Tampere University

Jay Rajubhai Patel

A DEMAND FORECASTING FRAMEWORK FOR INTERMITTENT SPARE PARTS DEMAND

Master's Thesis
Faculty of Management and Business
Examiners: Professor Teemu Laine and
University Lecturer Jouni Lyly-Yrjänäinen
December 2025

ABSTRACT

Jay Rajubhai Patel: Developing a framework to forecast the spare part demand
Master of Science Thesis
Tampere University
Master´s Degree Program in Business and Technology
December 2025

In the after-market business, it is vital to minimise obsolete products and maximise product availability and sales. To make it possible, OEMs must know what the proper order quantity is when making procurement decisions. Hence, the demand forecasting plays a crucial role in after-sales business. However, generating an accurate forecast is challenging, especially when it comes to spare parts, where demand is often zero. Moreover, the most used forecasting methods in industry are outdated, and they cannot capture the accurate demand. Consequently, there is a clear need for an accurate and efficient method for forecasting spare part demand.

This thesis proposes a framework to generate the forecast and provide a comparison with theoretical demand to make informed decisions. The forecasting problem is formulated as a classification task in which the objective is to forecast the magnitude of quantity for spare part products. The framework begins by classifying spare parts based on demand interval and demand variability. The classification of spare parts gave four different categories of spare parts. The models are explored to find the best-performing model for the case company dataset. Each model is applied to each category of spare parts, and based on performance metrics, they are compared to find the best-performing model. Furthermore, in the framework, the theoretical demand is calculated based on machine population and maintenance records. This theoretical demand is compared with actual demand and forecasted demand to understand the aftermarket business and take informed decisions.

The results show that the majority of spare parts belong to the intermittent demand category, and the focus of the thesis is on predicting intermittent demand accurately. In the explored models based on literature, Teunter-Syntetos-Babai performed better in all demand categories compared to other models, which are prominent in industry, as well as models which are similar to Teunter-Syntetos-Babai. Even though the models explored are statistical, this study provides the framework to explore more advanced models with the same approach can give better results in other contexts as well. As a suggestion for future research, the models could be developed further, as well a machine learning models could be explored based on data availability. The new inventory policies and key performance indicators could be developed based on this research.

Keywords: spare parts, inventory management, procurement decisions, purchase decisions, demand forecasting, spare part demand forecasting, statistical models, moving average, exponential smoothing, Croston method, Syntetos-Boylan, Teunter, Syntetos & Babai.

USE OF AI IN THESIS

I have utilised AI tools in my thesis:

☐ No
☒ Yes

The AI tools utilised in my thesis and their purposes are described below:

In this thesis, ChatGPT-o3 and Microsoft 365 Co-pilot were utilised. The purpose of the utilisation of AI tools was to enhance the quality and structure of the text. The tools were used during the whole process and the thesis document, and specific chapters where the tools were used can't be named.

I acknowledge that I am fully responsible for the entire content of my thesis, including the parts generated by AI, and accept accountability for any violations of ethical standards in publications.

PREFACE

I would like to thank everyone who supported me during the process of working on my master's thesis. First, I'm very thankful for my supervisors, Teemu Laine and Jouni Lyly-Yrjänäinen, who both gave me valuable insights and advice during the process.

Second, I want to thank the case company for making it possible to solve such an interesting real-life problem. I am thankful to my manager for his guidance and time to explore the thesis directions.

Helsinki, 12th December 2025

Jay Patel

TABLE OF CONTENTS

1 Introduction

1.1 Background and Need for Research

Modern companies are focusing on providing spare parts and after-sales service to their customers, which generates highly profitable revenue. According to Van der Auweraer and Boute (2019), many modern industries have also shifted towards controlling their supply chain from the main equipment to their spare parts in the equipment lifecycle, which increases the need to predict what number of spare parts needed to run equipment without downtime and increase the productivity of the customer operations.

