are based on the spare part characteristic demand. However, spare parts can also be classified based on their characteristics like value, stockout costs, lead times, criticality, or supply uncertainty (Bhalla et al., 2021). Especially, criticality seems like a relevant characteristic to use in the classification the spare parts of NS as we saw in Subsection 2.2.3. Another classification method, based on more subjectiveness, is the VED-analysis. This is an abbreviation for the vital, essential, and desirable analysis. Spare parts are classified in one of these classes based on consultation with experts. There are no quantitative factors contributing to the assessment of the class which makes the classification a difficult and less relevant (Bacchetti and Saccani, 2012).

3.2 Forecast Methods

The main reasons for an inventory control system to look ahead with forecasting are the presence of a lead time for ordering items and the need to order in batches instead of product for product. Approaches of forecasting are estimating based on historical data or based on other factors like demand of complementary items, sales campaigns, or the weather forecast (Axsäter, 2006). Currently, SCO uses a forecast based on historical demand. Figure 3.3 visualises a forecast framework using historical demand (Silver et al., 2016). The figure shows the components the mathematical model, the actual forecast and the forecast errors. These components are the main input of this section.

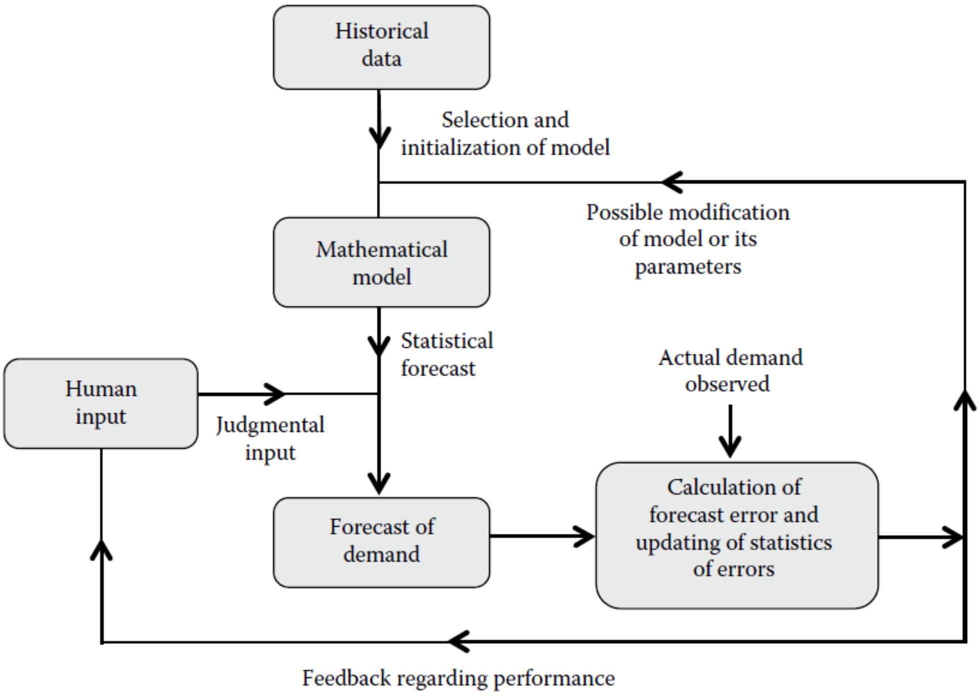

FIGURE 3.3: A Forecasting Framework (Silver et al., 2016)

3.2.1 Mathematical model

The mathematical model in the context of forecasting consists of five components. Level (a), trend (b), seasonal variations (F), cyclical movements (C), and irregular random fluctuations (ϵ). Level represents the scale of demand and trend is the rate of growth or decline of demand. Seasonal variations identify as the variations of natural forces or those resulting from human decisions. Next, cyclical variations are the ups and downs of economic activity resulting from business cycles. The last component is the irregular fluctuation which is the residue that remains after the effects of the other components are removed. This residue is important for inventory control which will be explained in Section 3.3. Equation 3.1 shows an example of a forecast formula using all explained concepts, where x_t represents the demand in period t (Silver et al., 2016). This example shows an additive model, but there are also multiplicative models where the parameters are multiplied instead of summed.

$$x_t = a + bt + F_t + C_t + \epsilon_t \tag{3.1}$$

A constant demand model has the least parameters to estimate as it only involves level. The trend-seasonal demand model has more parameters to estimate as it also involves trend and seasonal variations. The trend-seasonal demand model can become very specific due to its parameters and could cover a wide class of demand. Axsäter (2006) advises to use a simple demand model with few parameters unless there is evidence that a model with more parameters clearly shows certain advantages. This is why we look into identifying trend and seasonality in Subsection 3.2.7

There are three steps involved in the use of a mathematical model. First, the selection of a general form of the model. Next, the choice of initial values for parameters within the model. Finally, the use of the mathematical model and chosen parameters to make the forecast (Silver et al., 2016). This will be the structure of the following subsections that represent forecast methods.

3.2.2 Moving Average

The moving average forecast method has as demand model the level model: $x_t = a + \epsilon_t$. The idea of this method is that it uses the average demand of the last N observations as estimate for a. Reducing N gives more weight to the recent data as older historical data is not included anymore. If $N = 12$ (in months) seasonal variations have no effects on the forecast (Axsäter, 2006). Every included data point has the same weight however, it is possible to modify the moving average approach with different weights for each period. An example of such a method is the weighted average forecast method (Silver et al., 2016).

3.2.3 Single Exponential Smoothing

Exponential smoothing has also a level demand model with $x_t = a + \epsilon_t$. The difference is the updating procedure for the level coefficient. This procedure includes a smoothing constant α. An α close to the value one leads to a high forecast dependency of the recent demand. When it is close to zero, the level coefficient will not change much as the level coefficient of the previous time period has the smoothing constant $1 - \alpha$. The initialization works by using the average demand in the first periods as estimate for the level (Silver et al., 2016). If the period length is one month, it is common to use a smoothing constant alpha between 0.1 and 0.3. When the forecast is updated every week, a smaller alpha should be used (use $\alpha = 2/(N + 1)$). For small values of alpha, it can take long before forecasts are reliable as it converts slowly (Axsäter, 2006). Another way of selecting the smoothing factor is by doing a traditional grid search (Van der Heijden, 2022). This method divides the existing data into two subsets, uses one for the initialisation, and the other subset for forecasting. By maximising the forecast performance, it obtains a local optimal smoothing factor. This exponential smoothing method is also called the simple or single exponential smoothing method (SES) as it only includes a level coefficient (Maretania et al., 2021).

3.2.4 Holt Method

An extension on the single exponential smoothing is the inclusion of trend and is called the double exponential smoothing method, or Holt method (Maretania et al., 2021). This method uses the underlying demand model $x_t = a + bt + \epsilon_t$. For this model, updating accounts for both the level and trend coefficient in the same way as it does for single exponential smoothing. The least squares regression is used as a method to initialise the level and trend coefficients (Silver et al., 2016). This model has two coefficients and therefore also two smoothing constants, the α and β. Common values for monthly updates are $\alpha = 0.20$ and $\beta = 0.05$ (Van der Heijden, 2022). For small values of alpha or beta it can take long before forecasts are reliable as they convert slowly(Axsäter, 2006).

3.2.5 Holt Winters' Method

The Holt Winters' trend-seasonal model is also a form of exponential smoothing. As it is including level, trend and seasonality, it is a form of triple exponential smoothing (Omar and Kawamukai, 2021). The demand model of this method is $x_t = (a + bt)F_t + \epsilon_t$. The seasonal coefficient F_t has the length of P periods. When $P = 4$ (in a year), it means there are quarterly seasons. When $P = 12$, there are monthly seasons. All three coefficients each have similar updating procedures as for single or double exponential smoothing. For the initialisation of the coefficients, the model uses the ratio to moving average procedure. Step 1 is a rough estimation of level with moving average over P periods. Step 2 is estimating the seasonal factors. Step 3 is estimating the final level and trend coefficients with use of regression. Some typical values for the smoothing factors α, β, and γ are 0.19, 0.053, and 0.10 respectively (Silver et al., 2016).

3.2.6 Other Forecast methods

When a spare part follows an intermittent demand pattern, it has periods of zero demand and periods with positive demand, it could be hard to forecast with use of exponential smoothing. Exponential smoothing will react too slow for peaks in demand (Axsäter, 2006). The Croston method does differentiate between the size of positive demand and the time between two periods of positive demand. After each period with positive demand, the method updates the demand size and demand interval with use of exponential smoothing. One disadvantage of the method is that it has a positive bias because it over-forecasts the mean demand (Teunter et al., 2011). The forecast package in R combines the size of positive demand and the time between two periods of positive demand to create an average forecast R Core Team (2024). Lindsey and Pavur (2013) propose a way to include seasonality. This makes the Croston method relevant for intermittent seasonal spare parts. However, this option is not included in the forecast package in R.

In the previously described forecast methods, we assume independence of demand. However, it could be the case that this is not true and demand is correlated. A forecasting technique that handles correlated stochastic demand variations is the autoregressive integrated moving average (ARIMA) model. This model involves parameters for the autoregression order, difference order, and moving average order (Axsäter, 2006). The seasonal ARIMA (SARIMA) model includes seasonal components and involves double the number of parameters compared to ARIMA as every parameter also has a seasonal parameter. Also, it has a seventh parameter representing the number of time steps for a single seasonal period (Malki et al., 2022).

To see whether machine learning methods are suitable for demand forecast, Moroff et al. (2021) compares the statistical Holt Winters' method and an extended version of the SARIMA with the machine learning methods XGBoost, random forest, long-term short-term memory, and multilayer perceptron. They conclude that two factors must be considered in order to compare the methods. The forecast performance but also the implementation effort of the methods are relevant to compare the methods. The statistical methods do perform well but have a high implementation effort, while the machine learning methods are relatively easy to implement, but sometimes have a lower performance. Makridakis et al. (2020) also comment on the performance of different forecast methods. 40 years ago, the first forecasting competitions took place, also known as the M forecasting competition. In the first competitions, a simple exponential smoothing method outperformed ARIMA models. They also found that human judgement did not improve the accuracy of forecasting. In the latest competitions from 2020 onwards, the conclusions shifted towards the machine learning models outperforming all others.

The combination of multiple forecast methods is a well-established procedure to improve the performance of forecasting. For example, by taking the mean or median of multiple point forecasts. However, the mean of point forecast is sensitive for outliers and the median needs about 30 forecasts to function well (Barrow and Kourentzes, 2016).

3.2.7 Identifying Trend and Seasonality

Now, we have seen forecast methods including level, trend and seasonality. The inclusion of these factors requires more parameters to estimate and if there is no significant effect present, it will only cause more noise. There are several ways to identify trend or seasonality. The most simple method is with an exploratory data analysis using graphs or pivot tables. For example, deseasonalising demand with the moving average over a year and plotting this can help identifying trend. Or, plotting the demand per month for every year in a graph can help identifying seasonality. However, this involves creating visualisations and therefore requires visual recognition of trend or seasonality which is time-consuming (Van der Heijden, 2022).

There are also statistical test for identifying trend and seasonality. Use regression on a large dataset and check whether the slope is significant to identify the trend. A statistical approach to identify seasonality is using a regression model with dummy variables for each seasonal factor. First, remove the trend. Then, use linear regression models for the de-trended demand with seasonal dummy variables. At last, perform t- and f-tests to see whether the seasonal effects and model are significant (Van der Heijden, 2022). The statistical methods work in a more systematic way then the exploratory data analysis and are easier to use when considering many spare parts.

Seasonality might be deterministic and the seasonal factors can be modelled with use of seasonal dummy variables. However, seasonality might be non-stationary as it evolves over time. Assuming the one type when the other is dominant can lead to high bias, so testing and considering whether seasonality is deterministic or stochastic can lead to better forecasting performance. Darné and Diebolt propose the HEGY (Hylleberg, Engle, Granger, and Yoo) test procedure to determine the nature of seasonality (Darné and Diebolt, 2002). For the scope of the research, we proceed with the deterministic assumption as we assume spare parts are constantly sensitive for seasonality.

3.2.8 Forecast Accuracy

Measuring the accuracy of a forecasting process helps three different purposes. First, it helps identifying the height of safety stock. Additionally, it makes it possible to monitor the validity of the forecasting model and parameters. Lastly, the measurements will count as feedback to the forecasts (Van der Heijden, 2022).

The most common procedure to calculate the forecast accuracy within traditional forecast evaluation is the last block evaluation (Bergmeir and Benítez, 2012). Here, the available demand data is split in a train and test set. The train set is used to train the forecast model. This trained model will make predictions and this will be tested along the test set. A negative side effect of this simple approach is that there could be dependency in this procedure. Performing multiple validation enhances the robustness of the measured performance and prevents overfitting. Possibilities are cross-validation or forward-validation, where forward-validation performs better. An example is the growing-window forward-validation (Figure 3.4), using five validation iterations, each with a growing train set and a changing test set (Schnaubelt, 2019).

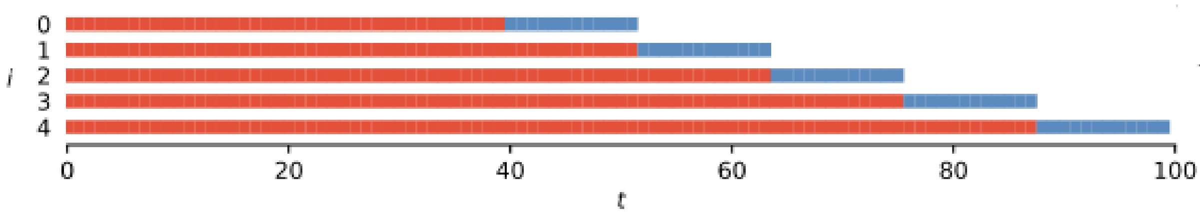

FIGURE 3.4: Growing-Window Forward-Validation Visualisation (Schnaubelt, 2019)

A common measurement of performance is the bias. If the bias value is close to zero, the forecast is unbiased which indicates a well balanced forecast. A high positive or high negative bias indicates that a forecast is constantly too high or low. The standard deviation of forecast errors gives an indication of the spread of forecast errors. The bias and standard deviation both give an indication of the performance but do not give a complete description that could help decide how to allocate resources (Silver et al., 2016). In contrary, the mean squared error (MSE), the root mean squared error (RMSE), the mean absolute deviation (MAD), the mean absolute percent error (MAPE) and the symmetric mean absolute percent error (sMAPE) are performance measures for variability that are widely used. The remarks on these methods are that the MAD is seen as more robust than the MSE as the MSE is sensitive to outliers. When aggregating the forecast performance measures MSE, RMSE, and MAD, it does not include the influence of demand size. When dividing these measures with the demand size, it is possible to aggregate. Dividing the MAD with the demand size creates the formula for the MAPE. However, the MAPE is not useful for low demand. Therefore, the sMAPE is the most appropriate performance measure ranging from 0-200%. The formula for the sMAPE is shown in Equation 3.2 respectively.

$$\text{sMAPE} = \frac{100\%}{h} \sum_{t=1}^{h} \frac{|\hat{x}_t - x_t|}{\frac{|x_t|+|\hat{x}_t|}{2}} \tag{3.2}$$

where:

- x_t is the actual demand in time period t
- $\hat{x}_t$ is the forecasted demand in time period t
- h is the forecast horizon

3.3 Inventory Control

Forecasting alone is not enough for inventory management. Controlling the inventory and designing how to replenish, are as important. Inventory control considers different ordering systems involving order size, reorder points, order up to levels, and safety stock.

3.3.1 Policies

When looking at different inventory control policies, there are policies with a continuous or periodic review, and policies with a fixed or variable order size.

The (s, Q)-policy uses a fixed lot size with continuous review. When the inventory position is at or below the reorder point s, the fixed lot size Q is ordered. Reorder point s consists of the expected demand during lead time plus the safety stock which should cover the uncertainty in demand during the lead time (Van der Heijden, 2022).

The (R, s, Q)-policy is the same as the (s, Q) policy, except it does not review the inventory position continuously but periodically with parameter R. There is more chance of a undershoot due to waiting for the next review period which should be included when calculating the safety stock (Van der Heijden, 2022).

The (s, S)-policy uses a variable lot size with continuous review. This means if the inventory position is at or below the reorder point s, an order is placed. The size of the order depends on the current inventory position since a lot size is used to reach the order up to level S. The safety stock should cover the uncertainty in undershoot plus the uncertainty in lead time demand (Van der Heijden, 2022).

The (R, s, S)-policy makes use of a variable lot size with a periodic review. This means that according this policy, every R periods the inventory position is checked. If the inventory position is at or below reorder point s, up to reorder level S is ordered. The (R, S)-policy also uses a variable lot size with a periodic review. However, every R periods an order is placed up to the reorder level S. The safety stock should cover the uncertainty in demand during lead time plus the review period (Van der Heijden, 2022). The review period is often dependent on external factors like the frequency of truck deliveries (Silver et al., 2016).

3.3.2 Performance Measures

There are multiple ways to measure the performance or to set objectives for inventory control. These can be divided into shortage costs and service levels. These can be calculated or these could be set as goal. Based on the cost, there are four measures (Van der Heijden, 2022).

1. B_1: Costs per stockout occasion
2. B_2: Fractional charge per unit short
3. B_3: Fractional charge per unit short per unit time
4. B_4: Costs per customer order line short

Based on the service level, there are five measures.

1. P_1: Cycle service level
2. P_2: Volume fill rate
3. P_3: Ready rate
4. Time between stockout occasion
5. Order (line) fill rate

These measures are used to calculate the actual performance, but also to calculate the expected performance. Organisations can set targets based on a specific measure. Subsection 3.3.3 gives examples to calculate inventory parameters based on a target volume fill rate.

3.3.3 Reorder Point & Safety Stock

One of the main parameters is the reorder point. The reorder point indicates when an order should be placed. There are multiple ways to calculate the reorder point dependent on the demand distribution. Before looking into the calculations, we take a look at determining the demand distribution.

Silver et al. (2016) propose considering the gamma or lognormal distribution when the CV is greater than 0.5, else the normal distribution is fitting. Also, when the expected demand over lead time is below 10 units, other distributions are more appropriate. This reasoning and determination of the demand distribution follow the decision tree in Figure 3.5 (Van der Heijden, 2022).