Spare part demand forecasting is crucial for modern industries, which is tied to the cost of inventory and the availability of critical parts needed for equipment. The spare part demand forecasting tackles the unique challenge that comes from erratic and highly variable demand consumption patterns, which is more complicated for commonly used traditional models to predict the demand in a modern competitive market (Pinçe et al., 2023). Effective demand forecasting has visible impacts on service levels, cost control and supply chain resilience, which makes it an important area for operations management research (Hu et al., 2018). The current research is focusing on data-driven approaches and bridging the gap between academic research and real-life applications.

In practice, many organisations still heavily rely on traditional time series models, which are easy to understand and implement on a large scale. However, these models give less accurate results on highly variable demand like spare parts demand, which leads to an increase in stockouts and excess inventory for some parts (Syntetos et al., 2005). According to Auweraer and Boute (2019), there are several excesses of data available with companies, such as installed base information, maintenance records and environmental conditions, which could be utilised in a modern framework to offer more precise predictions and inventory optimisation. Therefore, the research is needed to validate, compare and integrate advanced theoretical and empirical forecasting methods while developing a framework that accommodates the changing needs of the industries and technological advancements in the supply chain.

1.2 OBJECTIVE

The problem motivating the research lies in a forest machine manufacturing company. The company started an initiative to move towards more data-driven processes, where demand forecasting is one of them. The case company has a lower revenue percentage from its after-sales business in total sales compared to industry benchmarks and competitors. The case company wants to use data data-driven approach in demand forecasting to get an accurate demand to plan other processes. The case company does not have any current process for that. Currently, they rely on expert judgement from their extensive experience in industry.

The case company wants to develop a demand forecasting tool to use internally, which they want to use to determine when material should be ordered to make efficient replenishment. In practice, this means that reorder points are variable results of calculations based on safety stock values, known upcoming consumption and forecasted demand over the material lead time, and thus change as new forecasts are generated.

The forecasted demand by tool is shared with different stakeholders in the supply chain. By sharing material needs with vendors, the case company allows them to optimise their material flows by having better visibility into what the company will be ordering from them and when. In addition to being able to optimise material flows, the vendors can also utilise the information in their production planning, once again improving performance. As more and more vendors are given access to the forecast data, it becomes even more valuable to improve forecasting accuracy as much as possible.

Therefore, this research addresses the gap between the forecasting literature and its empirical application. The research would solve such problems in the company by answering the questions below:

Q.1) What are the best-performing methods to forecast demand for the case company’s spare part business?

Q.2) How can theoretical demand be calculated and included in the decision-making?

1.3 Research Methods

Research done in this thesis can be divided into three phases. The first phase consists of a literature review about general spare parts, spare part classification and forecasting methods, as well as the metrics used for measuring forecast accuracy. Information in this phase has been gathered from scientific, peer-reviewed articles and books which present best practices.

The second phase contains a case study in which multiple forecasts are generated with the statistical methods discussed in the literature review. After generating the forecasts, their accuracy is measured using different metrics to determine the most suitable forecasting method. Software solutions utilised in this study were Jupyter Notebook and Excel.

The demand data used for this study is the actual demand of spare parts recorded between 2019 to 2025. The data was collected from the enterprise resource planning system of the case company. The maintenance records documents were collected from the company database. The accuracy of each forecast is then compared to other forecasts in the same demand category. This phase also offers suggestions on which system should be used in the future to have as accurate forecasts as possible. Also, parameters which affect the outcome of forecasting models available are considered, and optimisation possibilities are identified.

Phase three consists of calculating the theoretical demand and comparing of theoretical demand, actual demand and forecasted demand. This phase provides the information related to the performance of after-sales strategies and dealer performance.

2 Spare Parts

2.1 Introduction to Spare Parts

Equipment or technical products require scheduled maintenance and repairs in case of failures. In most maintenance and repair scenarios, replacement parts are necessary to substitute defective components (Fortuin & Martin, 1999). Buker (2001) defines service parts, also known as spare parts or spares, as components used to maintain the operating condition of the products or equipment that a company sells or services. These parts may be stored at the production location as finished goods or distributed across service locations, dealers, or areas closely involved in the repair or maintenance of the end product. Vuorela (2012) adds that spare part inventories are strategically placed at locations close to customers to ensure timely availability.