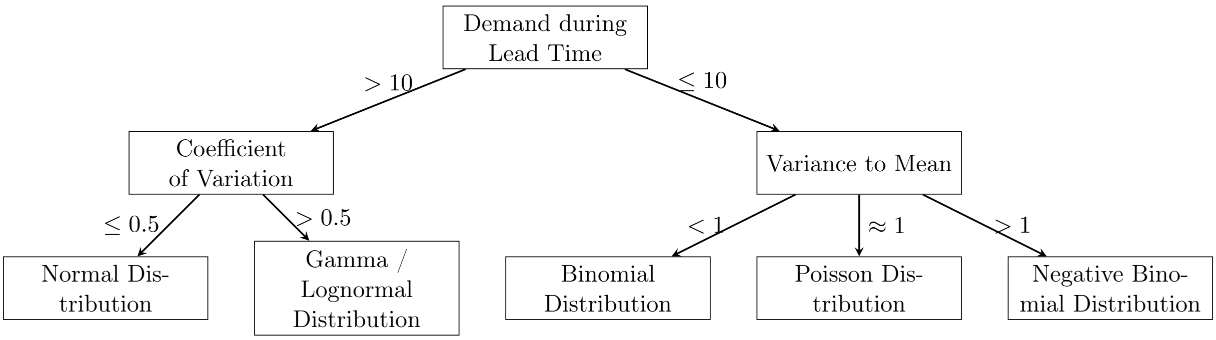

FIGURE 3.5: Demand Distribution Decision Tree

When considering the normal distribution, the reorder point is calculated with the use of the demand over lead time and the safety stock (see Equation 3.3). The standard deviation is dependent on the uncertainty of demand and lead time. The equation assumes that the lead time (L) and annual demand (D) are independent from each other (Silver et al., 2016). The standard deviation (σ) is calculated over the errors (ϵ), rather than over the demand itself because we already considered explainable variation in demand due to seasonality and trend when forecasting. The standard deviation should only account for the unexplainable variation in demand: the uncertainty. This equation also assumes that the forecast errors are not correlated, while these are. Prak et al. (2017) state that this assumption underestimates the variance over lead time. They correct this underestimation by applying the law of total variance. As this in-depth standard deviation calculation goes beyond the scope of this research, we will stick to Equation 3.3.

$$s = \hat{x}_L + k \times \sqrt{E(L) \times \sigma_\epsilon^2 + E(D)^2 \times \sigma_L^2} \tag{3.3}$$

Furthermore, the safety stock is calculated with the safety factor k. This is calculated differently according to the service level measure. As the volume fill rate P_2 is important for NS, Equation 3.4 and the standard normal loss function are appropriate for calculating the safety factor, where the order quantity Q is used. Note that these equations do not include undershoot which is not relevant for the (s, Q)- or (s, S)-policy (Van der Heijden, 2022).

$$G(k) = \frac{Q \times (1 - P_2)}{\sqrt{E(L) \times \sigma_\epsilon^2 + E(D)^2 \times \sigma_L^2}} \tag{3.4}$$

When considering the other distributions, we need different equations for calculating the reorder points. Equation 3.5 is the basis for using the volume fill rate in the calculation of the expected shortage per replenishment cycle ($ESPRC$). Silver et al. (2016) write about calculating the reorder points in case of gamma and poisson distributed demand. With the use of Equation 3.6 and the cumulative distribution function of a gamma distribution, we can find s for gamma distributed demand. Here, the alpha and beta are the shape and scale parameters, respectively. For poisson distributed demand, Equation 3.7 holds to determine the reorder point. This is calculated with the cumulative distribution function of a poisson distribution and the lambda as mean arrival rate. There are no sources found for the calculation of the ESPRC in case of the binomial and negative binomial distribution. However, by transposing the ESPRC formulas in Appendix B, we find the formulas for binomial and negative binomial demand in Equation 3.8 and 3.9, respectively.

$$ESPRC = Q \times (1 - P_2) \tag{3.5}$$

$$ESPRC = \alpha\beta \times [1 - F_{gamma}(s; \alpha + 1, \beta)] - s \times [1 - F_{gamma}(s; \alpha, \beta)] \tag{3.6}$$

$$ESPRC = \lambda \times [1 - F_{pois}(s - 1; \lambda)] - s \times [1 - F_{pois}(s; \lambda)] \tag{3.7}$$

$$ESPRC = np \times [1 - F_{binom}(s - 1; n - 1, p)] - s \times [1 - F_{binom}(s; , n, p)] \tag{3.8}$$

$$ESPRC = \frac{r \times (1 - k)}{k} \times [1 - F_{nbinom}(s - 1; r + 1, k)] - s \times [1 - F_{nbinom}(s; r, k)] \tag{3.9}$$

3.3.4 Order Size

When having a constant demand, it is possible to use the EOQ formula to determine the order size of an inventory policy. However, it is also possible to adapt to time-varying demand when determining the order size. Basically, there are three approaches to deal with a deterministic, time-varying demand pattern. When the variability is low, the EOQ formula still makes sense to use. Silver et al. (2016) propose using the EOQ formula when the coefficient of variation is smaller than 0.2. Also, the Wagner-Whitin method can be used as exact best solution. Lastly, an approximate or heuristic model, like the Silver-Meal heuristic, is an approved method to determine the order size. There are two situations where the use of a heuristic can lead to significant costs. When the demand drops rapidly and when there are a lot of periods having no demand. This argument in combination with simplicity, draws the conclusion to choose for the EOQ method.

3.4 Combining Forecast Methods with Inventory Control

The research of this thesis considers both forecasting and inventory control. A lot of articles describe solving a problem regarding either of these subjects. Therefore, it is insightful to gather literature sources about articles that describe a solution to both subjects. This section is dedicated to this.

Bacchetti and Saccani (2012) state that the elements of integrated spare parts management follow a closed loop, see Figure 3.6. The loop goes from spare parts classification, to demand forecasting, then inventory management and ending with the performance assessment after which it returns to spare parts classification . Mobarakeh et al. (2017) follow a similar approach. The approach starts with classifying the SKUs, then it compares a bootstrap forecast function with other existing forecast methods like moving average, exponential smoothing, and Croston. The best scenario is chosen for each demand class. In the end, it calculates the inventory costs based on a predefined inventory strategy based on intermittent demand. This results in a matrix of the best forecasting method and best inventory management strategy for each SKU.

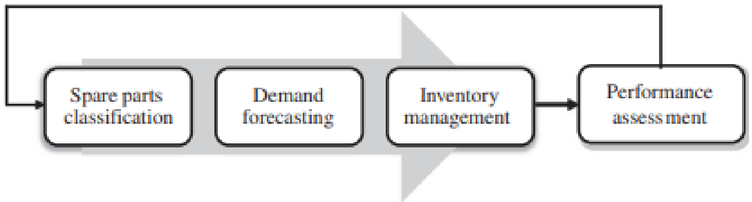

FIGURE 3.6: Integrated Approach to Spare Parts Management (Bacchetti and Saccani, 2012)

Table 3.1 shows what literature combines which forecast methods with which inventory control policies. Here, MA means moving average. Relevant outcomes of this literature are:

- The Croston method is used for infrequent (intermittent) demand data in combination with the (s, (n)Q)-policy (van Wingerden et al., 2014). This is relevant for the problem of NS, as NS also has a lot of infrequent demand data.
- In a similar industry, the aircraft service, the SES method is used in combination of the (R, S)-policy (Syntetos and Boylan, 2006).
- Normal distributed demand including seasonality is forecasted using the Holt Winters' method in combination with the (s, Q)-policy (Alstrøm and Madsen, 1994).
- In Table 3.1, we find the same forecast methods as we find in Section 3.2. So, it is worth evaluation all of these in Chapter 4.
- In Table 3.1, we find mostly articles using the (s, Q)- and (s, (n)Q)-policy. This is a simple approach similar to the policy currently used. We want to use a continuous policy as this is possible within the capacities of NS and reduces the safety stock needed. Also, a fixed lot size is preferable as spare parts could come in boxes or on pallets. To make sure the inventory position increases above the reorder level when reordering, we choose the (s, (n)Q)-policy for Chapter 4.
- Our research expands the literature by identifying spare parts with seasonal variation and evaluating many forecast methods on a niche area, namely spare parts management in a railway company. The integration with an inventory policy is not a theoretical contribution but it is relevant for NS.

TABLE 3.1: Related Research Solution Approaches

Problem Context	**Forecast Method**	**Inventory Control**	**Comment**	**Article**
Infrequent, stochastic demand & LT	Croston & bootstrapping	(s, (n)Q)-policy	-	van Wingerden et al. (2014)
All demand patterns	MA, Croston & bootstrapping	(s, (n)Q)-policy	Policy recommendation for each SKU category	do Rego and de Mesquita (2015)
High variable demand patterns (seasonality)	MA & maximum value	(R, s, S)-policy	Coordinated replenishment policy is considered	Tiacci and Saetta (2009)
Seasonality and cyclic demand patterns	Ensemble deep learning	(R, s, S)-policy	Using machine learning	Seyedan et al. (2023)
Demand history of aircraft service	SES	(R, S)-policy	Addresses error of auto-correlation	Syntetos and Boylan (2006)
Normal distributed demand with seasonality	Holt Winters'	(s, Q)-policy	Includes shortage costs	Alstrøm and Madsen (1994)
Fast-moving consumer goods	Tobit Kalman filter	Newsvendor replenishment policy	Includes lost sales	Trapero et al. (2023)
Seasonal Spare Parts	MA, Croston, Holt Winters' method	(s, (n)Q)-policy	This is our research	Van Sambeek (2024)

3.5 Conclusion of the Literature Research

This chapter answers the second research question of this thesis. The findings of this chapter give an overview of the possibilities for the model approach in the next chapter.

> RQ2: *"What are possible forecast methods and inventory policies for the spare parts to apply in the context of SCO?"*

First, we identified several classification methods present in literature. Combining multiple classification methods addresses limitations of a single classification method. Combining ABC- and XYZ-classification including relocation of critical spare parts is a robust approach and suggest target service levels. Besides classification based on ADI and the CV is a commonly accepted method to determine suitable forecast methods. The FSN- and VED-classification method are not relevant for this research.

Furthermore, multiple regression is a time-efficient method to identify the presence of seasonal variations in comparison to an exploratory data analysis. When seasonal variations are present, forecast methods, like the Holt Winters' trend-seasonal method, SARIMA, or machine learning methods can be used. The Croston method performs well when observing an intermittent demand pattern which could be performing well as seasonal spare parts could have intermittent demand. Other mentioned methods, like the moving average, SES, and Holt method could be relevant in case of smooth demand. There is not one single performance measure for forecasting that captures the complete accuracy. Multiple measures like the bias, mean absolute deviation, root mean squared error, and symmetric mean absolute percent error give an insight in the forecast performance. However, the symmetric mean absolute percentage error gives the best insight as it is not sensitive to outliers, zero demand, and possible to aggregate over multiple spare parts. To prevent overfitting, we should separate training and testing data and include the growing-window forward-validation method when forecasting. The growing-window forward-validation method is a method appropriate for time-series data.

Additionally, the chapter explained different inventory policies and corresponding performance measures. The most important measure for NS is the volume fill rate. Setting a target for this measure is used to calculate a corresponding reorder point. Also, we should consider different demand distributions for calculating the reorder points as these are relevant. The EOQ formula is suggested to determine the order quantity as the Wagner-Whitin method and Silver-Meal heuristic are complex and sensitive to a high ADI, respectively.

Lastly, we looked into literature about combining forecast methods with inventory control. Here, the literature confirms the approach of first classifying spare parts, then forecasting, inventory, and ending with performance assessment. Also, we gave an overview of examples of combining forecast methods with inventory control policies. We conclude that it is worth evaluating all forecast methods from Section 3.2 in Chapter 4. Also, we conclude that we will be controlling inventory using the (s, (n)Q)-policy as this is the most used policy.

The end of this chapter completes step 4 of the MPSM, formulating solutions. Figure 3.7 shows the progress of the research so far.

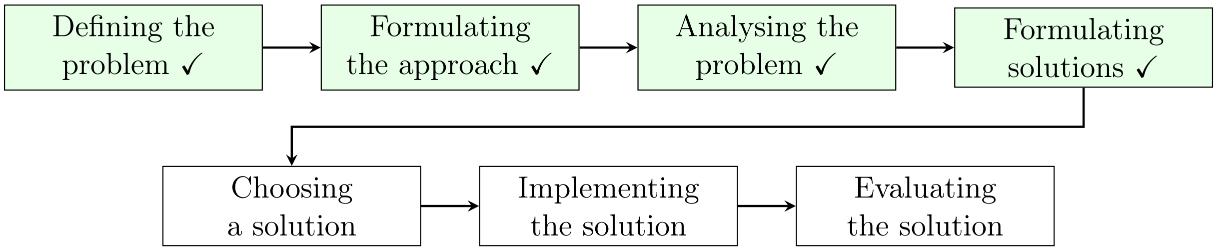

Figure 3.7: MPSM Flow Chart - Step 4 Completed

Chapter 4

Model Design

The fourth chapter presents the model approach to answer the third research question (RQ) of this thesis partly.

> RQ3: *"What are the best fitting forecast methods and inventory policies for the seasonal spare parts at SCO?"*

We visualize the model design in Figure 4.1. This chapter follows this flow chart step by step. First, this chapter describes identifying seasonal spare parts, where the non-seasonal spare parts are excluded from the research. Next, it will explain the classification method for the seasonal spare parts. Then, several forecast methods will be implemented as well as a method to measure the performance of these forecasts. The best performing forecast methods will be chosen as input for the inventory model. The last section describes the inventory control that will be used and how to measure its performance. All parts of the model are coded in the programming software R (R Core Team, 2024). Summarising, this chapter will explain the model after which Chapter 5 will use this model to perform numerical experiments and present the results.

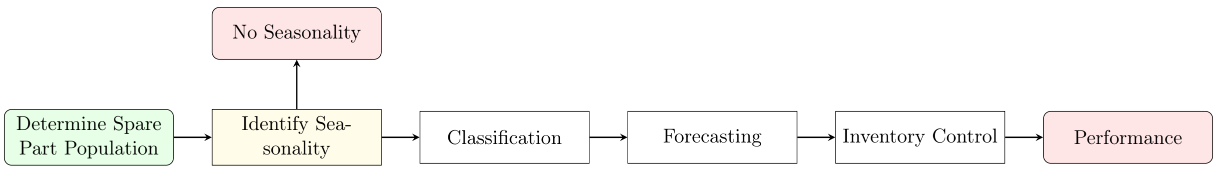

FIGURE 4.1: Flow Chart of Model Approach

4.1 Identifying Seasonality

To identify the presence of seasonality in spare part demand, we use the statistical analysis outlined in Subsection 3.2.7. This analysis investigates the demand pattern of a particular spare part to see if there is a relationship between the demand and the corresponding months. With the analysis goal in mind, we specify the null and alternative hypotheses. These are:

H_0 : There is no relationship between the months of a year and the demand.

H_A : There is a relationship between the months of a year and the demand.

This test is a correlational analysis where we explore the possible relationships between the variables demand and month. The demand is a numerical variable whereas the month is a categorical one. The input for this approach requires demand data of more than a year as we consider yearly returning seasons. Our objective is to perform a regression analysis on the demand data which can conclude that there are yearly recurring seasonal variations in certain months. Before performing the regression analysis, we have to make sure we are dealing with stationary data. This means that the demand data is constant over time. Because seasonal variations could be influenced by the trend of demand, we should remove this influence to make the data stationary.

To remove the trend, we should estimate it. First, we calculate the twelve months moving average using the rollmean function (R Core Team, 2024) to disregard any seasonality. This is not possible for the first eleven months, so we exclude these data points from the identification of seasonality. Then, we perform linear regression on the moving averages of the selected demand points using the stats package (R Core Team, 2024). The regression provides a slope which is considered as the estimated trend b_i. Equation 4.1 shows the formula to calculate the de-trended demand $\tilde{D}_{it}$ per time period t for each spare part i with use of the original demand data D_{it} and subtracting the estimated trend multiplied by the time period.

$$\tilde{D}_{it} = D_{it} - b_i \times t \tag{4.1}$$

The next step is the regression analysis on the de-trended demand to identify seasonality. The month variable is a categorical variable and should be transformed into a numerical variable to include in the regression analysis. A solution for this problem is creating dummy variables. Dummy variables are binary, they only take two possible numerical values, zero or one (James et al., 2022). In the context of the variable month, we create eleven dummy variables. An example of a dummy variable is:

$$Jan_t = \begin{cases} 1 & \text{if the } t\text{-th demand point is in January,} \\ 0 & \text{if the } t\text{-th demand point is not in January.} \end{cases}$$

Table 4.1 shows three examples of demand points. Each demand point now has twelve variables, one demand variable and eleven month dummy variables. Note that there is no dummy variable for the month December. When all dummy variables are set to zero, this represents a demand point in December as shown in the third example. The level with no dummy variable (December) is known as the baseline (James et al., 2022). As there are multiple variables, we use multiple regression with the demand data as y-input (dependent variable) and the dummy variables as x-input (independent variables).

Table 4.1: Example of Three Demand Points with Corresponding Values of Demand and Month Dummy Variables

Demand Point	**Demand**	**Jan**	**Feb**	**Mar**	**Apr**	**May**	**Jun**	**Jul**	**Aug**	**Sep**	**Oct**	**Nov**
January 2023	22	1	0	0	0	0	0	0	0	0	0	0
July 2023	32	0	0	0	0	0	0	1	0	0	0	0
December 2023	42	0	0	0	0	0	0	0	0	0	0	0

To test whether the multiple regression model is significant or does rely on chance, we calculate the significance of the F-test statistic. This calculation is known as the Fisher test and is performed with the use of the stats package (R Core Team, 2024). The significance of the F-test statistic is the chance that any association between the predictors and response occurs by chance, in the absence of any real association between the predictors and the response (James et al., 2022). This is also known as the type I error (see Table 4.2). To identify the spare parts that include a form of seasonality within the demand, there is a threshold. This is a chosen significance level. The threshold is used to determine which spare part demand patterns include seasonality. When the calculated significance is lower than the chosen threshold level, we reject the H_0 hypothesis and identify seasonality. It is unlikely that all parameters of the model are zero. The identified seasonal spare parts are the input for the further model.

TABLE 4.2: Seasonal Spare Part Confusion Matrix

Concluded \Reality	**No Seasonality**	**Seasonality**
No Seasonality	Confidence Level $(1-\alpha)$	Type II error (β)
Seasonality	Type I error (α)	Statistical Power (1-β)

In this analysis, we make several assumptions because of simplicity reasons, these are:

- We assume seasonality to be monthly dependent as we have monthly data while it could also be possible that seasonality is quarterly dependent,
- We assume a consistent cycle length of twelve months, the length of a year,
- Using the multiple regression analysis, we assume additive seasonality,
- We assume that the seasonal variations are stationary.

4.2 Classification

To make deliberate choices in forecasting and inventory control, we perform a three-dimensional classification on seasonal spare parts. This classification consists of the (i) ABC-classification, (ii) XYZ-classification, and the (iii) classification based on the average demand interval (ADI) and coefficient of variation (CV) which are explained in Section 3.1. The ABC- and XYZ-classification will give guidance in decisions regarding inventory control, while the classification method based on the average demand interval and coefficient of variation helps choosing fitting forecast methods as specific forecast methods perform well on specific demand patterns. This section explains the method for the classification, Section 5.2 describes which cut-off values are used and what the results are.

The ABC-classification (i) method requires unit prices and historical demand from the seasonal spare parts. We multiply the unit price and summed demand to calculate the usage value. After sorting the spare parts according to the usage value, the spare parts are split into three classes. The classification will provide a list of spare parts with an A-, B-, or C- classification. An addition to the ABC-classification is that we also include the criticality of spare parts. When a spare part has the characteristic of being critical, we decide to move it to the A class.

The XYZ-classification (ii) only requires historical demand data from the selection of spare parts. Based on this data, the coefficient of variation (CV) is calculated with Equation 4.2 and spare parts are classified in the X-, Y-, or Z-class. Low CV values represent stable demand and are classified with an X, where high CV values indicate irregular demand and are labeled as Z.

$$CV_i = \frac{\sigma_i}{\mu_i} \tag{4.2}$$

where:

- σ_i is the standard deviation of the demand data for spare part i,
- μ_i is the mean of the demand data for spare part i.

Table 4.3 shows qualitative target service levels (TSLs) based on the ABC- and XYZ-classification. The high usage value and predictable spare parts (class AX) receive a higher TSL as this will benefit the inventory costs. The low usage value and less predictable spare parts (class CZ) receive the lowest TSL. The TSLs are chosen based on an aggregate TSL. Compared to the current classification of SCO, explained in Subsection 2.2.4, we are combining instead of separating the usage and part value. Therefore, the rule that spare parts with a low value receive a high TSL does not hold anymore. Now, spare parts with a low usage value receive a low TSL.

TABLE 4.3: Qualitative TSLs based on ABC- and XYZ-Classification

ABC\XYZ	**Class X**	**Class Y**	**Class Z**
Class A	Highest	High	Medium
Class B	High	Medium	Low
Class C	Medium	Low	Lowest

The third classification method (iii) also uses the CV. But next to the CV, the calculation of the ADI, requiring historical demand data, is needed. Equation 4.3 shows the corresponding formula. The ADI provides insight into the intermittency of the demand pattern. Higher ADI values indicate more intermittent demand. Both calculations will result in a division of four classes, namely the intermittent, smooth, erratic, or lumpy classes.

$$ADI_i = \frac{T_i}{D_i} \tag{4.3}$$

where:

- T_i is the total number of periods observed for spare part i,
- D_i is the number of periods with non-zero demand for spare part i.

Table 4.4 provides the classification based on the CV and the ADI.

TABLE 4.4: Classification based on CV and ADI

CV \ADI	**Low ADI**	**High ADI**
Low CV	Smooth	Intermittent
High CV	Erratic	Lumpy

The three-dimensional classification approach using the ABC-classification, XYZ-classification, and the classification based on the ADI and CV leads to 36 classes in total as the methods represent three, three, and four classes, respectively ($3 \times 3 \times 4 = 36$).

4.3 Forecasting

Forecasting involves choosing a method, initialising the parameters of that method, the demand forecast, and updating the parameters. To choose an appropriate forecast method, we develop a model that uses historical corrective demand data to test various forecasting methods. We should consider the best performing forecasting methods as appropriate forecast methods.

In our model, we implement the current used forecast methods of SCO, but also additional forecast methods that include seasonality and are suitable for different demand patterns. These forecast methods are explained in Section 3.2 and implemented with the use of the forecast package (R Core Team, 2024). This package also makes estimates of the initial values and smoothing factors of a forecast method. This means we neglect the proposed smoothing parameters within the literature research in Section 3.2 as these are possibly not optimal for our set of data. The forecasts are one period ahead which means that the forecasts are updated every month. The currently used forecast methods are:

- Moving average,
- Weighted average,
- Simple exponential smoothing (SES),
- Holt,
- No forecast (assuming zero demand every time period).

The additional proposed forecast methods are:

- Holt-Winters,
- Croston,
- (Seasonal) Auto Regressive Integrated Moving Average ((S)ARIMA).

We say (S)ARIMA as we use the auto.arima function (R Core Team, 2024) which automatically selects the best model based on the input data. This means it can in- and exclude the seasonal parameters depending on what fits the data the best. In our model, we are not using any machine learning methods. This is because of the low amount of data points. Machine learning methods are known to perform well when they have access to a lot of historical data.

We implement the growing-window forward-validation method, explained in Subsection 3.2.8. This validation method makes the results more robust and prevents overfitting. We calculate the bias, mean absolute deviation, and the symmetric mean absolute percentage error per validation set with the formulas in Subsection 3.2.8. By averaging the performance of the different validation iterations, we can compare the forecast methods and provide a substantiated conclusion for the best forecast method. One remark is that we test the forecast methods on the same data as we tune the parameters on. This is not the best approach. Ideally, we would have enough data to separate tuning and testing. However, Section 5.3 explains that we do not have this amount of data which made us decide to tune and test on the same data.

The assumptions we make in this model are that we assume that we will capture seasonal patterns within at least one year of historical data.

4.4 Inventory Control

We will use the realized demand data of the most recent year to test the continuous (S, (n)Q)-policy inventory policy as this covers a year of demand and includes all seasonal variations. This means that we have another test set. It is important that this is not used in the forecast selection procedure of Section 4.3 as we want to test our inventory policy with forecasted demand on unseen data. Due to time limitations and therefore simplicity reasons:

- We generalize exchange parts as wear parts,
- We assume no variability for preventive maintenance demand,
- We assume infinite supplier capacity,
- We assume constant unit prices, holding rates and ordering costs, so also no bulk discounts,

- We assume no spare parts in the pipeline at the start of testing the inventory model.
- We determine the mean and variance of the lead time based on historical records. If not present, we assume a mean of 222 business days and the maximum known variance. We do not consider the lead times within the contracts with suppliers,
- We use static inventory parameters,
- Some spare parts have expire date which means they expire after being hold on stock for too long resulting in more inventory costs. We exclude this from the model.

4.4.1 Forecasting Demand over Lead Time

With the use of the best performing forecast method and the lead times (L), we forecast the corrective demand over lead time for every time period of the test set. Furthermore, we add the preventive demand over lead time for every time period of the test set to the forecast to end up with the total forecasted demand over lead time ($\hat{x}_L$).

4.4.2 Demand Distribution Division

An important parameter is the reorder point. How to calculate this, is determined by the demand distribution. Subsection 3.3.3 explains with the use of a decision tree how to determine the demand distribution. We look at the forecasted demand over lead time, CV, and the variance over mean (V/M) ratio. In Subsection 4.4.1, we explained how to obtain the forecasted demand over lead time. In Section 4.2, we described the formula for the CV. However, instead of the mean and standard deviation of demand, we use the forecasted demand over lead time and the standard deviation of the forecast error over lead time respectively, as seen in Equation 4.4. The same holds for the calculation of the V/M ratio, shown in Equation 4.5 (Van der Heijden, 2022).

$$CV_i = \frac{\sqrt{E(L) \times \sigma_\epsilon^2 + E(D)^2 \times \sigma_L^2}}{\hat{x}_{Li}} \tag{4.4}$$

$$V/M_i = \frac{E(L) \times \sigma_\epsilon^2 + E(D)^2 \times \sigma_L^2}{\hat{x}_{Li}} \tag{4.5}$$

where:

- $\sigma_{\epsilon i}$ is the standard deviation of the forecast error for spare part i.

After calculating the CV and V/M ratio, we divide the spare parts into distribution groups. How we do this is explained in Subsection 3.3.3. The assigned distribution forms the basis for the reorder point calculations.

4.4.3 Normally Distributed Demand

To calculate the reorder point (s) of the normally distributed spare parts, we first calculate the safety stock (SS) which is explained in Subsection 3.3.3. Since we are considering uncertain demand and lead time, the equation is this extensive. With the calculated safety stock, we can use the formula in Equation 3.3 to determine the reorder point. We are updating the reorder point for every time period in the test set as we have a forecasted demand over lead time for all these time periods.

4.4.4 Other Distributed Demand

To calculate the reorder point (s) of the spare parts with a non-normal distribution of demand, we use Equation 3.5. Here, we calculate the maximum estimated shortage per replenishment cycle ($ESPRC$) given the TSL. With use of Equations 3.6, 3.7, 3.8, and 3.9 we find s where we use $\hat{x}_L$ for the mean and $E(L) \times \sigma_\epsilon^2 + E(D) \times \sigma_L^2$ for the variance. Again, the calculations for the gamma and poisson distribution come from Silver et al. (2016). For the gamma distribution, we will estimate α and β with the use of Equation 4.6 and 4.7. We will estimate λ with Equation 4.8 for the poisson distribution. For the binomial distribution, we will estimate n and p with Equation 4.9 and 4.10 respectively where we round n to an integer. Lastly, for the negative binomial distribution, we will estimate r and k with Equation 4.11 and 4.12 respectively (Mandal, 2023). There is one remark, we can not find s in these equations algebraically. We have to estimate s with the use of an iterative process in R where we will increase s until we reach the desired volume fill rate (Silver et al., 2016).

$$\alpha = \frac{\hat{x}_L^2}{E(L) \times \sigma_\epsilon^2 + E(D)^2 \times \sigma_L^2} \tag{4.6}$$

$$\beta = \frac{E(L) \times \sigma_\epsilon^2 + E(D)^2 \times \sigma_L^2}{\hat{x}_L} \tag{4.7}$$

$$\lambda = \hat{x}_L \tag{4.8}$$

$$n = round[\frac{\hat{x}_L^2}{\hat{x}_L - (E(L) \times \sigma_\epsilon^2 + E(D)^2 \times \sigma_L^2)}] \tag{4.9}$$

$$p = \frac{\hat{x}_L}{n} \tag{4.10}$$

$$r = \frac{\hat{x}_L^2}{E(L) \times \sigma_\epsilon^2 + E(D)^2 \times \sigma_L^2 - \hat{x}_L} \tag{4.11}$$

$$k = \frac{\hat{x}_L}{E(L) \times \sigma_\epsilon^2 + E(D)^2 \times \sigma_L^2} \tag{4.12}$$

4.4.5 Order Quantity

Furthermore, we can calculate the order quantity with use of the economic order quantity (EOQ), see Equation 4.13.

$$Q_i = \sqrt{\frac{2D_iS_i}{h_ic_i}} \tag{4.13}$$

where:

- D_i is the annual demand for spare part i,

- S_i is the ordering cost per order for spare part i,
- h_i is the holding rate per unit per year for spare part i,
- c_i is the unit price for spare part i.

To prevent that the EOQ suggests a quantity that does not increase the inventory position above the reorder point, we use the option to order $n \times EOQ$ where n should be chosen such that the inventory position surpasses the reorder point. However, we also take into account the minimum order quantity (MOQ). This is a restriction from the supplier. Therefore, we do $Orderquantity = max(nQ, MOQ)$.

4.4.6 Material Resource Planning

With all input parameters, we set up a material resource planning (MRP) based on a specific inventory policy. An MRP gives an overview of the materials in the supply chain. It calculates the on hand balance, fulfilled demand, backorders, replenishments, pipeline, and inventory position for the test period. Based on the MRP and realized demand, we calculate the performance measures of the inventory control. These are the inventory holding and order costs, but also the realized service levels. The performance measures give an indication of how the inventory costs and service levels relate.

4.5 Conclusion of the Model Approach

This chapter answers the third research question of this thesis partly. The findings of this chapter give the approach for the solution.

> RQ3: *"What are the best fitting forecast methods and inventory policies for the seasonal spare parts at SCO?"*

First, we use multiple regression to determine whether the demand pattern of a spare part is correlated to the months of the year. This helps us identify the seasonal spare parts. Then, we will classify the seasonal spare parts. We use a three-dimensional classification with ABC- and XYZ-classification including relocation of critical spare parts and the classification based on the ADI and CV. The ABC- and XYZ-classification will give guidance in choosing the target service levels, while the third classification method helps choosing fitting forecast methods as specific forecast methods perform well on specific demand patterns.

Furthermore, we presented a forecast model that tests different tuned forecast methods with the use of the growing-window forward validation. The model calculates the forecast accuracy using multiple measures like the bias, mean absolute deviation, and symmetric mean absolute percent error.

Lastly, the inventory control model calculates the reorder points per demand distribution and the order sizes which are used in an MRP. We test our inventory model and obtain the inventory performance. This performance consists of the inventory holding costs, order costs, and the achieved service level.

The model approach of this chapter will be used to experiment with in Chapter 5.

Chapter 5

Results Analysis

The fifth chapter presents the results of the model to answer the third research question (RQ) of this thesis partly.

> RQ3: "*What are the best fitting forecast methods and inventory policies for the seasonal spare parts at SCO?*"

First, this chapter identifies the seasonal spare parts, whereas the non-seasonal spare parts are excluded from the research. Next, it will classify the seasonal spare parts to use as input for the forecast and inventory model. Then, the results of the forecasts and inventory model are analysed. The last section elaborates on the individual value of every model component. Summarising, this chapter will perform numerical experiments and present the results.

5.1 Determine Seasonal Spare Parts

In Section 4.1, we explain the statistical analysis to identify yearly seasonal variations in the demand pattern of spare parts. In this section, we describe the experiments and results of this analysis.

5.1.1 Identified Seasonal Spare Parts

To start, we use the demand data from March 2019 until March 2024 resulting from corrective maintenance of all spare parts within the current train series. This is demand data from 19,483 spare parts. We prepare the data by removing all spare parts with a total demand of less than 60 parts over the past five years. If we do not make this adjustment, there is a possibility that spare parts with low usage will be identified as seasonal spare parts. SCO does not want to handle these as seasonal spare parts as these spare parts do not require intensive monitoring. We do not use a higher demand boundary as spare parts that were introduced two or three years ago could show yearly recurring seasonal variations. With a higher demand boundary, these are possibly not selected. This limitation excludes 95.0% of the spare parts from the statistical analysis. This is a huge proportion as NS never removes historical data from spare parts. What remains are 971 spare parts that we analyse.

Before looking at the results of the multiple linear regression analysis on all demand data, we should make a decision for the value of the threshold for the significance of the F-test statistic. Figure 5.1 shows the relation between the number of identified seasonal spare parts and the significance threshold. Approximately, it follows a linear graph which indicates that increasing or decreasing the threshold would mean a linear increase or decrease of identified seasonal spare parts.

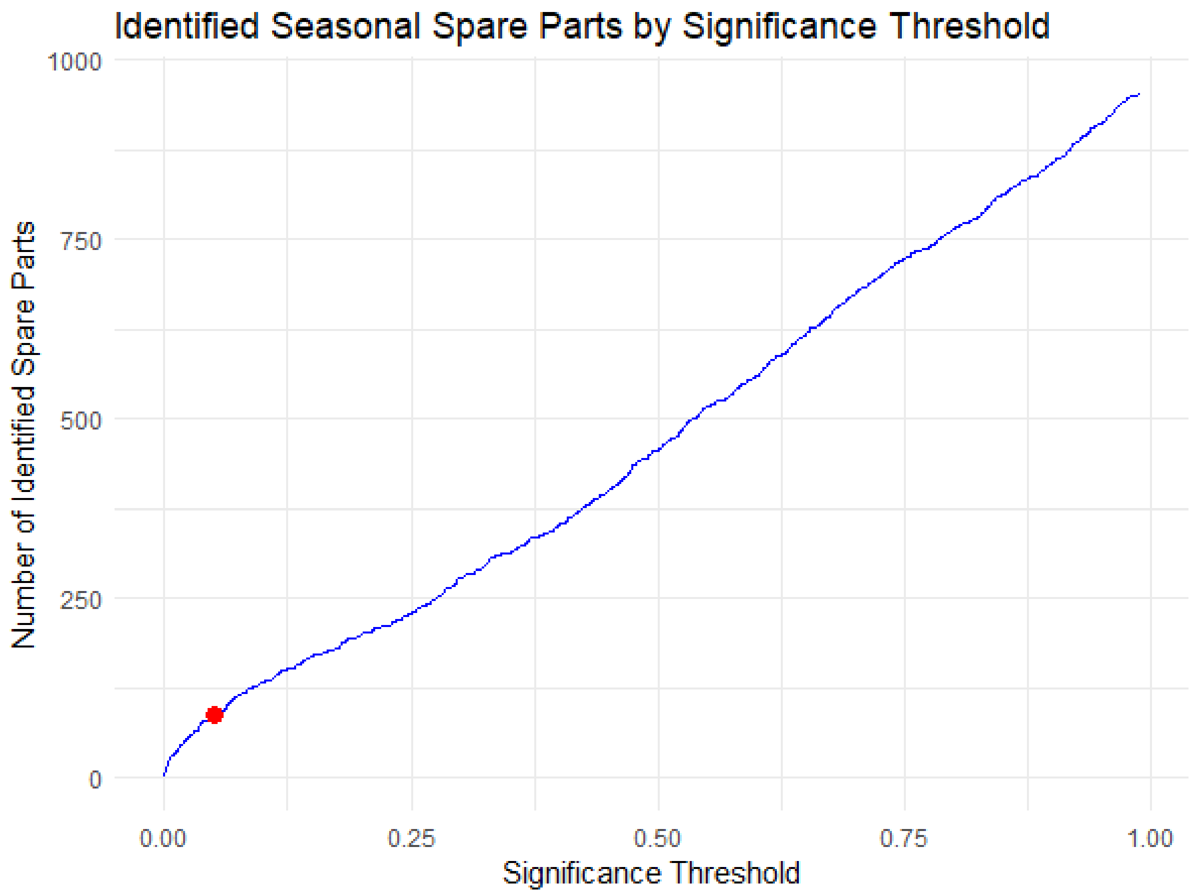

FIGURE 5.1: Line Graph with Number of Identified Seasonal Spare Parts for Different Significance Thresholds

We set the threshold for the significance of the F-test statistic to 0.05 as this is a commonly used threshold to accept statistical models and NS agrees on this. Also, we see that for a significance threshold lower than 0.05, the number of spare parts increases relatively more than for a significance threshold higher than 0.05. With a significance threshold higher than 0.05, we should sacrifice more significance to include extra spare parts. Choosing a threshold of 0.05 means that if we identify a spare part as seasonal, it has a probability of at maximum 5% that any association between the predictors and response occurs by chance. Using this threshold, the analysis identifies 88 seasonal spare parts.

5.1.2 Validation Statistical Identification Method

To validate the outcomes of the seasonality identification procedure, we compare the list of previously classified seasonal spare parts with the current identified seasonal spare parts of Subsection 5.1.1. We find nine spare parts that are present in both the previous and the current list. This also means that 50 spare parts are not identified by our procedure. Furthermore, we perform an explanatory data analysis using historical demand graphs to assess the presence of seasonality. Subsection 3.2.7 explains that this procedure is time-consuming. However, the list of previously classified seasonal spare parts is not extensive which makes it possible to perform this explanatory analysis to validate the identified seasonal spare parts.

Table 5.1 shows the results after analysing the demand graphs. By the explanatory data analysis using historical demand graphs (validation), we identify six seasonal spare parts (see also Subsection 2.2.2). This is a lower amount than when using the statistical test. We already discussed some of the demand graphs in Subsection 2.2.2 and Appendix A. In Figure 2.2, NS identifies a clear seasonal demand pattern, in Figure A.1, NS identifies an unclear seasonal demand pattern but they do classify it as seasonal, and in Figure A.2, NS does not identify a seasonal demand pattern. Interesting to know is that the six seasonal spare parts of the validation are also identified as seasonal in the statistical test. We can conclude that this observation validates the statistical test as all visually identified seasonal spare parts are also selected by the statistical test.

TABLE 5.1: Validation Results of Identifying Seasonal Spare Parts based on Previous Classified Seasonal Spare Parts

	NS Classification	**Statistical Classification**	**Validation**
Seasonality	59	9	6
No Seasonality	N/A	50	53

5.2 Classification Analysis

After identifying the seasonal spare parts, we classify them according to the approach described in Section 4.2. We use the demand data from the previous twelve months resulting from the corrective maintenance of all 88 seasonal spare parts. We choose for twelve months as it captures a full year capturing all possible seasons. If we use more or less months of demand data, possibly seasonal variations are captured multiple times or are not captured at all. We could choose to include demand data from the past two years. However, this detracts from the classification as we want to identify the current importance and variability of the spare parts and not the importance and variability of two years ago. Furthermore, we use the unit price of all seasonal spare parts. We calculate the annual usage value by multiplying the demand and unit price and also the coefficient of variation (CV) by dividing the standard deviation by the mean. We set up the following classification boundaries:

- Class A: the spare parts with the highest annual usage value covering 50% of the total annual usage value.
- Class B: the spare parts with the highest annual usage value after class A covering 30% of the total annual usage value.
- Class C: the spare parts with the lowest annual usage value covering 20% of the total annual usage value.
- Class X: the spare parts with a CV lower than 0.20.
- Class Y: the spare parts with a CV higher than 0.20 but lower than 0.40.
- Class Z: the spare parts with a CV higher than 0.40.