Gopalakrishnan and Banerji (2004) define spare parts as components that are similar to those being replaced due to wear and tear during the equipment's operational life. Spare parts can be categorised into three types: first, they can include materials which are small products in big installations, such as nuts, bolts, and hoses. Second, spare parts could include equipment which is already assembled with many parts, such as engines, compressors, and alternators. Third, they may also consist of complete equipment, such as pumps and gears, which are directly installed into machines. While spare parts seem smaller and less expensive than main products, they play a crucial role in maintaining and enhancing the reliability of equipment. According to Teixeira et al. (2018), efficiency in spare parts management is important for maintaining operational efficiency of the equipment and minimising disruptions, which leads to better cost control and improved customer satisfaction.

Botter and Fortuin (2000) categorise spare parts into two main groups: repairable and consumables. Repairable components are those that can be replaced with other functional units and sent for restoration at a repair facility when they have breakdown, being both technically and economically viable for repair. However, consumables are parts that are replaced with new units and discarded upon failure, as they are neither technically nor economically suitable for repair. This classification approach is also discussed in other literature, albeit with slight terminological

variations. For instance, Fortuin and Martin (1999) use terms like recoverable items or recyclable parts for repairable, and disposables, throwaway parts, or expendables for consumables. Similarly, Driessen et al. (2010) use the same conceptual division but opt for the more straightforward terms repairable parts and non-repairable parts. The consistency of this classification principle across multiple studies underscores its significance in spare parts management, forming a foundation for developing effective inventory control strategies and maintenance planning in various industries.

Spare parts are integral to equipment maintenance, which can be categorised into different types based on timing and purpose. Several authors have given similar classifications with slight variations. For instance, Kennedy et al. (2002) and Cavalieri et al. (2008) distinguish the difference between preventive maintenance and corrective maintenance or unplanned maintenance. Preventive maintenance allows for more predictable spare parts demand, which often enables just-in-time ordering practices and more modern practices. In contrast, unplanned repairs necessitate safety stock policies to mitigate potentially significant stock-out costs, at the same time increasing the overstocking risk.

Expanding on this binary classification, Driessen et al. (2010) propose a tripartite categorisation: preventive, corrective, and modificative or improvement maintenance. Preventive maintenance aims to preempt failures and is typically scheduled within predetermined time frames when equipment is non-operational. Delays in such maintenance can increase the equipment downtime, reducing operational availability and increasing cost. Corrective maintenance addresses unexpected failures, requiring immediate action to minimise downtime. The third category, modificative or improvement maintenance, focuses on enhancing equipment performance and can be more flexibly scheduled.

Huiskonen (2001) emphasises that understanding these maintenance types is crucial for efficient and customer-centric spare parts management. Preventive maintenance provides more efficient inventory control due to its predictable nature of maintenance. Corrective maintenance demands efficient safety stock strategies to ensure a fast response to unforeseen failures. Improvement maintenance, being more discretionary, allows for strategic planning in spare parts procurement and allocation. This understanding of maintenance categories and their distinct characteristics is essential for developing spare parts management strategies that balance cost efficiency with operational readiness.

2.2 Spares Parts Business

The spare part business plays a critical part in the manufacturing and service industries. Naapila (2024) states that spare part business involves the manufacturing, logistics and after-market sales of parts, components and accessories, which are important for operating the machines, vehicles and equipment.

The spare parts business is often regarded as one of the most profitable functions within manufacturing and engineering-driven firms. Suomala et al. (2002) describe it as "the most profitable function of a corporation," owing to its significant contribution to both revenue and profit margins. Recent studies have revealed that after-sales services, including spare parts sales, contribute approximately 25% to the revenues and 40–50% to the profits of such firms (Dennis & Kambil, 2003; Cohen, 2005). Furthermore, Cohen (2005) states that the spare parts and after-sales services contribute around 8% of the annual gross domestic product (GDP) of the USA, showing its economic importance and how important it is in modern business and securing more profits in the spare and service business.