The boundaries of the ABC-classification are adopted from the literature in Subsection 3.1.1. However, the boundaries of the XYZ-classification are not adopted but adjusted based on the literature in Subsection 3.1.2. These boundaries are adjusted to create a more even distribution of the spare parts over the classes. Table 5.2 summarises the ABC- and XYZ-classification.

TABLE 5.2: Results of ABC- and XYZ-Classification for Seasonal Spare Parts

ABC\XYZ	**Class X**	**Class Y**	**Class Z**
Class A	4	5	2
Class B	1	2	11
Class C	7	27	29

In Subsection 2.2.3, we describe the difference between critical and non-critical spare parts. As we do not desire to control critical spare parts using a low target service level (TSL), we manually transfer those spare parts to class A so these receive a higher TSL. As 81 of the 88 seasonal spare parts are critical, the adjustment based on criticality makes a lot of impact on the final classification. Table 5.3 shows how the critical spare parts influenced the final ABC- and XYZ- classification. Almost no spare parts appear in class B or C. This raises the question about the value of the ABC-classification. However, it could be the case that for another selection of spare parts, there are less critical spare parts which would make the ABC-classification useful. Besides, including this classification is not time-costly. This makes it worth including the ABC-classification in the analysis.

TABLE 5.3: Adjusted Results of ABC- and XYZ-Classification for Seasonal Spare Parts based on Criticality

ABC\XYZ	**Class X**	**Class Y**	**Class Z**
Class A	12	32	37
Class B	0	0	1
Class C	0	2	4

With the final ABC- and XYZ-classification, we can determine the TSLs for each class. We use an aggregated TSL of 97.0% based on the annual demand as this is the norm for NS. Table 5.4 shows the annual demand for each class while Table 5.5 shows how the division of TSLs over the different classes. These TSLs will be the input for Section 5.4.

TABLE 5.4: Annual Demand of ABC- and XYZ-Classification

ABC\XYZ	**Class X**	**Class Y**	**Class Z**
Class A	13,144	6,046	2,411
Class B	0	0	715
Class C	0	115	105

TABLE 5.5: TSLs of ABC- and XYZ-Classification for Aggregate TSL of 97.0%

ABC\XYZ	**Class X**	**Class Y**	**Class Z**
Class A	99.0%	96.0%	92.0%
Class B	96.0%	92.0%	90.0%
Class C	92.0%	90.0%	85.0%

Using the recommended threshold from literature in Subsection 3.1.3, Table 5.6 provides the classification based on the CV and the ADI. This leads to the outcomes shown in the same table. Most spare parts follow a smooth demand pattern. However, half of the seasonal spare parts follow an intermittent, erratic, or lumpy demand pattern. Possibly, specific forecast methods will perform well on specific demand patterns. Therefore, we will use this classification only for the forecasting part. We will test the forecast methods for each of these classes to see which method fits the best.

TABLE 5.6: Classification based on CV and ADI

CV \ADI	$ADI < 1.32$	$ADI \geq 1.32$
$CV < 0.49$	Smooth: 45	Intermittent: 20
$CV \geq 0.49$	Erratic: 16	Lumpy: 7

5.3 Forecast Analysis

After classifying the seasonal spare parts, we forecast the demand according to the approach described in Section 4.3. We use all available historical demand data resulting from corrective maintenance of the identified seasonal spare parts. However, we exclude the most recent twelve months of demand data as this data is reserved for testing the inventory model in Section 5.4. This means that 48 months (four years) of demand data remain for the forecasting model. As explained, we split the data in training and testing sets using the growing-window forward validation. This validation method prevents overfitting and is appropriate for time-series data. We will use the available data as shown in Figure 5.2. Because we only have 48 months of data, there are some limitations in using the data. It is not possible to use five independent testing sets, so these overlap.

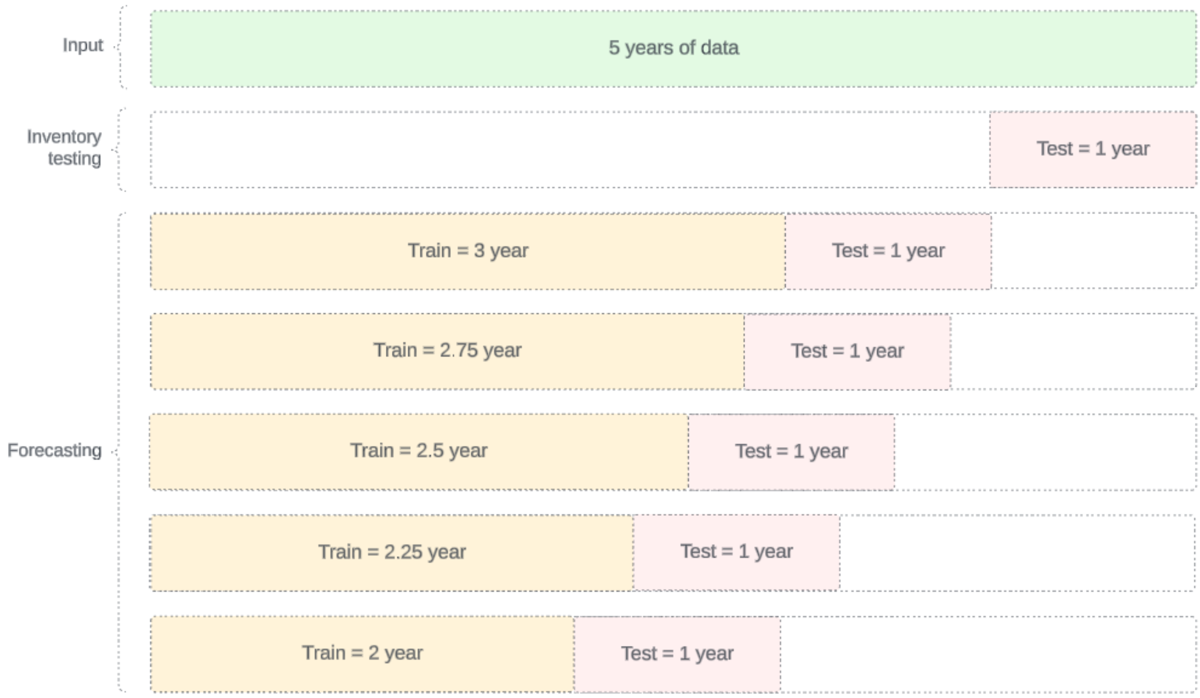

FIGURE 5.2: The Model Validation Splits for Forecasting and Inventory Control

Subsection 3.1.3 explains how each demand class has its suitable forecast methods. To test these hypotheses, we experiment with our forecast methods per demand class. Besides, making a conclusion about the best suitable forecast method for all spare parts will be very generalising. Therefore, it is also better to split up the spare parts in groups with the same demand patterns. This section will elaborate on the performance of the forecast methods based on the smooth, intermittent, erratic, and lumpy demand.

5.3.1 Forecasting Smooth Demand

There are 45 smooth demand seasonal spare parts with an average monthly demand of 34. Forecasting with the eight proposed forecasting methods for smooth demand leads to the results in Table 5.7. These results are the average results of the seasonal spare parts for the five test sets. When looking at the symmetric mean absolute percentage error (sMAPE), the Croston method performs the best, shortly followed by the SES method. This is in line with the hypothesis since smooth demand does not have many variation and a lot of forecast methods could fit well for this type of demand. While the Croston method is considered to be the most effective for intermittent demand, it can still perform well on other types of demand. This method combines the ADI and the demand level and generates an average forecast. When dealing with smooth demand, which has a low ADI, the Croston method can predict this correctly and then it is also likely to perform well.

TABLE 5.7: Average Forecast Performance of Smooth Demand Seasonal Spare Parts

Forecast Method	**sMAPE**
Moving Average	51.60%
Weighted Average	52.08%
SES	51.26%
Holt	52.57%
No Forecast	191.78%
Holt-Winters	57.61%
Croston	51.25%
(S)ARIMA	55.34%

Figures 5.3 and 5.4 show that we use a lot of methods that are forecasting an average demand. However, the Holt-Winters method is trying to predict the peaks. Based on the performance, the forecast methods without much fluctuations score the best for smooth demand.

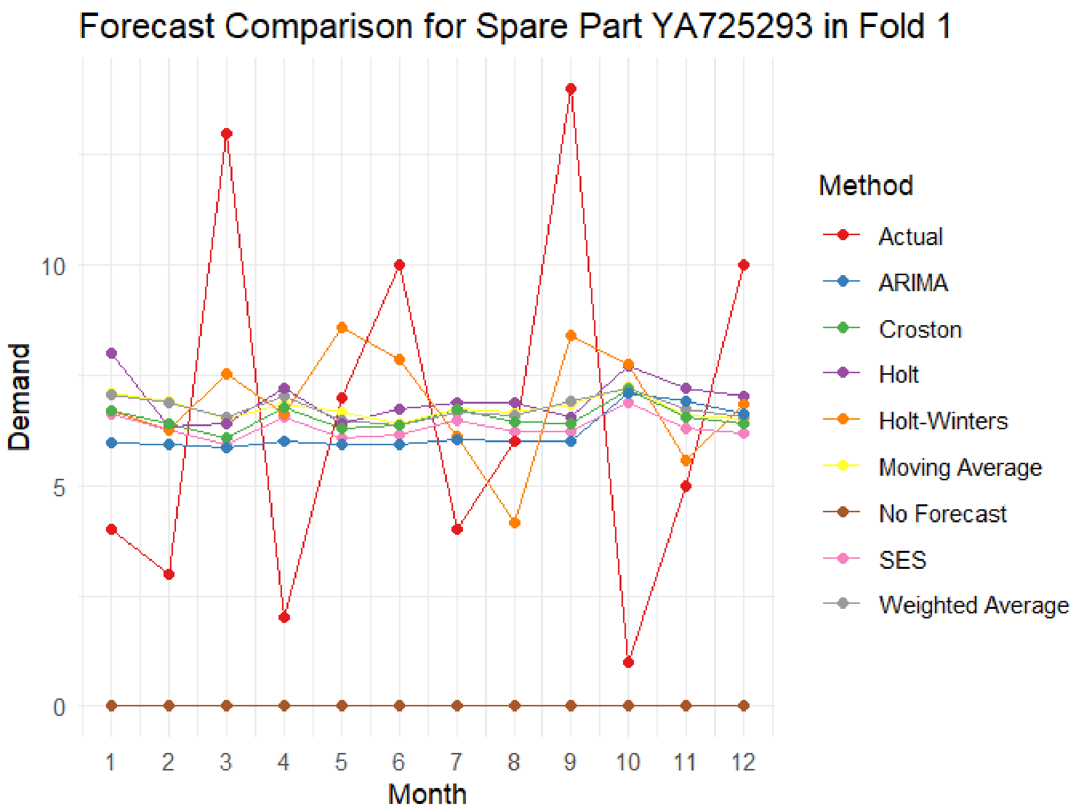

FIGURE 5.3: Forecast Comparison for a Smooth Spare Part

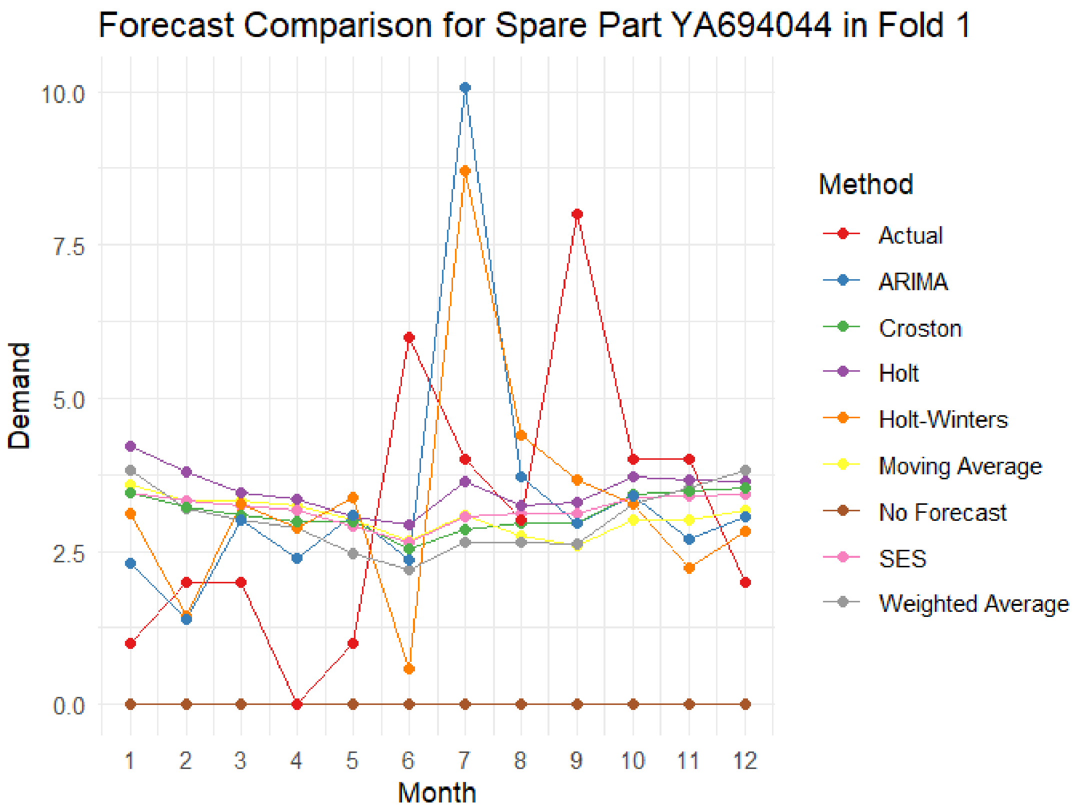

Figure 5.4: Forecast Comparison for another Smooth Spare Part

5.3.2 Forecasting Intermittent Demand

There are 20 intermittent demand seasonal spare parts with an average monthly demand of 7. Forecasting with the eight proposed forecasting methods for intermittent demand leads to the results in Table 5.8. These results are the average results of the seasonal spare parts for the five test sets. All forecast methods perform bad as they have a sMAPE over 100%. This is because the average monthly demand is quite low. When a forecast differs even a very small bit from low demand, the sMAPE can become large easily. The (S)ARIMA method performs the best. This is possibly because (S)ARIMA involves trend and seasonality. It is performing way better than other forecast methods. This is not in line with the hypothesis as the exponential smoothing methods and the Croston method should perform the best. When looking at the demand graphs of intermittent demand, we do not see strange demand patterns. However, we find that a lot of the intermittent spare parts have an ADI between the threshold 1.32 and the value 2.00. 17 of the 20 spare parts are in this area which could be considered as grey area. These spare parts are classified as intermittent parts but show behavior that is similar to smooth demand. This could be a reason why the Croston method is not performing well. However, in Subsection 5.3.1 we see that the Croston method is performing well on smooth demand. Another reason for the bad Croston performance could be the influence of trend and seasonality which we detect in some of the demand graphs. The Croston method does not consider these components.

TABLE 5.8: Average Forecast Performance of Intermittent Demand Seasonal Spare Parts

Forecast Method	sMAPE
Moving Average	112.61%
Weighted Average	112.20%
SES	111.97%
Holt	116.47%
No Forecast	124.67%
Holt-Winters	110.88%
Croston	114.46%
(S)ARIMA	100.04%

Figures 5.5 and 5.6 show that the forecast methods cannot predict the fluctuations of intermittent demand very well. Hence, the bad forecast performance. Here, it becomes visible that the Croston forecast method is often forecasting lower than other forecast methods. The bad forecast performance indicates that in the test set, there are less zero and low demand periods than in the training set. The performance of (S)ARIMA shows that this forecast method has been the best in predicting the peaks and low demands.

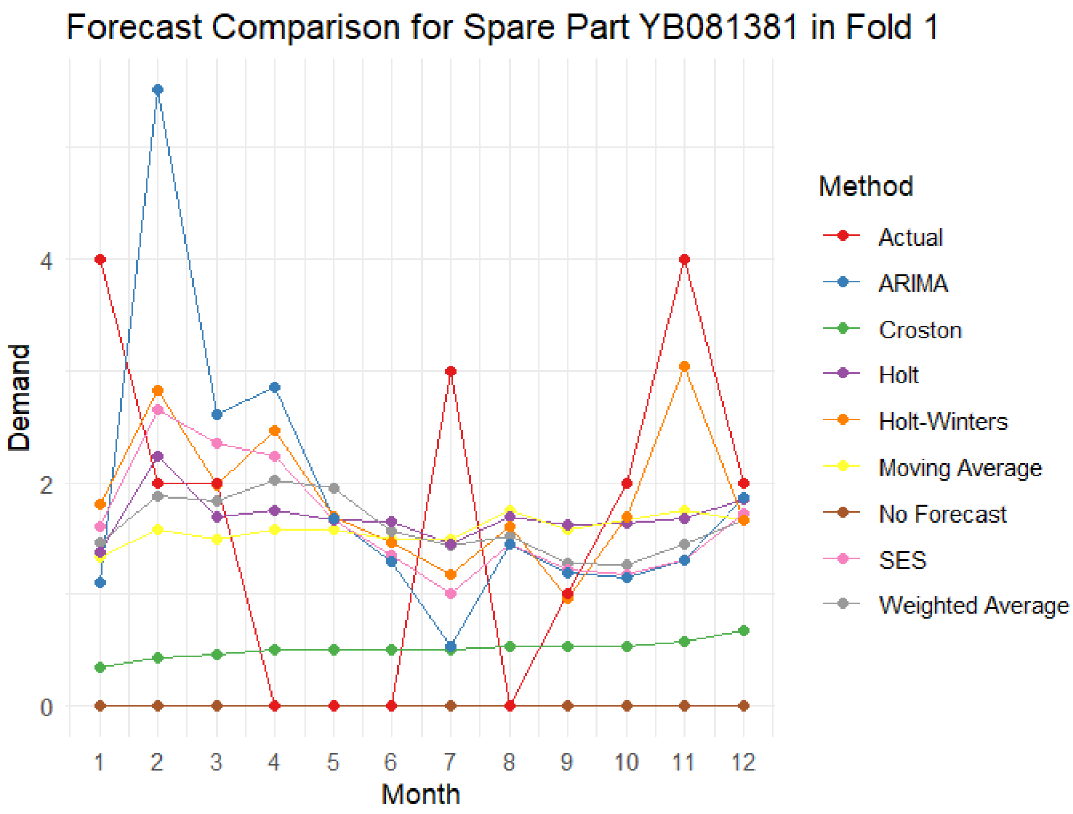

FIGURE 5.5: Forecast Comparison for an Intermittent Spare Part

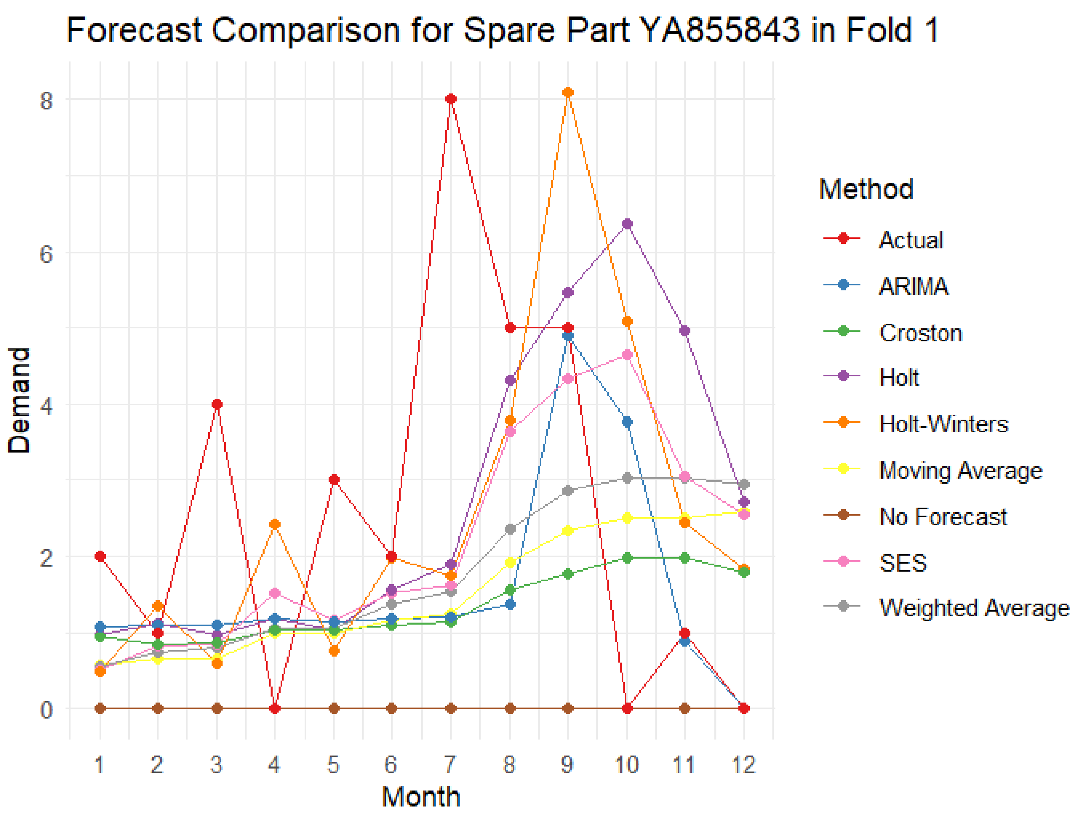

FIGURE 5.6: Forecast Comparison for another Intermittent Spare Part

5.3.3 Forecasting Erratic Demand

There are 16 erratic demand seasonal spare parts with an average monthly demand of 7. Forecasting with the eight proposed forecasting methods for erratic demand leads to the results in Table 5.9 These results are the average results of the seasonal spare parts for the five test sets. When looking at the best performing method based on the sMAPE, the SES scores the best, shortly followed by the Croston and moving average methods. This type of demand does not have one clear best fitting method because of the irregular demand. This is in line with the hypothesis as linear forecast methods do not specifically perform well as machine learning methods are more suitable.

TABLE 5.9: Average Forecast Performance of Erratic Demand Seasonal Spare Parts

Forecast Method	sMAPE
Moving Average	75.87%
Weighted Average	76.37%
SES	75.23%
Holt	77.98%
No Forecast	174.38%
Holt-Winters	79.63%
Croston	75.45%
(S)ARIMA	80.38%

Figure 5.7 is visualising the demand pattern of a spare part with higher demand than the spare part in Figure 5.8. Figure 5.7 shows that the (S)ARIMA and Holt-Winters can follow peaks and low demands well when a spare part has a high average demand. However, according the forecast performance and Figure 5.8 they are not doing well. There are a lot of spare parts with

a low average demand within the erratic demand class. Here, on average the SES and Croston perform better.

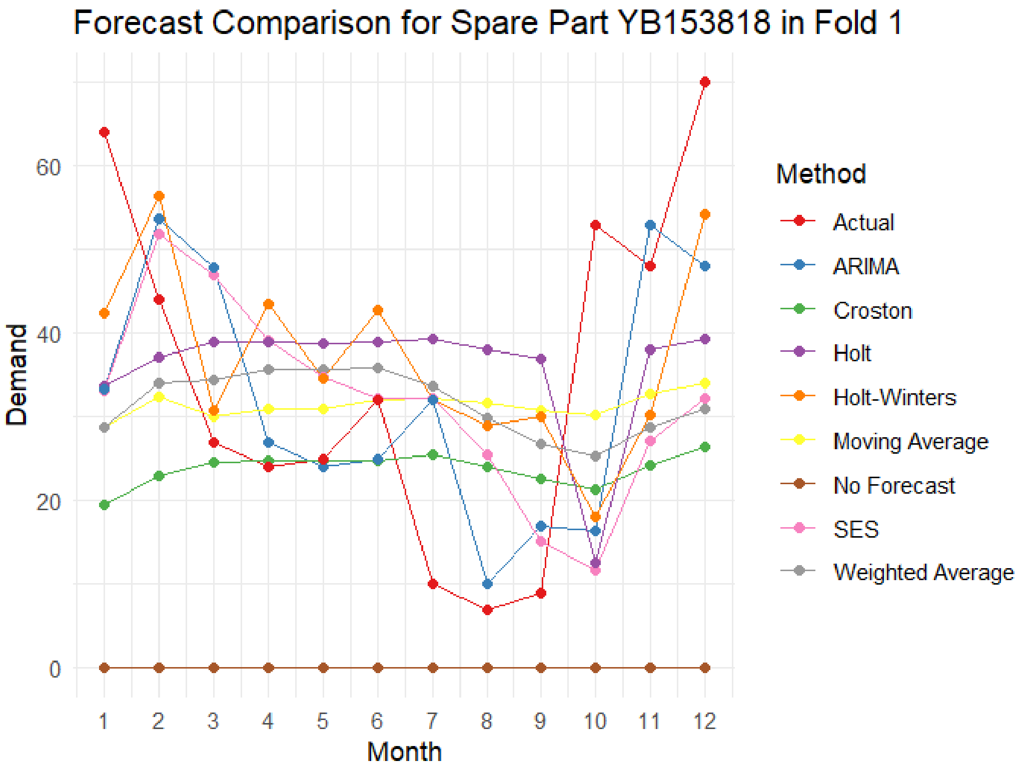

FIGURE 5.7: Forecast Comparison for an Erratic Spare Part

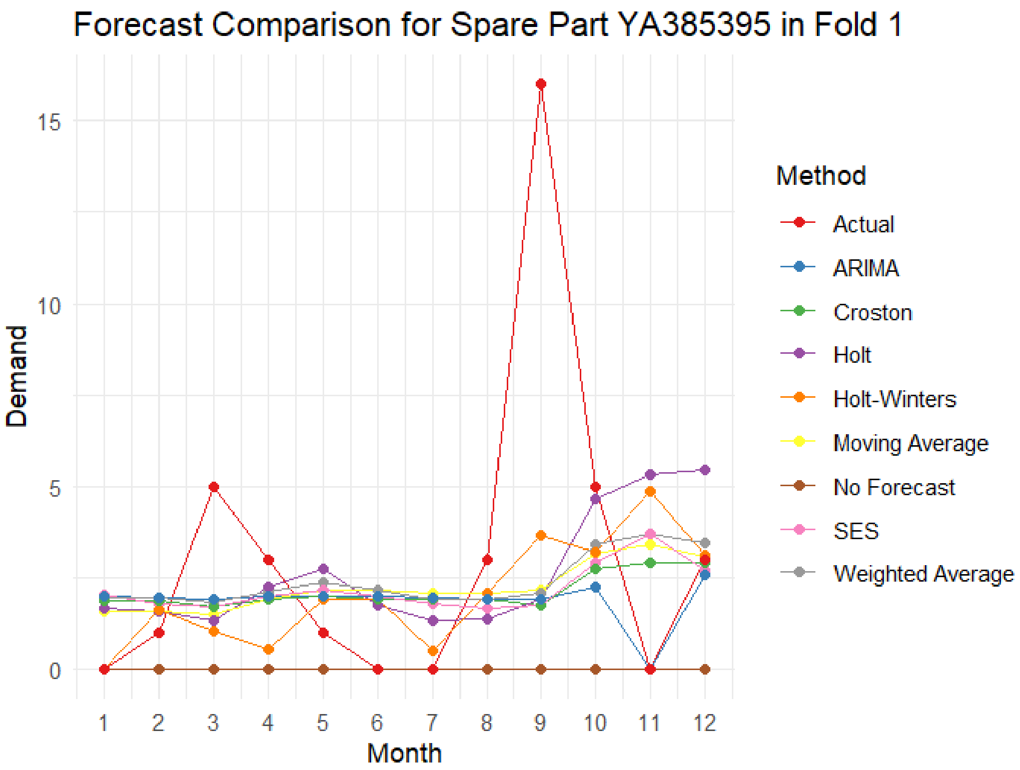

FIGURE 5.8: Forecast Comparison for another Erratic Spare Part

5.3.4 Forecasting Lumpy Demand

There are 7 lumpy demand seasonal spare parts with an average monthly demand of 5. Forecasting with the eight proposed forecasting methods for lumpy demand leads to the results in Table 5.10. These results are the average results of the seasonal spare parts for the five test sets. The results are very bad but this is due to the low average demand. One small mistake influences the SMAPE heavily. When looking at the best performing method based on the sMAPE, (S)ARIMA, shortly followed by Holt-Winters and moving average, score the best. This is not in line with the hypothesis as the moving average and Croston should perform the best. The Croston method is performing almost the worst. Probably, spare parts with noisy demand patterns and differences in train and test sets are causing this.

TABLE 5.10: Average Forecast Performance of Lumpy Demand Seasonal Spare Parts

Forecast Method	sMAPE
Moving Average	117.21%
Weighted Average	118.14%
SES	117.24%
Holt	121.57%
No Forecast	118.57%
Holt-Winters	116.75%
Croston	121.12%
(S)ARIMA	116.61%

Figures 5.9 and 5.10 show that identified seasonal spare parts within a demand class can differ a lot. Figure 5.9 shows a very low average demand. Here, the forecast methods including factors like trend and seasonality do not seem to be accurate. Figure 5.10 has a very strange demand pattern as it has a lot of zero demand periods and in three periods a demand of 100 or 200. This is probably not realistic demand, an error, and noise for the forecast performance. This could be the case why the best performing forecast method is not in line with the hypothesis.

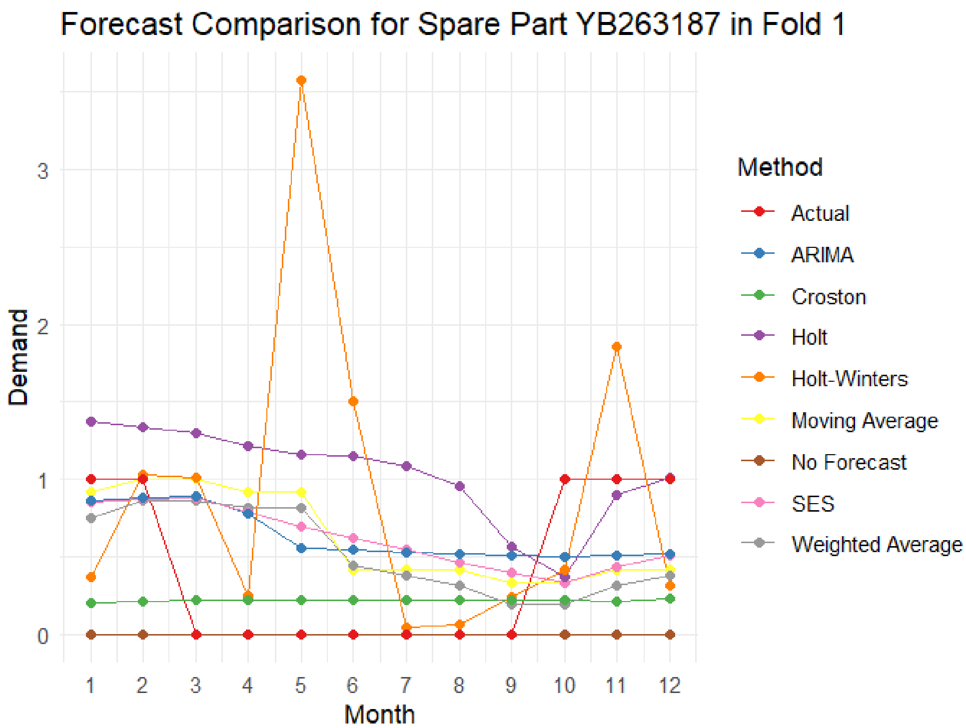

FIGURE 5.9: Forecast Comparison for an Lumpy Spare Part

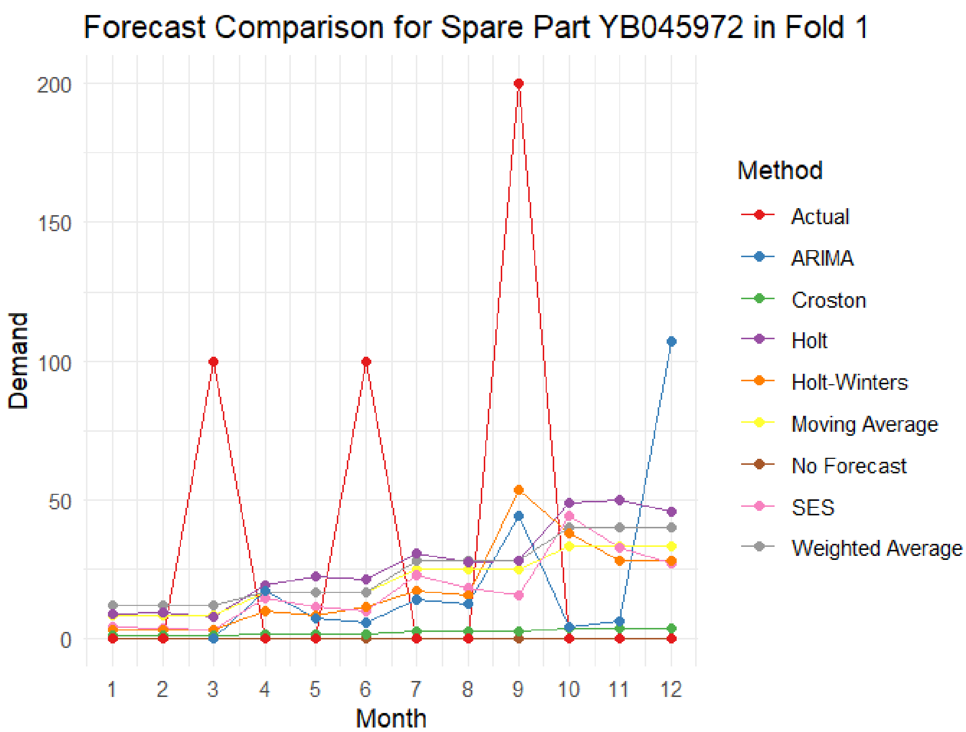

FIGURE 5.10: Forecast Comparison for another Lumpy Spare Part

5.3.5 Comparing Forecasting Demand

We compare the performance of this section with the current performance measured in Table 2.4. At first glance, we see that the tested forecast methods perform worse for all measures except the sMAPE for smooth demand. This does not validate the forecast model as we use better fitting forecast methods. However, we should not forget that we state that the historical demand should be bigger than 60. When applying this criterion to the current selection of 59 seasonal spare parts, only 25 spare parts remain. Table 5.11 evaluates the performance of the 25 seasonal spare parts with a historical demand higher than 60. These performance measures are similar to the performance measures of this chapter as the average forecast performance is 76.18% (without the forecast prediction zero demand) which indicates a valid forecast model. However, we cannot compare the results in Table 5.11 one on one with the results in Table 2.4 as it involves different spare parts.

TABLE 5.11: Average Forecast Performance of NS' Seasonal Spare Parts with a Historical Demand Higher Than 60

Measure	2020	2021	2022	2023
sMAPE	66.28%	73.82%	71.11%	74.75%

To overcome the issue that we cannot compare the forecasts one on one, we calculate the historical NS forecast performance on our 88 selected seasonal spare parts. We compare this performance with the average performance of our model for the same selection of spare parts, for the same time period, in Table 5.12. It shows that our best performing forecast methods perform better, based on the sMAPE, than the historical forecast of NS.

Table 5.12: Comparison Historical Forecast Performance of NS to our Forecast Methods on new Seasonal Spare Parts in 2022

Measures 2022	sMAPE
NS' Forecast	78.81%
Moving Average	74.35%
Croston	76.78%
(S)ARIMA	74.85%

5.3.6 Extra Validations Forecasting Demand

In the earlier subsections of this section, we validate the best performing forecast methods for our selection of seasonal spare parts with the use of the growing-window forward validation method. We have come up with this selection of seasonal spare parts by using a significance threshold of 0.05 in Section 5.1. We want also to validate if the same forecast methods perform the best if we choose another significance threshold. We change the threshold of the identification procedure of Subsection 5.1.1 from 5.0% to 2.5% and 7.5% and find two new sets of seasonal spare parts. Of course, the two validation sets are similar to the original set, so the results will also be similar. However, these validation sets should also have the seasonality characteristic.

Table 5.13 shows how the two validation sets differ in demand class division compared to the original set. For both the validation sets, the division is similar to the division of significance level 0.05. It is logical that less spare parts are selected when a lower alpha is chosen because there are less spare parts with a lower F-test statistic, and vice versa for a higher alpha.

Table 5.13: Validation of Forecasting Seasonal Spare Parts based on Demand Class

Demand Class	Threshold = 0.05	Threshold = 0.025	Threshold = 0.075
Smooth	45 (51.1%)	30 (52.6%)	55 (47.8%)
Intermittent	20 (22.7%)	13 (22.8%)	25 (21.4%)
Erratic	16 (18.2%)	12 (21.1%)	18 (15.4%)
Lumpy	7 (8.0%)	2 (3.5%)	17 (14.5%)
Total	**88 (100%)**	**57 (100%)**	**115 (100%)**

This section already explained which forecast methods suit best for which demand class in case of using the seasonal spare parts resulting from a significance level of 0.05. Table 5.14 shows the conclusions based on the forecasts of all three cases.

Table 5.14: Validation of Forecast Methods based on Demand Class

Demand Class	Threshold = 0.05	Threshold = 0.025	Threshold = 0.075
Smooth	Croston	Croston	SES
Intermittent	(S)ARIMA	(S)ARIMA	(S)ARIMA
Erratic	SES	SES	SES
Lumpy	(S)ARIMA	Moving average	Holt-Winters

When looking at the results of the forecast methods in case of a significance level of 0.025, we have some other conclusions. For almost all demand classes, the same forecast methods perform the best. Only for lumpy demand, the moving average is performing the best. However, these results are not representative as the lumpy demand group only consists of two spare parts.

When looking at the results of the forecast methods in case of a significance level of 0.075, we have some other conclusions. For smooth demand, the SES method performs the best. However, the Croston method performs similar to the SES method which is in line with the significance level of 0.05 case as the Croston does perform the best. For intermittent demand and erratic, the same forecast methods perform the best. For lumpy demand, the Holt-Winters forecast method performs the best. This is not in line with the other cases. However, it stresses that lumpy demand is very chaotic.

We can conclude that the results are representative for the demand pattern group. When these groups differ and contain other spare parts, the results do not differ a lot. This indicates that the results are valid.