Various companies have changed their focus from optimising manufacturing processes to enhancing after-sales services and customer support. This transition highlights the growing importance of spare parts management as a profitable business area. Legnani and Cavalieri (2010) say that a well-practised service business ensures stable, long-lasting cash flows while fostering strong customer relationships and loyalty. According to Driessen et al. (2010), the availability of spare parts during machine breakdowns significantly minimises downtime, which increases customer satisfaction and also reduces costs related to lost revenues, possible associated claims, and even public safety hazards.

Despite the mentioned benefits about spare parts business, there are several challenges faced by companies in this business area. According to Boone et al. (2008), the following top challenges are faced by many companies:

- Absence of a comprehensive or systems-based perspective
- Not accurate spare part forecasts
- Lack of information system integration among stakeholders
- Variability in lead time
- Maintaining the product revision data
- Obsolescence of parts

- Planning for the service needs of older equipment and the maintenance of worn-out components
- Planning and execution for new product introduction
- Ensuring adherence to the repair cycle process
- Planning the location and physical distribution of spare parts

According to Boone et al. (2008), the absence of a comprehensive or systems-based perspective has been widely noted, as it is faced by many managers. This gap is intensified by a lack of collaborative relationships and information sharing across the stakeholders in the supply chain. The accuracy of spare parts forecast remains a critical challenge, unanimously identified by managers. An accurate forecast of spare parts is crucial to production planning and inventory control (Hasni et al., 2019). As Hua et al. (2007) mention, the commonly cited adage, “the forecast is always wrong,” is relevant in the context of spare parts due to their intermittent demand pattern and minimal historical data.

Lack of information among stakeholders is the most crucial challenge in the spare part business. It results in customer dissatisfaction when the lack of information about the part availability or the status of parts within the repair process is not shared with customers. Another challenge is variability in lead time, which results in the need for buffer inventory for desired customer satisfaction; however, these buffer stocks lead to high costs to maintain the availability and risks of the product going into an obsolete state (Christopher & Lee, 2004). In the spare part business, maintaining accurate configuration management and product revision data is a key problem. Patton & Steele (2003) mention that service providers are often unaware of the configuration until technicians are on site. To tackle this situation, Fortuin & Martin (1999) provide the solution that service providers have to keep additional spare parts in stock, which results in an overall cost.

One of the greatest risks in the spare parts business is the obsolescence of spare parts. Parts quickly become obsolete because of rapid advancement in technologies, changes in regulations or market conditions (Zheng et al., 2015). Even losing the service contract makes the inventory linked to those contracts redundant. Another Challenge mentioned by Boone et al. (2008) is planning for the service requirements of ageing products and the repair of ageing parts. Determining the last buy of service parts, keeping the suppliers and repairers for those parts and policies related to the repair and refurbishments make this challenge important.

The new product introduction results in the need for new service parts. It needs the planning of the service cycle and spare parts availability planning. Challenges lie in maintaining the service cycle and deciding the distribution of the spare parts. The modern spare parts business struggles to distribute the spare parts and manage maintenance work in every area. These are all the challenges faced by managers in the spare parts business, which makes it hard to get the targeted revenue from these sectors.

In recent years, there have been many technological advancements to help companies in their after-market service, such as predictive maintenance, blockchain technology, and additive manufacturing (Balakrishnan & Ramanathan, 2021). Additive manufacturing helps in reducing inventory holding costs by enabling on-demand production of slow-moving or obsolete parts, which also increases the supply chain responsiveness (Ahlsell et al., 2023). Similarly, data-driven predictive maintenance helps companies in anticipating parts failure, optimising spare parts inventory and reducing downtime, which results in customer satisfaction. According to Balakrishnan & Ramanathan (2021), the adoption of modern digital technologies in the supply chain improves resilience and performance.

2.3 Spare Parts Classification

Spare part classification is an important part to guide the whole management process. Classification is essential to know which parts are important and which should receive more attention (Cavalieri et al., 2008). According to Boylan & Syntetos (2008), spare parts for consumer products are highly varied with different costs, service requirements, and demand patterns. Therefore, companies must classify spare parts into different classes or types of demand patterns.