5.4 Inventory Analysis

After deciding the best fitting forecast methods, we run the inventory control model according to the approach described in Section 4.4. In this section, we describe the experiments and results of this analysis.

5.4.1 Performing Inventory Analysis

We use all the historical demand data from the previous twelve months resulting from corrective and preventive maintenance of the identified seasonal spare parts. To determine how the safety stock of each spare part should be calculated, we estimate the demand distribution of each spare part. Table 5.15 shows the number of spare parts per demand distribution.

TABLE 5.15: Demand Distributions of Seasonal Spare Parts

Demand Distribution	**# Spare Parts**
Normal	25
Lognormal / Gamma	35
Binomial	13
Negative Binomial	1
Poisson	14

To calculate the reorder points, and order quantities the formulas from Section 4.4 are used. We also test the inventory policy for the last twelve months of the available data. When inventory levels drop below the reorder points, the model orders new spare parts. We use the holding rate, ordering cost, and aggregate target service level for the model shown in Table 5.16. In Table 5.5, we show how the individual target service levels are determined to reach the aggregate service level target of 97.0%. These inputs are based on the parameters that NS uses.

TABLE 5.16: Input Parameters

Input Parameter	**Amount**
Holding Rate	30% per year
Ordering Cost	€25.- per order
Aggregate target service level	97.0%

Table 5.17 shows the performance measures of the inventory model including all relevant costs. These are the summed costs for all seasonal spare parts. Interesting to see is that almost all costs consist of holding costs. This is a result from the fixed order costs per order compared to the relative holding costs.

Table 5.17: Resulting Costs from Inventory Control

Costs	**Amount**	**Proportion**
Holding Costs	€422,555.20	97.47%
Ordering Costs	€10,975.-	2.53%
Total Costs	**€433,530.20**	**100%**

Furthermore, we achieve a certain service level after simulations. The TSL input does not mean that it will actually be achieved. The achieved service levels per class are shown in Table 5.18 where only class CY achieved its TSL as shown by the figures in parenthesis with pp as percentage points. Not every class has the same weights as some classes have a higher usage than others. Incorporating the weight, shown in Table 5.19, leads to an aggregate service level of 87.9%. This is lower than the desired aggregate service level of 97.0%. The obtained service level is probably lower than the TSL because of the starting on hand balance. When we start testing the inventory policy, 46 seasonal spare parts have an inventory position lower than the reorder point. The chance of backorders for this short horizon of twelve months is relatively high influencing the achieved service level.

Table 5.18: Achieved Service Levels (& Difference to Targets) based on ABC- and XYZ-Classification

ABC\XYZ	**Class X**	**Class Y**	**Class Z**
Class A	93.7% (-5.3pp)	85.4% (-10.6pp)	90.0% (-2.0pp)
Class B	-	-	74.4% (-15.6pp)
Class C	-	97.9% (+7.9pp)	75.2% (-9.8pp)

Table 5.19: Weights of ABC- and XYZ-Classification

ABC\XYZ	**Class X**	**Class Y**	**Class Z**
Class A	0.548	0.264	0.116
Class B	-	-	0.029
Class C	-	0.004	0.039

5.4.2 Validation Inventory Analysis

We want to validate our model as we do not obtain our TSL. We hypothesize that we do not achieve the TSL because of the starting on hand balance and our short test period. By increasing the test period, we create a more stable inventory environment. The achieved service level will probably increase towards the target. We simulate the realized demand with the earlier obtained demand distribution and parameters for each spare part.

Figure 5.11 and 5.12 show how the service level and cost performance evolve over time when simulating for different horizon lengths up to 192 months (16 years). The figures show that after 12 months, the model is not yet stable. Thus, the results of Subsection 5.4.1 are indeed influenced by the starting inventory position. In approximately after 48 months, which is called the warm-up period, the model is stable and scores an average service level of 96.6% over periods 49 until 192 which is just below the target. The corresponding costs are on average €846,654.80. A possibility, why the service level does not reach the target, could be a difference between the used and the actual demand distributions. Given that the model approximates the target service level, we conclude that the used inventory model is valid.

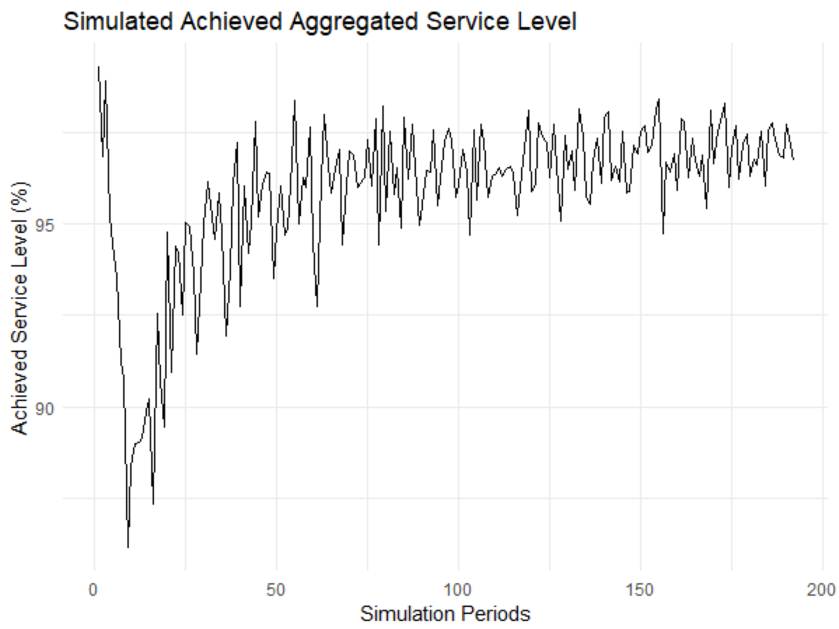

FIGURE 5.11: Achieved Service Level for Different Simulation Runs

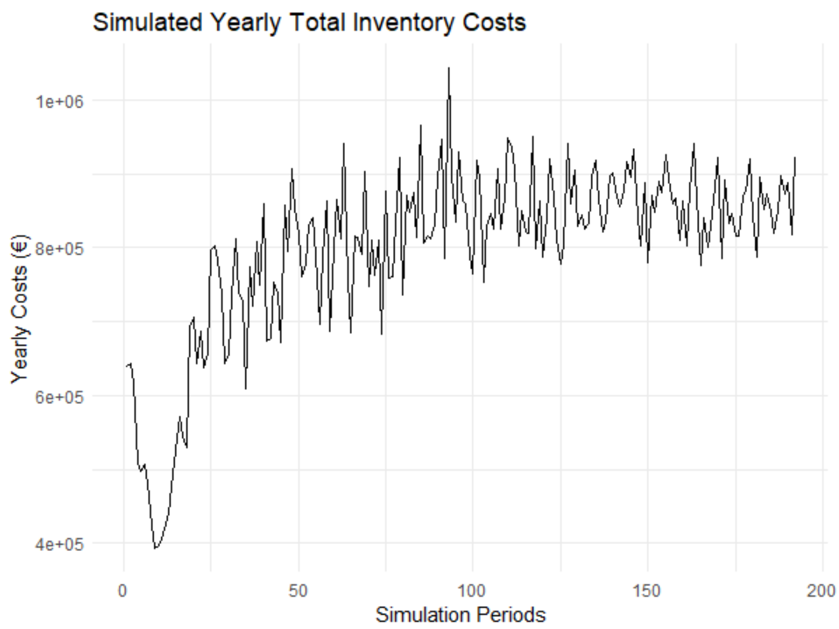

FIGURE 5.12: Inventory Costs for Different Simulation Runs

5.4.3 Comparison Inventory Control

We want to compare the performance of this section with the current performance of NS. Table 2.6 can help us compare the obtained solution with the current situation. However, this is not a one on one comparison as we identified new seasonal spare parts. To overcome the issue that we cannot compare the forecasts one on one, we calculate the historical NS inventory performance on our 88 selected seasonal spare parts. We compare this performance with the performance of our model after the warm-up period for the same selection of spare parts, for the same time period, in Table 5.20. It shows that our inventory model performs better on average stock value (28.3% less) and the usage/value ratio, and performs closer to the service level KPI target. Therefore, we can say that our inventory model performs better.

TABLE 5.20: Inventory Comparison of New Seasonal Spare Parts

Inventory Control KPIs	**NS' Inventory: in 2023**	**Our Model: Testing in 2023**	**Our model: After Warm-Up (Simulation)**
Network Service Level	99.9%	87.9%	96.6%
Total Usage	22,326	22,326	22,000
Average Stock Value	€3.9M	€1.4M	€2.8M
Usage/Value	5,724.6	15,947.1	8,148.1

5.5 Individual Value of Model Components

Currently, we have tested all model components together. To see if every component of this model has an added value, we will test these individually and elaborate on the results in this section.

5.5.1 Identification Seasonality

The identification of seasonality is the first model component, this does not rely on any of the other model components and just generates the input for the other model components. When we would neglect this model component, the output of the model obviously differs as other spare parts are used as input. However, due to this model component, we analyse 88 seasonal spare parts. Therefore, we would say that this model component creates added value. In Section 6.1, we will dive deeper into the contribution of this model component for NS.

5.5.2 Classification

To see if our ABC- and XYZ-classification method results in added value, we experiment with the target service levels. In Subsection 5.4.1, we differentiate the target service levels of the spare parts based on their ABC- and XYZ-class. In this subsection, we set the target service levels of all spare parts to 97.0%.

When simulating the inventory control, it scores an average service level of 96.7% over periods 49 until 192 which is just below the target. The corresponding costs are on average €875,193.4. This means that the service level has gone up with 1.0%, but the costs also increased with a total of 3.0%. We conclude that using the ABC- and XYZ-classification method has an added value.

Furthermore, we look into the added value of the demand pattern classification based on ADI and CV. If we do not use this classification, this means that we will look at one best performing forecast method for all seasonal spare parts. Table 5.21 shows the aggregate forecast performance of the forecast methods based on all 88 seasonal spare parts with the SES method as best performing method with an average sMAPE of 72.43%. When averaging the sMAPE of Croston for smooth demand, (S)ARIMA for intermittent and lumpy demand, and SES for erratic demand, we achieve a SMAPE of 71.92%. This means that choosing a forecast method per specific demand group is beneficial for the overall forecast performance. Therefore, the classification based on ADI and CV is valuable to this research.

TABLE 5.21: Average Forecast Performance of All Seasonal Spare Parts

Forecast Method	**sMAPE**
Moving Average	73.51%
Weighted Average	72.77%
SES	72.43%
Holt	74.18%
No Forecast	169.51%
Holt-Winters	77.70%
Croston	74.62%
(S)ARIMA	75.52%

5.5.3 Forecasting

To analyse whether the forecasting model has added value, we test a forecast method for all seasonal spare parts which is most used in the current forecast process. This is the moving average forecast method. We use this forecast method to forecast demand over lead time as input for the inventory control model. When using the moving average forecast method for all seasonal spare parts in the inventory simulation, it scores an average service level of 96.9% over periods 49 until 192 which is just below the target. The corresponding costs are on average €845,840.4. This means that the service level has gone up with 3.1%, but the costs decrease with a total of 0.9%. We conclude that choosing the forecast method as input for the inventory control based on the performance measure sMAPE does not give guarantee for a better inventory performance. This could be because the performance results of the different forecast methods do not differ much.

5.5.4 Inventory Control

For the inventory control, we should implement an inventory policy with a reorder point and an order size determined by the periodic order quantity to compare our model with the model of NS. With input of the forecasts of NS, we can compare the inventory performance of both inventory policies and determine whether our inventory model has added value. We did not implement this because of time limits.

5.6 Conclusion of Comparing Solutions

This chapter answers the third research question of this thesis. The findings of this chapter give the results of the research approach.

> RQ3: *"What are the best fitting forecast methods and inventory policies for the seasonal spare parts at SCO?"*

First, we find that the demand pattern of 88 spare parts is correlated to the months of the year. Only 9 out of the 88 seasonal spare parts are similar to the selection of NS. Then, we use a three-dimensional classification where the spare parts are divided into ABC- and XYZ classes and the demand classes smooth, intermittent, erratic, and lumpy. Here, most spare parts are classified as class AZ and with a smooth demand pattern. Based on this classification, we set the TSLs and divide the spare parts for forecasting.

Furthermore, we forecasted demand using eight different forecast methods and looked at the results for each demand class. The following conclusions are drawn based on the forecast performance:

- For smooth demand, the Croston forecast method performs the best,
- For intermittent demand, the (S)ARIMA forecast method performs the best,
- For erratic demand, the SES forecast method performs the best,
- For lumpy demand, the (S)ARIMA forecast method performs the best.

For smooth and erratic demand, the outcomes are in line with the theoretical hypotheses. However, for intermittent and lumpy demand, this is not the case. This is probably due to the difference between train and test data and some noisy, unrealistic demand patterns. When we compare our forecast performance with the current performance, our best performing forecast methods perform better based on the sMAPE than the historical forecast of NS. Therefore, we can conclude that our forecast model performs better than the current process.

Our best performing forecast methods are used as input for the inventory control model. Using the demand distributions, the reorder points are calculated as well as the order quantities. Testing the MRP for twelve months results in €433,530.20 costs with an aggregate service level of 87.9%. After simulating the inventory model for several years, the performance stabilizes with an average €846,654.80 of inventory costs and a service level of 96.6%. When we compare this inventory performance with the current performance, our inventory model performs better on average stock value (28.3% less) and the usage/value ratio and is closer to the service level KPI. Therefore, we can conclude that our model is better.

Lastly, we looked at the added value of all of our model components. We conclude that our identification of seasonality and classification methods have added value. But, choosing our best performing forecast methods based on sMAPE, as input for our inventory model, does not give guarantee for the best inventory performance. Furthermore, we do not test the individual value of our inventory control model due to time limits.

The end of this chapter completes step 5 of the MPSM, choosing a solution. Figure 5.13 shows the progress of the research so far.

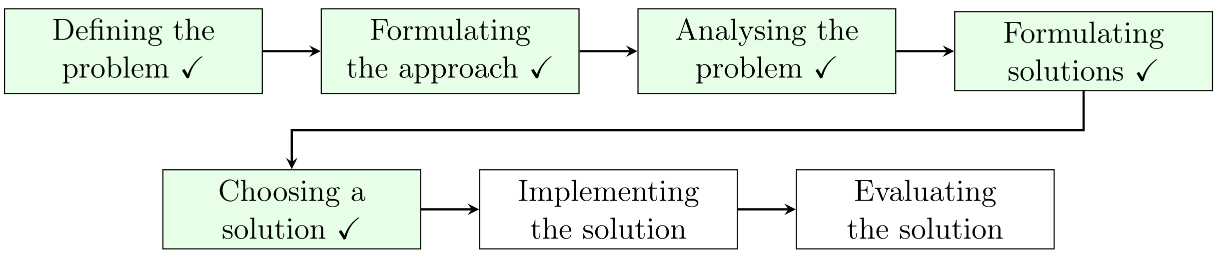

FIGURE 5.13: MPSM Flow Chart - Step 5 Completed

Chapter 6

Implementation and Evaluation Plan

The sixth chapter presents the implementation and evaluation actions to answer the fourth and fifth research questions (RQ) of this thesis.

> RQ4: *"What should SCO do to implement the forecast and inventory tools?"*
>
> RQ5: *"How should SCO evaluate the new forecast methods and inventory control policies?"*

First, this chapter explains the implementation and evaluation steps for the identification of seasonal spare parts. Next, we highlight the adjustments for the classification of spare parts. The chapter ends with the implementation and evaluation plan for the forecast and inventory control of seasonal spare parts.

6.1 Adjust Procedure Seasonal Spare Parts

In Section 5.1, we identify and validate the results of our seasonality analysis. Instead of instantly copy-pasting this method into the working methods of NS, it is important to take a look at the current process flow and see how it fits. The theoretical and practical procedure of the selection and evaluation of the seasonal spare parts list show a discrepancy between the two.

In the theoretical process flow, the supply chain planner (SCP) evaluates the usage of previously identified seasonal spare parts with the use of an evaluation report. In this evaluation report, the SCP proposes seasonal spare parts based on the actual usage of spare parts in previous seasons. This is a quantitative analysis. Besides, the reliability engineer (RE) provides a list of new spare parts that are seasonally relevant and introduced in the operations since the last identification. This is a qualitative analysis. These two actions are the input for the SCP to select the seasonal spare parts for the next season. However, this selection is still a concept list until the RE judges and agrees with the list. As NS considers two seasons, summer and winter, this process is followed twice a year.

In reality, the theoretical process does not comply as described. The difference is that the list of seasonal spare parts is not evaluated with use of the evaluation report. Also, the process is not executed twice a year. This has led to a list where spare parts are present that do not have seasonal or even any demand. Besides, there are spare parts not identified as seasonal, while they show statistical seasonal variations. Conclusively, the list is not up to date.

To improve the process, the theoretical process should be restored in the working methods. Also, the list of current identified seasonal spare parts should be revised. Our statistical analysis will be an addition to this. The idea is to include the statistical analysis within the evaluation report as currently in the report only the interpretation of the SCP is involved. The concept list of seasonal spare parts will still be judged by the RE to see if there could be a qualitative correlation. Additionally, it is interesting to see which spare parts the statistical tool does not select, while the RE selects them according to qualitative characteristics. The addition of the tool creates a more complete process of evaluating seasonal spare parts.

Currently, we evaluated our seasonal concept list with some of the REs at NS. In 17 of the 88 identified seasonal spare parts, the RE saw technical arguments to add them to the seasonal spare part list. Six of these 17 are already within the current selection of NS. This means that our solution creates a contribution of 11 seasonal spare parts for the seasonal list.

Adding the tool is quite an investment as RStudio, which is used as a programming environment, is not integrated into the working methods of the SCO department. This would mean, that one or possibly more employees should get to know the programming language. Additionally, the implementation requires working instructions for the tool. The employees should know how the tool works, and when something does not work, they should know how to fix it. Currently, the used packages of RStudio are not available on the normal working computers of NS, but an employee needs a development computer. However, it is also possible to transpose the tool to Microsoft Excel. This will take less time and makes it easier to adapt. One disadvantage of Excel is that it slows down when analysing many spare parts. Excel could even stop working. A solution to this is a limitation of the size of the input. For example, an user can only analyse 100 spare parts at one time. This would make the identification of seasonal spare parts time-consuming compared to the tool in RStudio. The analysis in RStudio takes less than 10 minutes.

6.2 Adjust Classification Seasonal Spare Parts

We want to highlight one part of the classification in this section. The current classification uses 18 months of data to classify spare parts based on demand and value. This is not appropriate for seasonal spare parts. Seasonal spare parts have a higher demand in a certain period compared to the rest of the year. When classifying based on 18 months of demand data, just after the peak season, the high demand is included twice and spare parts could be classified as more important. However just before the season starts, the seasonal demand is included only once and spare parts could be classified as not important. Changing the classification based on data of the past 12 or 24 months would be more appropriate. This is a small adjustment in the inventory software Servigistics XelusParts. We recommend evaluating this classification constantly. Despite the fact that classifying demand pattern classes has added value for the average forecast performance, we do not recommend using it as the proposed forecast methods never clearly outperformed others. Additionally, these best performing methods were no guarantee for a good inventory performance.

6.3 Adjust Forecasting Seasonal Spare Parts

Currently, the SCP analyses historical data and looks at only four methods out of the fourteen possibilities of Servigistics XelusParts. The process of choosing a suitable method involves altering between forecast methods and assessing which would fit the best. Due to a lack of knowledge about using forecast methods that account for seasonal variations, there is an evaluation step in the forecast process where the planner can manually adjust the forecast to meet seasonal demand.

Currently, there is no forecast performance available for the SCP. This makes it difficult to know how the SCP is performing and if the appropriate forecast method is used. To create a more complete process, we recommend to implement this measure. Furthermore, based on the results of Section 5.3, forecast methods like the Croston and (S)ARIMA method show us that more forecast methods should be involved within the forecast process. This requires an investment in knowledge for the SCP.

The forecast process should also change regarding the manual adjustments. Currently, the sentiment is that there must not be any stockouts leading to a standstill under any circumstances. Therefore, the manual adjustment option is used a lot to increase the inventory level resulting in forecasting too much demand throughout the whole year. This is not a correct way to intervene in the inventory process. Adjusting the reorder levels and reorder quantity is more appropriate. The underlying problem of the sentiment is that there are no clear costs associated with a stockout. With a widely accepted approximation of the stockout costs, a good equilibrium can be established between stockouts and the inventory level. Overall, the sentiment should be that a stockout is accepted under good consideration of costs.

Again, adding the tool or the performance measurement is quite an investment as RStudio is not integrated into the working methods of the SCO department. However, it is also possible to transpose the performance measurement to a standard format in Microsoft Excel. When exporting and using the correct data, the SCP can easily evaluate the performance. Building this tool in Excel will cost a lot of time as this involves many forecast methods that work differently. If NS wants to use this analysis, we suggest evaluating this twice a year to stay up to date, in advance of the seasons.

A more convenient implementation is suggesting some requirements for the new inventory software that NS wants to use within two years. These requirements could be used for the market consultation of the new software. These are the requirements that we would suggest adding to the current requirements:

- The software should give insight in the forecast performance,
- The software should include the forecast methods that we tested.

6.4 Adjust Inventory Control Seasonal Spare Parts

Currently, one inventory method is used with a certain set of inventory parameters. We do not recommend changing the current way of controlling the inventory. However, the current inventory program Servigistics XelusParts does not show any insights into the possible inventory costs resulting from the current parameters and policy. The tool that we created gives insights into these costs. Implementing our tool would create the opportunity to test different inventory parameters and see what works best.

Again, adding the tool or the performance measurement is quite an investment as RStudio is not integrated into the working methods of the SCO department. It is possible to transpose the performance measurement to a standard format in Microsoft Excel. However, constructing this tool in Excel will be time-consuming. If it works, the SCP can easily evaluate the performance. If NS wants to use this analysis, we suggest evaluating this twice a year to stay up to date, in advance of the seasons.

A more convenient implementation is suggesting some requirements for the new inventory software that NS wants to use within two years. These requirements could be used for the market consultation of the new software. These are the requirements that we would suggest adding to the current requirements:

- The software should include identifying demand distributions with corresponding inventory parameter calculations,
- The software should include direct feedback on the inventory performance, for example feedback on the costs and service level.

6.5 Conclusion of Implementation and Evaluation Plan

This chapter answers the fourth and fifth research questions of this thesis. The findings of this chapter give recommendations about the implementation and evaluation plan.

> RQ4: *"What should SCO do to implement the forecast and inventory tools?"*

> RQ5: *"How should SCO evaluate the new forecast methods and inventory control policies?"*

First, we state that the identification tool should be implemented within the seasonal evaluation process. The outcomes of the tool should be evaluated by a RE and supplemented by qualitative reasoning. After this evaluation, 17 of the 88 seasonal spare parts are acknowledged by the REs, from which 11 seasonal spare parts are an addition to the current list. The investment in implementing the identification tool in RStudio tool is quite big. However, it is also possible to transpose the tool into Excel. The tool should be evaluated twice a year, just after a season ends.

Furthermore, we recommend changing the demand classification based on the previous 18 months of data to using data from the past 12 or 24 months as this considers a complete cycle of a year. Adapting the classification based on ADI and CV is not necessary as the corresponding forecast method never outperformed other forecast methods.

The new proposed forecast methods and performance measurements should be considered in the forecasting process which requires an investment in knowledge and in transposing the measurement in Excel. Also, changing the forecast process regarding manual adjustment is necessary. A widely accepted estimation of stockout costs could help this change.

Lastly, the implementation of the performance measurements of the inventory control is a big investment. However, it is worth the insights that will be obtained. The SCP will gain insight into how different inventory parameters will influence the inventory performance which is currently not present.

The end of this chapter completes steps 6 and 7 of the MPSM, implementing the solution and evaluating the solution. Figure 6.1 shows the completed research flow chart.

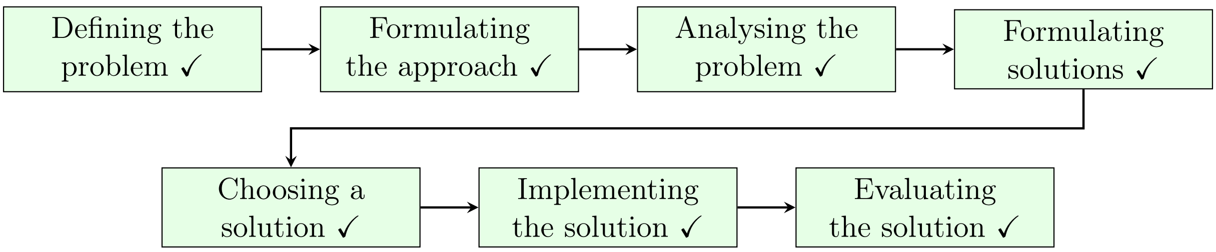

Figure 6.1: MPSM Flow Chart - Step 7 Completed

Chapter 7

Conclusions and Recommendations

This chapter includes the conclusions and recommendations of this research. First the answers to the research questions are presented, followed by the research limitations. Subsequently the recommendations to NS and regarding the future research are described.

7.1 Conclusions

This section draws conclusions of the research and answers the main research question:

> *"How can the forecasting methods and the inventory policies of seasonal spare parts be changed to improve the performance of the SCO department at NS?"*

To solve the main research question, we will discuss all research questions one by one. In Chapter 2, we perform a context analysis to solve the research question *"What does the current forecast process and inventory control of SCO look like?"*. The forecast and inventory control involves corrective and preventive maintenance which is divided per train series. For maintenance, there are exchange and wear spare parts that differ in inventory management. Currently, SCO identifies 59 spare parts as seasonal. Those parts show seasonal variation in either winter or summer. However, this list needs a reclassification because of unclear seasonal demand patterns, the lack of demand, and unknown lead times. Also, SCO tracks whether these parts are critical, indicating that when they fail, trains are not allowed to drive any further. The forecast process does not use many extensive forecast methods and is often manually adjusted to cover seasonal demand. However, the performance of the forecasts including manual adjustments is worse than the forecasts excluding manual adjustments which indicates that the supply chain planner often forecasts the demand too high. Furthermore, the network service level performance of the inventory control for the seasonal spare parts is similar to the performance for all spare parts. The service level is above the norm of SCO, while the BWOM is below the norm. However, if we let the fill rate get worse, the average train units waiting on materials KPI would also get worse. Another service level measure, like the order line fill rate, would be more appropriate as the shortage of one spare part within an order causes a delay of the complete order. The order line fill rate will be lower than the fill rate and therefore the order line fill rate would represent the average train units waiting on materials KPI better.

In Chapter 3, we perform a literature study to solve the research question *"What are possible forecast methods and inventory policies for the spare parts to apply in the context of SCO?"* We find that multiple regression is a time-efficient method to identify the presence of seasonal variations in comparison to an exploratory data analysis. When seasonal variations are present, forecast methods, like the Holt Winters' trend-seasonal method, or (S)ARIMA can be used while the Croston method performs well when observing an intermittent demand pattern. Other mentioned methods, like the moving average, SES, and Holt method could be relevant in case of smooth demand. Also, we conclude that we will be controlling inventory using the (s,Q)-policy as this is the most used policy in an overview of examples of combining forecast methods with inventory control policies. Lastly, the literature confirms the approach of first classifying spare parts, then forecasting, inventory, and ending with performance assessment.

In Chapter 4 and 5, we describe the model and present the results of this model to solve the research question *"What are the best fitting forecast methods and inventory policies for the seasonal spare parts at SCO?"* We use multiple regression to determine whether the demand pattern of a spare part is correlated to the months of the year. We find that the demand pattern of 88 spare parts is correlated to the months of the year. We call these the seasonal spare parts. Only nine out of the 59 original seasonal spare parts reappear in this list. With use of the ABC- and XYZ-classification, we gave guidance in choosing the target service levels. The classification method based on ADI and CV helps choosing fitting forecast methods as specific forecast methods perform well on specific demand patterns. Furthermore, we tuned parameters and forecasted demand using eight different forecast methods with the use of growing-window forward validation. We interpret the results for each demand class. The following conclusions are drawn based on the forecast performance:

- Smooth: The Croston forecast method performs the best,
- Intermittent: The (S)ARIMA forecast method performs the best,
- For erratic demand, the SES forecast method performs the best,
- For lumpy demand, the (S)ARIMA forecast method performs the best.

For smooth and erratic demand, the outcomes are in line with the theoretical hypotheses. However, for intermittent and lumpy demand, this is not the case. This is probably due to the difference between train and test data and some noisy, unrealistic demand patterns. When we compare forecast performances, we can conclude that our model performs better than the current forecast process. Several forecast methods like the moving average, Croston and (S)ARIMA score better than the current forecast performance. Therefore, we recommend to keep on testing all of our forecast methods and using them according to the corresponding performance. Additionally, the inventory control model calculates the reorder points, and order quantities which are used to simulate the inventory performance with unseen test data. After simulating the inventory model for several years, the performance stabilizes with an average €846,654.80 of inventory costs and a service level of 96.6%. When we compare inventory performances, our inventory model performs better on average stock value (28.6% less) and the usage/value ratio and is closer to the service level KPI. Therefore, we can conclude that our model is better. NS performs too high on service level. So, decreasing towards the norm will save costs.

Lastly, the identification of seasonality and classification methods have proved added value. But, choosing our best performing forecast methods based on sMAPE, as input for our inventory model, does not give guarantee for the best inventory performance. This means that another forecast performance measure might have been more appropriate.

7.2 Recommendations

In Chapter 6, we make an implementation and evaluation plan to solve the research questions *"What should SCO do to implement the forecast and inventory tools?"* and *"How should SCO evaluate the new forecast methods and inventory control policies?"* This section describes all relevant recommendations to the SCO department of NS as they are the problem owners. These are recommendations based on the results of the research and the lessons learned during the research.

It is interesting to see that forecast methods including seasonal factors do not necessarily perform well on our selection of seasonal spare parts. Possibly, the demand patterns are too volatile to predict accurately. The question arises as to whether the identification of seasonal spare parts is necessary. We think it is beneficial to identify and monitor them as there are inventory decisions where this characteristic is helpful. First of all, the SCP could have information in advance of the season that demand will increase or decrease. This could have a crucial impact on the inventory performance if SCO is not prepared for this. Besides, knowledge of seasonal spare parts is interesting for supplier agreements. SCO can make agreements about decreasing lead times in the summer which can decrease the mean and variance of the lead times. This results in less uncertainty in the supply chain. Less uncertainty means less safety stock and less holding costs while achieving the same service level.

As suggested in Section 6.1, we recommend identifying seasonal spare parts twice a year just after the summer and winter seasons. Currently, there is no fixed identification process and the seasonal spare part list is outdated. Using the identification tool will help to analyse a big population of spare parts at once and is an addition to the current procedure. Combining the outcomes of this tool together with the technical knowledge of reliability engineers, we found 11 new seasonal spare parts.

Furthermore, use a classification method based on the previous 12 or 24 months of demand data as this is more appropriate for seasonal spare parts. The current use of 18 months of demand data gives a distorted view as a peak season can appear once or twice within the 18 months depending on the time the classification takes place.

Additionally, we recommend implementing measurements for the forecast and inventory process of seasonal spare parts. The problem of SCO is about the forecast and inventory process of seasonal spare parts. As we analysed the current situation in Chapter 2, we concluded that there is no forecast performance evaluation available for any spare part. This way supply chain planners can't see in the long term which forecast methods perform well for which kind of spare parts. Because there is no insight into this performance, supply chain planners will manually adjust the forecast methods which increases the forecast error. A supply chain planner will not learn from mistakes as they have no insights into their performance. The inventory control performance for seasonal spare parts is also not present. This makes the consequences of the actions of supply chain planners invisible. Forecasting too many spare parts has become a common practice as there are no triggers that there is too much inventory. On the subject of which forecast methods to use, we see no single forecast method performing the best. Also, forecast methods including a seasonal factor are no guarantee for success for the seasonal spare parts of NS. We recommend to keep on testing all forecast methods and using them according to the corresponding performance.

Lastly, we recommend learning the supply chain planners how forecast and inventory control parameters work and how to use these. Currently, supply chain planners can not access parameters and do not know how these work. In the fixed working environment of supply chain planners, it is difficult to be adaptive to changes. When being more adaptive and proactive, the KPIs will be easier to manage.

7.3 Limitations and Further Research

In this section, we will highlight the limitations of this research and areas of improvement or further research.

Currently, we identified seasonal spare parts based on monthly correlation. For further research, it is more accurate to look at quarterly seasonality as this is more in line with the climate and weather conditions. Also, it may be interesting to see if including weather variables within the model as we think the weather is the biggest influencing factor of seasonality. Furthermore, we assumed that seasonality is deterministic while seasonality can be non-stationary and evolve over time. Testing and considering whether seasonality is deterministic or stochastic, as proposed by the HEGY test procedure, could lead to better forecasting performance. Hence, this is something for further research.

Furthermore, we forecasted with several forecast methods that are available in RStudio. Possibly, another programming language can test different forecast methods more effectively. For example, there are Croston variations that are improved and can perform better than the current implemented method in RStudio. Also, we did not evaluate machine learning methods as there is little demand data available. In the future, when there is more demand data available, it would be interesting to see whether machine learning methods would perform better than the methods we evaluated in this research.

In the inventory model, we assumed the total system with one inventory and one demand. In reality, the system is a multi-echelon system with multiple locations in hierarchy as Chapter 2 explains. Considering a multi-echelon system would give a more realistic view of the inventory process. Additionally, we used a basic inventory model. In practice, there are interventions like emergency orders to speed up the delivery process decreasing the backorders. Lastly, we did not include the inventory process of exchange parts within our model while this is a characteristic of a large proportion of the spare part population. Including the flow of exchange parts, will generate a more realistic model and more realistic performance measures.

Bibliography

Poul Alstrøm and Per Madsen. Evaluation of forecast models used for inventory control during a product's life cycle: A simulation study. *International Journal of Production Economics*, 35 (1):191–200, 1994. ISSN 0925-5273. doi: https://doi.org/10.1016/0925-5273(94)90081-7. URL https://www.sciencedirect.com/science/article/pii/0925527394900817.

Sven Axsäter. *Inventory Control*. Springer New York, NY, 2006.

Andrea Bacchetti and Nicola Saccani. Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. *Omega*, 40(6): 722–737, 2012. ISSN 0305-0483. doi: https://doi.org/10.1016/j.omega.2011.06.008. URL https://www.sciencedirect.com/science/article/pii/S0305048311001484. Special Issue on Forecasting in Management Science.

Devon K. Barrow and Nikolaos Kourentzes. Distributions of forecasting errors of forecast combinations: Implications for inventory management. *International Journal of Production Economics*, 177:24–33, 2016. ISSN 0925-5273. doi: https://doi.org/10.1016/j.ijpe.2016.03.017. URL https://www.sciencedirect.com/science/article/pii/S0925527316300226.

Christoph Bergmeir and José M. Benítez. On the use of cross-validation for time series predictor evaluation. *Information Sciences*, 191:192–213, 2012. ISSN 0020-0255. doi: https://doi.org/10.1016/j.ins.2011.12.028. URL https://www.sciencedirect.com/science/article/pii/S0020025511006773. Data Mining for Software Trustworthiness.

Swapnil Bhalla, Erlend Alfnes, Hans-Henrik Hvolby, and Fabio Sgarbossa. Advances in spare parts classification and forecasting for inventory control: A literature review. *IFAC-PapersOnLine*, 54(1):982–987, 2021. ISSN 2405-8963. doi: https://doi.org/10.1016/j.ifacol.2021.08.118. URL https://www.sciencedirect.com/science/article/pii/S2405896321008703. 17th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2021.

Olivier Darné and Claude Diebolt. A note on seasonal unit root tests. *Quality and Quantity*, 36: 305–310, 08 2002. doi: 10.1023/A:1016032601197.

Dinesh Kumar Dhoka. "XYZ" Inventory Classification & Challenges. *IOSR Journal of Economics and Finance*, 2:23–26, 2013. URL https://api.semanticscholar.org/CorpusID:167870318.

José Roberto do Rego and Marco Aurélio de Mesquita. Demand forecasting and inventory control: A simulation study on automotive spare parts. *International Journal of Production Economics*, 161:1–16, 2015. ISSN 0925-5273. doi: https://doi.org/10.1016/j.ijpe.2014.11.009. URL https://www.sciencedirect.com/science/article/pii/S0925527314003594.

Hans Heerkens and Arnold van Winden. *Solving Managerial Problems Systematically*. Noordhoff Uitgevers, 2017. ISBN 9789001887957. Translated into English by Jan-Willem Tjooitink.

G. Heinecke, A.A. Syntetos, and W. Wang. Forecasting-based SKU classification. *International Journal of Production Economics*, 2013. URL https://doi.org/10.1016/j.ijpe.2011.11.020.

H. Hewage, H. Perera, and S. De Baets. Forecast adjustments during post-promotional periods. *EUROPEAN JOURNAL OF OPERATIONAL RESEARCH*, 300(2):461–472, 2022. ISSN 0377-2217. URL http://doi.org/10.1016/j.ejor.2021.07.057.

Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An Introduction to Statistical Learning*. Springer New York, NY, 2022. ISBN 978-1-0716-1420-4. doi: https://doi.org/10.1007/978-1-0716-1418-1.

P. Jiang, X. Liu, Y. Huang, and Y. Yuan. An adaptive data-driven model for erratic demand forecasting. *IFAC-PapersOnLine*, 50(1):9020–9025, 2017. ISSN 2405-8963. doi: https://doi.org/10.1016/j.ifacol.2017.08.1584. URL https://www.sciencedirect.com/science/article/pii/S2405896317321808. 20th IFAC World Congress.

Matthew Lindsey and Robert Pavur. *Assessing a modification to Croston's method to incorporate a seasonal component*, pages 185–195. Advances in Business and Management Forecasting. Advances in Business and Management Forecasting, 2013. ISBN 9781781903315. doi: 10.1108/S1477-4070(2013)0000009015.

Spyros Makridakis, Evangelos Spiliotis, and Vassilis Assimakopoulos. The m5 accuracy competition: Results, findings and conclusions. *International Journal of Forecasting*, page 45, 10 2020.

Amer Malki, El-Sayed Atlam, Aboul Ella Hassanien, Ashraf Ewis, Guesh Dagnew, and Ibrahim Gad. Sarima model-based forecasting required number of covid-19 vaccines globally and empirical analysis of peoples' view towards the vaccines. *Alexandria Engineering Journal*, 61(12): 12091–12110, 2022. ISSN 1110-0168. doi: https://doi.org/10.1016/j.aej.2022.05.051. URL https://www.sciencedirect.com/science/article/pii/S1110016822003714.