A classification of spare parts will help determine the service requirements for different parts classes and for forecasting and inventory management. Generally, this classification is based on operating efficiency considerations such as inventory cost, usage rate, and lead time derived from historical data (Teixeira et al., 2018). According to most literature, there are three approaches to classifying the spare parts: quantitative classification, qualitative classification, and demand-based classification.

2.3.1 Quantitative Classification

The first approach to the classification of spare parts inventory is quantitative techniques. The traditional method of ABC classification is the most used quantitative method, which is used to determine the service requirements based on the class of spare parts (Molenaers et al., 2012). According to Bacchetti & Saccani (2012), this method categorises parts based on the amount of demand volume and per-unit price for the item. This technique divides the inventory into three classes based on item demand and unit price criteria. Figure 1 shows the visual categorisation of inventory based on the Pareto law.

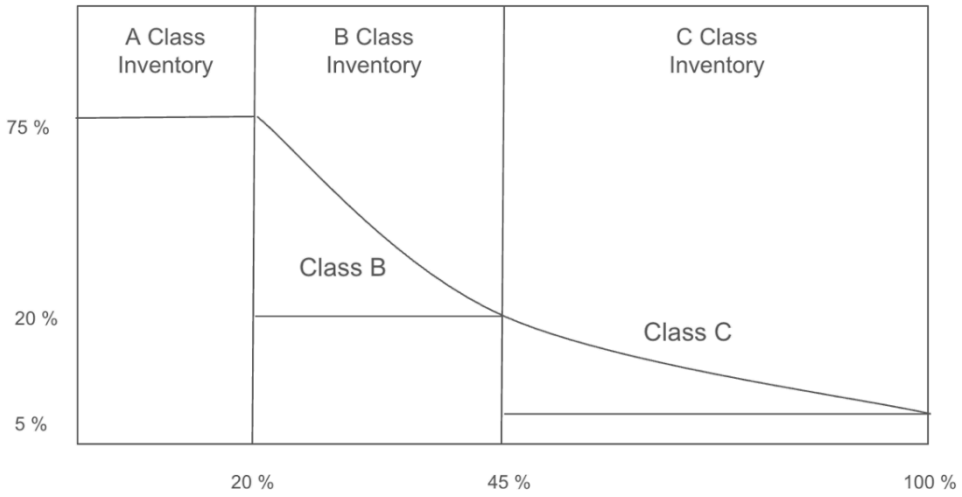

***Figure 1**. ABC Classification*

The above figure illustrates the ABC classification, which follows the Pareto law. Class A is the most important and valuable class based on Pareto law, with 60-80% of the total value, with 10-20% of the item inventory. Class C, with a value between 5% to 15%, has 50-60% of the inventory. Class B has 20% to 25% inventory with a value close to 30% (Hatefi et al., 2014). With such high variations, several authors agreed that each class should have a separate inventory policy. Class A needs the most attention whereas as a bit less control of class B items, and even less control of class C items (Oguji, 2013; Ramanathan, 2006).

However, Teunter et al. (2009) put forward the argument that since class C has the highest volume of the inventory, which means that the costs associated with handling back orders and urgent shipments for these parts can be significantly

higher than maintaining them in stock. Thus, they should also receive focus from management.

Syntetos et al. (2011) also contended that because class A represents the most critical segment, high-priced items have greater stocking costs due to elevated stock levels, ultimately leading to ineffectiveness in cost management. The ABC classification is a simple and convenient way to effectively manage inventory costs and service levels for inventories with relatively similar and consistent characteristics (Teixeira et al., 2017). However, according to Huiskonen (2001) ABC method relies solely on simple criteria of price and demand volume for classification, which fails to take different characterises of different inventory items.

Syntetos et al. (2011) also contended that because class A represents the most critical segment, high-priced items have greater stocking costs due to elevated stock levels, ultimately leading to cost ineffectiveness. The ABC classification is a simple and convenient way to effectively manage inventory costs and service levels for inventories with relatively similar and consistent characteristics (Teixeira et al., 2017). However, according to Huiskonen (2001) ABC method relies solely on simple criteria of price and demand volume for classification, which fails to take different characterises of different inventory items.