Pranab Mandal. Handout in Statistics and Probability Course, September 2023.

Indry Maretania, Muhammad Rummy Alfadjri, Pradieqta Uli Paramesywarie, and Rahmat Nurcahyo. Comparison of Double Exponential and Single Exponential Smoothing Accuracy in Krakatau Steel Demand Forecasting Fitted Model. *IEOM Society*, 2021. URL https://www.ieomsociety.org/proceedings/2021india/87.pdf.

Mitchell A. Millstein, Liu Yang, and Haitao Li. Optimizing ABC inventory grouping decisions. *International Journal of Production Economics*, 2014. URL https://doi.org/10.1016/j.ijpe.2013.11.007.

N. Ahmadi Mobarakeh, M.K. Shahzad, A. Baboli, and R. Tonadre. Improved forecasts for uncertain and unpredictable spare parts demand in business aircraft's with bootstrap method. *IFAC-PapersOnLine*, 50(1):15241–15246, 2017. ISSN 2405-8963. doi: https://doi.org/10.1016/j.ifacol.2017.08.2379. URL https://www.sciencedirect.com/science/article/pii/S2405896317331981. 20th IFAC World Congress.

Nikolas Ulrich Moroff, Ersin Kurt, and Josef Kamphues. Machine learning and statistics: A study for assessing innovative demand forecasting models. *Procedia Computer Science*, 180: 40–49, 2021. ISSN 1877-0509. doi: https://doi.org/10.1016/j.procs.2021.01.127. URL https://www.sciencedirect.com/science/article/pii/S1877050921001654. Proceedings of the 2nd International Conference on Industry 4.0 and Smart Manufacturing (ISM 2020).

Nederlandse Spoorwegen. Oorlog en Fusie - Geschiedenis van NS, 2024. URL https://www.ns.nl/over-ns/wie-zijn-wij/geschiedenis/oorlog-en-fusie.html. Accessed: 2024-01-26.

Irena Nowotyńska. An Application of Xyz Analysis in Company Stock Management. *Modern Management Review*, pages 77–86, 2013. URL https://api.semanticscholar.org/CorpusID:59036624.

Mwana Said Omar and Hajime Kawamukai. Prediction of ndvi using the holt-winters model in high and low vegetation regions: A case study of east africa. *Scientific African*, 14:e01020, 2021. ISSN 2468-2276. doi: https://doi.org/10.1016/j.sciaf.2021.e01020. URL https://www.sciencedirect.com/science/article/pii/S2468227621003215.

Dennis Prak, Ruud Teunter, and Aris Syntetos. On the calculation of safety stocks when demand is forecasted. *European Journal of Operational Research*, 256(2):454–461, 2017. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2016.06.035. URL https://www.sciencedirect.com/science/article/pii/S0377221716304568.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2024. URL https://www.R-project.org/.

A. Regattieri, M. Gamberi, R. Gamberini, and R. Manzini. Managing lumpy demand for aircraft spare parts. *Journal of Air Transport Management*, 11(6):426–431, 2005. ISSN 0969-6997. doi: https://doi.org/10.1016/j.jairtraman.2005.06.003. URL https://www.sciencedirect.com/science/article/pii/S0969699705000530.

Matthias Schnaubelt. *A comparison of machine learning model validation schemes for non-stationary time series data*. Number 11/2019 in 1867-6707. Friedrich-Alexander-Universität Erlangen-Nürnberg, Institute for Economics, Nürnberg, 2019. URL https://hdl.handle.net/10419/209136.

Mahya Seyedan, Fereshteh Mafakheri, and Chun Wang. Order-up-to-level inventory optimization model using time-series demand forecasting with ensemble deep learning. *Supply Chain Analytics*, 3:100024, 2023. ISSN 2949-8635. doi: https://doi.org/10.1016/j.sca.2023.100024. URL https://www.sciencedirect.com/science/article/pii/S2949863523000237.

Edward A. Silver, David F. Pyke, and Douglas J. Thomas. *Inventory and Production Management in Supply Chains*. CRC Press, 4th edition, 2016. doi: 10.1201/9781315374406. URL https://doi.org/10.1201/9781315374406.

Aris A. Syntetos and John E. Boylan. Smoothing and adjustments of demand forecasts for inventory control. *IFAC Proceedings Volumes*, 39(3):173–178, 2006. ISSN 1474-6670. doi: https://doi.org/10.3182/20060517-3-FR-2903.00104. URL https://www.sciencedirect.com/science/article/pii/S1474667015358249. 12th IFAC Symposium on Information Control Problems in Manufacturing.

Ruud H. Teunter, Aris A. Syntetos, and M. Zied Babai. Intermittent demand: Linking forecasting to inventory obsolescence. *European Journal of Operational Research*, 214(3): 606–615, 2011. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2011.05.018. URL https://www.sciencedirect.com/science/article/pii/S0377221711004437.

Lorenzo Tiacci and Stefano Saetta. An approach to evaluate the impact of interaction between demand forecasting method and stock control policy on the inventory system performances. *International Journal of Production Economics*, 118(1):63–71, 2009. ISSN 0925-5273. doi: https://doi.org/10.1016/j.ijpe.2008.08.010. URL https://www.sciencedirect.com/science/article/pii/S092552730800248X. Special Section on Problems and models of inventories selected papers of the fourteenth International symposium on inventories.

Juan R. Trapero, Enrique Holgado de Frutos, and Diego J. Pedregal. Demand forecasting under lost sales stock policies. *International Journal of Forecasting*, 2023. ISSN 0169-2070. doi: https://doi.org/10.1016/j.ijforecast.2023.09.004. URL https://www.sciencedirect.com/science/article/pii/S0169207023000961.

Matthieu Van der Heijden. Lecture notes in Advanced Inventory Management, November 2022.

E. van Wingerden, R.J.I. Basten, R. Dekker, and W.D. Rustenburg. More grip on inventory control through improved forecasting: A comparative study at three companies. *International Journal of Production Economics*, 157:220–237, 2014. ISSN 0925-5273. doi: https://doi.org/10.1016/j.ijpe.2014.08.018. URL https://www.sciencedirect.com/science/article/pii/S092552731400276X. The International Society for Inventory Research, 2012.

Appendix A

Demand Graphs Seasonal Spare Parts

In this appendix, we will show two demand graphs of seasonal spare parts which are classified by NS. Figure A.1 shows an example of a seasonal spare part which is prone to failures in the summer months. However, this demand pattern is not very clear as in some winter months the demand is also higher than the trendline.

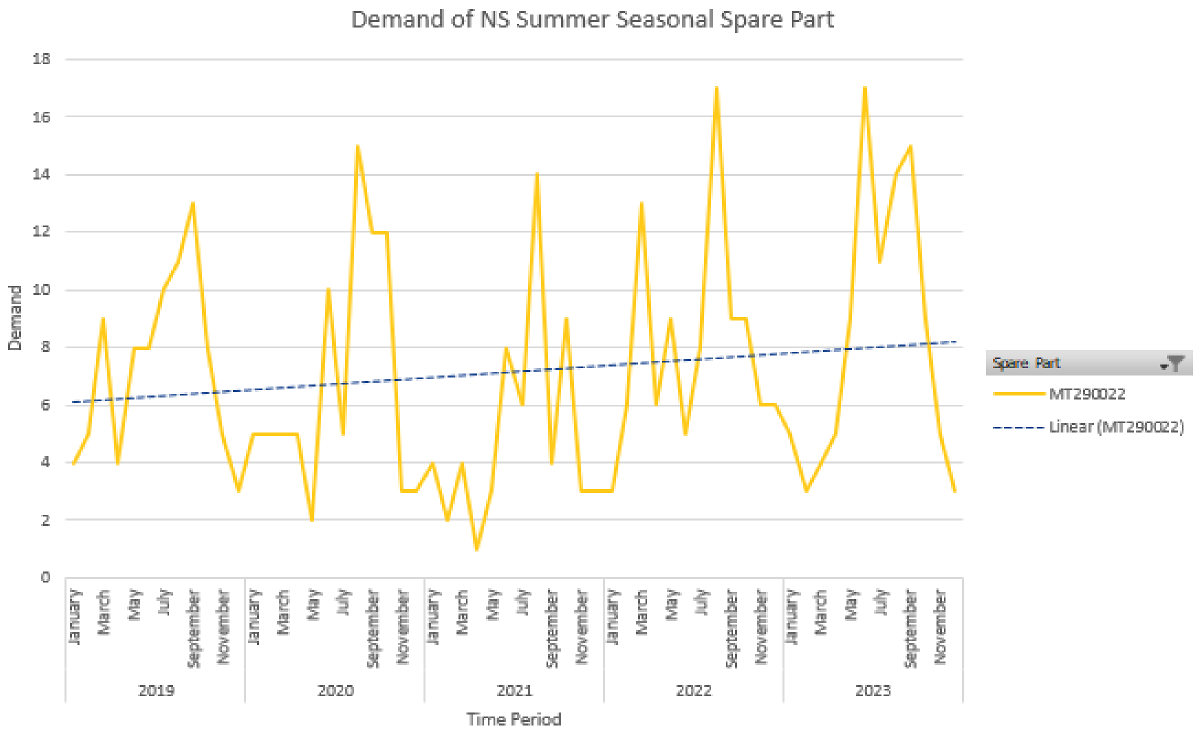

Figure A.1: Demand of a NS Summer Seasonal Spare Part over Time

Figure A.2 shows the demand pattern of a seasonal spare part, identified by NS, which does not have a clear seasonal demand pattern. During summer and winter months, the demand is higher and lower than the trendline. Evaluating the seasonal spare parts of NS showed that most spare parts do not have a clear seasonal demand pattern, only six seasonal spare parts do.

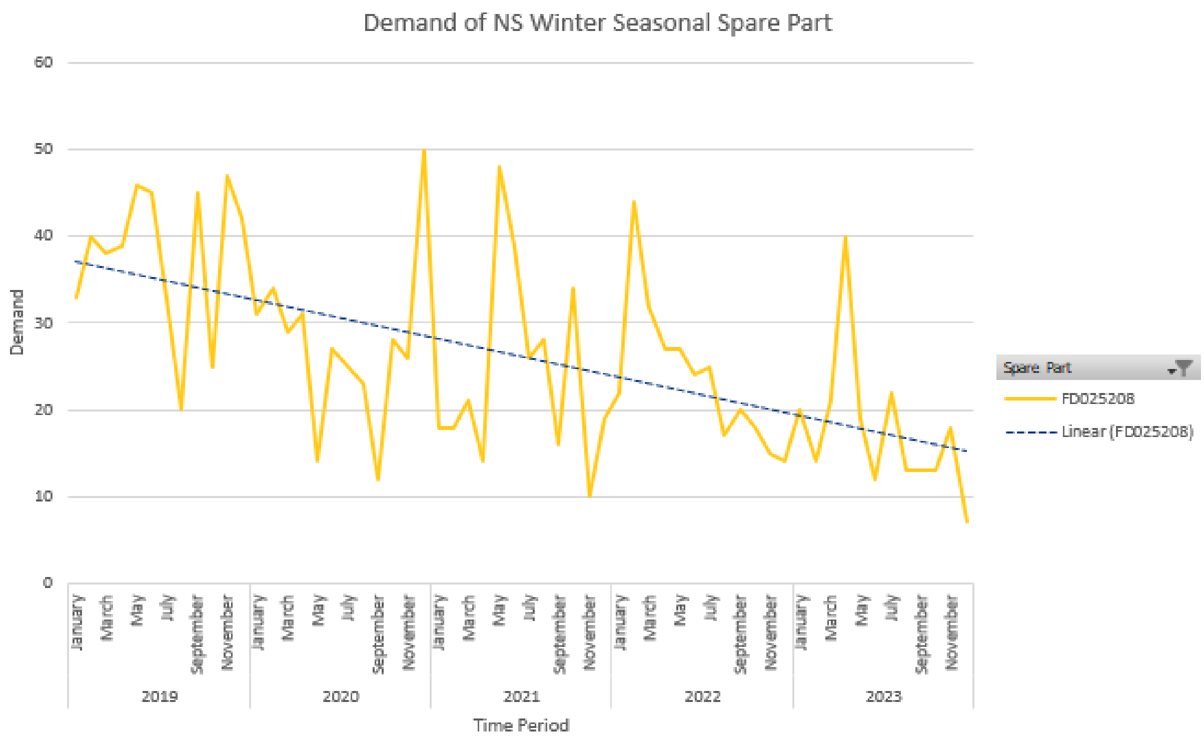

FIGURE A.2: Demand of a NS Winter Seasonal Spare Part over Time

Appendix B

Transposing ESPRC Calculations

In this appendix, we show transposing the expected shortage per replenishment cycle (ESPRC) for the binomial and negative binomial distributed demand. We need the formula for the ESPRC shown in Equation B.1.

$$\text{ESPRC}(s) = \int_{s}^{\infty} (x - s) f(x)\, dx \tag{B.1}$$

B.1 ESPRC Binomial Distribution

For binomial distributed demand, we use the summation of the binomial probability mass function. Using Equation B.1, we get Equation B.2. However, we want to simplify this equation to make the ESPRC calculations easier.

$$\text{ESPRC}(s) = \sum_{s}^{\infty} (x - s) \binom{n}{x} p^x (1-p)^{n-x} \tag{B.2}$$

To make the ESPRC calculations easier, we transform the Equation B.2:

$$\begin{aligned}
\text{ESPRC}(s) &= \sum_{s}^{\infty} (x - s) \binom{n}{x} p^x (1-p)^{n-x} \\
&= \sum_{s}^{\infty} x \binom{n}{x} p^x (1-p)^{n-x} - s \sum_{s}^{\infty} \binom{n}{x} p^x (1-p)^{n-x} \\
&= 1 - \sum_{0}^{s} x \binom{n}{x} p^x (1-p)^{n-x} - s \left(1 - \sum_{0}^{s} \binom{n}{x} p^x (1-p)^{n-x} \right)
\end{aligned}$$

The second summation is in the form of the cumulative distribution function of a binomial distribution $F(s; n, p)$. We also aim to transform the first sum:

$$\sum_{x=0}^{s} x \binom{n}{x} p^x (1-p)^{n-x}$$

to the form involving the cumulative distribution function of a binomial distribution F. Therefore, we will zoom in on the transformation of this sum.

Recall that the binomial coefficient is defined as:

$$\binom{n}{x} = \frac{n!}{x!(n-x)!}$$

Then, the sum can be rewritten as:

$$\begin{aligned}
\sum_{x=0}^{s} x\binom{n}{x}p^x(1-p)^{n-x} &= \sum_{x=0}^{s} x\frac{n!}{x!(n-x)!}p^x(1-p)^{n-x} \\
&= \sum_{x=0}^{s} n\frac{x}{n}\frac{n!}{x!(n-x)!}p^x(1-p)^{n-x} \\
&= n\sum_{x=0}^{s} \frac{x}{n}\frac{n!}{x!(n-x)!}p^x(1-p)^{n-x} \\
&= n\sum_{x=0}^{s-1} \frac{x+1}{n}\frac{n!}{(x+1)!(n-x-1)!}p^{x+1}(1-p)^{n-(x+1)} \\
&= n\sum_{x=0}^{s-1} \frac{(n-1)!}{x!(n-1-x)!}p^{x+1}(1-p)^{n-x-1} \\
&= np\sum_{x=0}^{s-1} \frac{(n-1)!}{x!(n-1-x)!}p^{x}(1-p)^{n-1-x} \\
&= np\sum_{x=0}^{s-1} \binom{n-1}{x}p^{x}(1-p)^{n-1-x} \\
&= np \times F(s-1; n-1, p)
\end{aligned}$$

Thus, we have shown that:

$$\sum_{x=0}^{s} x\binom{n}{x}p^x(1-p)^{n-x} = np \times F(s-1; n-1, p)$$

where $F(s-1; n-1, p)$ is the cumulative distribution function of the binomial distribution.

This means that we can rewrite Equation B.2 to the following:

$$\text{ESPRC}(s) = np \times [1 - F(s-1; n-1, p)] - s \times [1 - F(s; n, p)]$$

which we use in Subsection 3.3.3.

B.2 ESPRC Negative Binomial Distribution

For negative binomial distributed demand, we use the summation of the negative binomial probability mass function. Using Equation B.1, we get Equation B.3. However, we want to simplify this equation to make the ESPRC calculations easier.

$$\text{ESPRC}(s) = \sum_{s}^{\infty}(x-s)\binom{x+r-1}{x}(1-k)^x k^r \tag{B.3}$$

To make the ESPRC calculations easier, we transform the Equation B.3:

$$\begin{aligned}
\mathrm{ESPRC}(s) &= \sum_{s}^{\infty}(x-s)\binom{x+r-1}{x}(1-k)^x k^r \\
&= \sum_{s}^{\infty} x\binom{x+r-1}{x}(1-k)^x k^r - s\sum_{s}^{\infty}\binom{x+r-1}{x}(1-k)^x k^r \\
&= 1-\sum_{0}^{s} x\binom{x+r-1}{x}(1-k)^x k^r - s\left(1-\sum_{0}^{s}\binom{x+r-1}{x}(1-k)^x k^r\right)
\end{aligned}$$

The second summation is in the form of the cumulative distribution function of a binomial distribution $F(s;r,k)$. We also aim to transform the first sum:

$$\sum_{x=0}^{s} x\binom{x+r-1}{x}(1-k)^x k^r$$

to the form involving the cumulative distribution function of a binomial distribution F. Therefore, we will zoom in on the transformation of this sum.

Recall that the binomial coefficient is defined as:

$$\binom{x+r-1}{x} = \frac{(x+r-1)!}{x!(x+r-1-x)!} = \frac{(x+r-1)!}{x!(r-1)!}$$

Then, the sum can be rewritten as:

$$\begin{aligned}
\sum_{x=0}^{s} x\binom{x+r-1}{x}(1-k)^x k^r &= \sum_{x=0}^{s} x\frac{(x+r-1)!}{x!(r-1)!}(1-k)^x k^r \\
&= r\sum_{x=0}^{s}\frac{x}{r}\frac{(x+r-1)!}{x!(r-1)!}(1-k)^x k^r \\
&= r\sum_{x=0}^{s-1}\frac{x+1}{r}\frac{(x+r)!}{(x+1)!r!}(1-k)^{x+1} k^r \\
&= r\sum_{x=0}^{s-1}\frac{(x+r)!}{x!r!}(1-k)^{x+1} k^r \\
&= r\frac{1-k}{k}\sum_{x=0}^{s-1}\frac{(x+r)!}{x!r!}(1-k)^{x} k^{r+1} \\
&= r\frac{1-k}{k}\sum_{x=0}^{s-1}\binom{x+r}{x}(1-k)^{x} k^{r+1} \\
&= r\frac{1-k}{k}F(s-1;r+1,k)
\end{aligned}$$

Thus, we have shown that:

$$\sum_{x=0}^{s} x\binom{x+r-1}{x}(1-k)^x k^r = r\frac{1-k}{k}F(s-1;r+1,k)$$

where $F(s-1; r+1, k)$ is the cumulative distribution function of the negative binomial distribution.

This means that we can rewrite Equation B.3 to the following:

$$\mathrm{ESPRC}(s) = np \times [1 - F(s-1; r+1, k)] - s \times [1 - F(s; r, k)]$$

which we use in Subsection 3.3.3.