One more single criterion method, which takes the demand aspect while using the Pareto law, is the FSN method proposed by Larson (1980), where items are categorised into three parts: fast pace moving, slow pace moving, and non-moving. The parts having the highest demand (Top 25%) are classified as fast pace moving, and the lowest demand (bottom 25%) are classified as non-moving items. The remaining 50% items are classified as slow pace moving items (Hmida et al., 2014). According to Ghewari (2020), the FSN classification is preferable and used when the focus is on the moving rate of spare parts items and inventory turnover, by which it can be easily identified which products are more frequent and which products are non-moving or items that have become obsolete after a certain period.

2.3.2 Qualitative Classification

The second approach is a qualitative classification. According to Cavalieri et al. (2008), Qualitative methods try to assess the importance of keeping spare parts in stock based on different factors such as cost, downtime and usage of spare parts. Unlike quantitative methods, qualitative classification relies on expert judgement and contextual factors. These methods are based on the experienced people in the

company with expertise in managing the inventory (Roda et al., 2014). This method is particularly important for spare part management, where unpredictable demand patterns with a high risk of obsolescence are there.

The qualitative classification prioritises spare parts based on their operational impacts. The most common method in qualitative research is VED analysis, where items are classified in Vital, Essential, or Desirable based on the perception of the spare part importance (Mukhopadhyay et al. 2003). According to (Teixeira et al., 2017), VED classify parts into:

- Vital(V): spare parts that have a high critical impact on the production process.
- Essential(E): Less impact on the production process and moderate cost loss.
- Desirable (D): Less disruption in the production process and less risk of loss.

Like the ABC analysis, the VED classification is easy; however, the major drawback is that it depends on users understanding of parts, which brings biases into categorisation. Therefore, Gajpal et al. (1994) introduced the use of the Analytical Hierarchy Process (AHP) to overcome the problem of judgment by assisting in the criteria selection to limit the problem of subjective judgment. When the selection of multiple criteria comes, AHP is also considered an effective method to handle the situation (Teixeira et al., 2017). The method helps the decision makers to decide which alternative gives the optimal solution. It is designed as a hierarchy diagram with a goal on top and different levels, such as criteria, sub-criteria and alternatives, to decide the optimal solutions (Nurcahyo et al., 2017).

Other methods are hybrid frameworks which combine the quantitative and qualitative factors to address the limitation. Cavalieri et al. (2008) provide an approach to combine the criticality, lead time and demand patterns. While Roda et al. (2014) talk about hybrid models to balance the cost and operational costs. Ramanathan (2006) introduced ABC classification based on multiple criteria using weighted linear optimisation. In the model, all criteria are converted into a scalar score using weighted sums of measures under each criterion. And then, according to scores, items were grouped into different classes.

Bayraktar et al. (2008) proposed a criticality-based ABC classification, which has multiple criteria where criticality is the main criterion. Figure 2 shows criticality

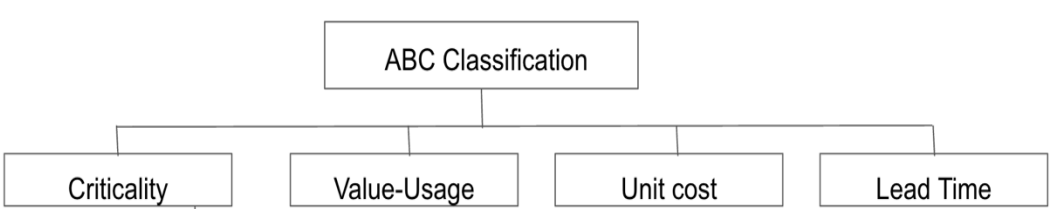

analysis and three sub-criteria of criticality, namely substitutability, penalty costs and commonality.

Figure 2. *ABC classification with Criticality as the main criterion (Bayraktar et al.,2008)*

As seen in the above figure, Bayraktar et al. (2008) include the value-usage, unit cost, lead time and criteria for classification. Based on criticality analysis, parts were categorised into three categories: very important, important and less important. This classification helps focus on the parts that are very important, based on their criticality to the equipment or the business overall. This method facilitates criticality-based maintenance planning, enhancing customer satisfaction and profitability. Where the penalty cost, substitutability and commonality are represented as weights in the formula to classify parts. When the model is solved repeatedly for each item, the value of the objective functions gives the optimal criticality scores, which are then used in the classification of the inventory items.

2.3.3 Demand Pattern-Based Classification

The third method of classification is based on the demand trend of the parts, which considers the demand variability, erraticness, and intermittence. Williams (1984) study put the demand into three categories, namely sporadic, slow-moving and smooth, based on the degree of lumpiness and intermittence. Huynh (2019) states that the degree of lumpiness is calculated by the squared cp-efficient of variation as a proportion of the mean inter-demand interval and mean lead time, and the degree of intermittence is the amount of lead time between two non-zero demands.

According to Williams (1984), the items with a high lumpiness level of more than 0.5 and a high level of intermittence are classified as sporadic demand items. Slow-moving demand items are those with a low level of intermittence, although the level of lumpiness might vary. And a high level of intermittence and a low level of lumpiness items are grouped into smooth demand items.

Eaves & Kingsman (2004) argue that many frameworks do not include lead time variability into account while categorising the items. They suggested considering the lead time variability to divide the intermittent demand into highly intermittent and mildly intermittent demand. Eaves and Kingsman (2004) studies lead time in the lead time demand, which is made up of three variables, namely transaction variability represented by squared coefficient of variation for demand size proportion

of lead time mean, demand size variability represented by squared coefficient of variation as a proportion of average number of transactions and lead time variability represented by squared coefficient of variation of lead time.

Eaves & Kingsman (2004) classify the demand into five categories being smooth, slow-moving, irregular, mildly intermittent and highly intermittent. Smooth demand parts are identified by low transaction variability and low demand size variability. The items which have low transaction variability and high demand size variability are termed irregular items. Whereas high transaction variability and low demand size variability define the slow-moving products. The difference between Williams (1984) and Eaves & Kingsman (2004) categorisation is the last demand, where Eaves & Kingsman (2004) divides the intermittent demand into mildly intermittent and highly intermittent, which have high transaction variability as well as high demand size variability, but low-time variability items are termed as mildly intermittent where as high lead time variability items are termed as highly intermittent. Figure 3 shows the classification proposed by Eaves and Kingsman (2004).

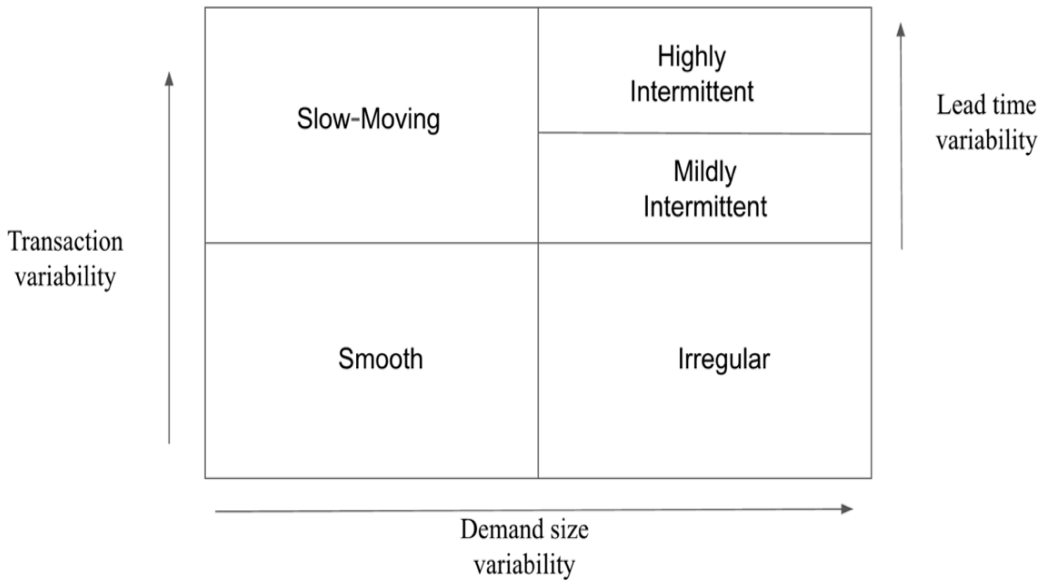

Figure 3. *Demand Classification by Eaves and Kingsman (2004)*

The above figure illustrates the classification of demand by Eaves and Kingsman (2004). The method gives management the freedom to set cut-off values for all three variables. Depending on the industry, management can decide on different values between category boundaries.

Further modification was made by Syntetos et al. (2005), where they classified items based on their level of erraticness and intermittence. Where the squared coefficient

of variation is the erraticness of items, and the average of inter-demand intervals is the intermittence. According to Syntetos et al. (2005), the items which have a high level of erraticness and a low level of intermittence are classified as erratic demand, where gaps between consecutive periods are low, and variation in demand size is high. Similarly, if the items have a high level of intermittence, then it is defined as lumpy demand. On the other hand, when the average gap between consecutive periods is low, it is possible that the level of intermittence is different, with a smaller number of zero-period items called smooth items and a larger number of zero-demand periods, called intermittent items (Syntetos et al., 2005). Figure 4 shows the classicisation with cut-off values based on erraticness and intermittence.

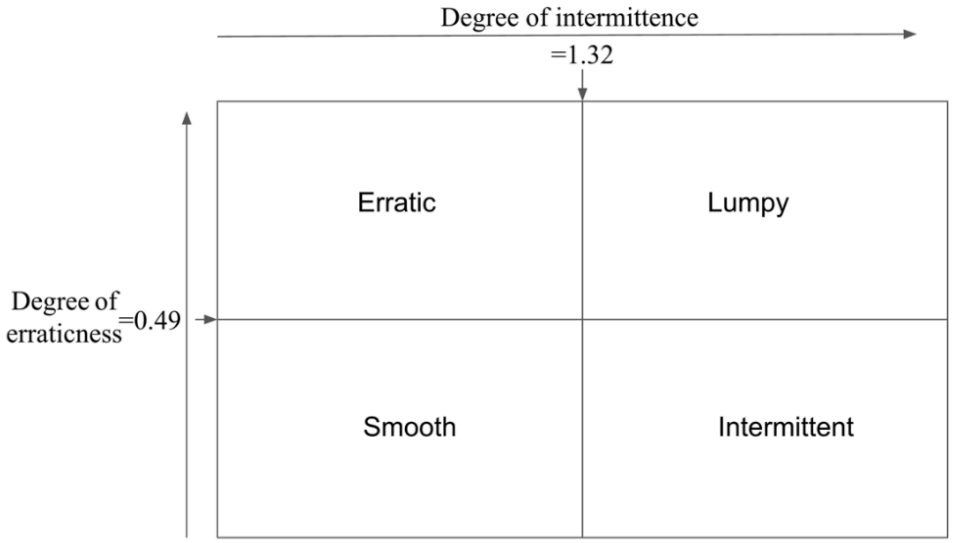

Figure 4. *Demand Classification by Syntetos, Boylan and Croston (2005)*

The above figure shows that a level of erraticness less than 0.49 defines the two demand patterns, where the degree of intermittence is less than 1.32, indicating smoothness, and more than 1.32 is defined as intermittent. And those items that have a degree of erraticness more than 0.49 and a degree of intermittence less than 1.32 are classified as erratic, and more than 1.32 are classified as lumpy demand. The classification of spare parts method provided by Syntetos, Boylan and Croston (2005) is practical and used in modern industry to manage each demand pattern differently. The classification of parts helps in choosing forecasting methods for each demand pattern to achieve accurate forecasting, as well as it can also help managers to prioritise and allocate resources to manage certain parts.

3 Inventory Management and Demand Forecasting

3.1 Inventory Management

Inventory management is an important function of any company, particularly in meeting the customer demand (Baron et al.,2010), as it helps to ensure a balance between the supply of the parts and the demand for the parts. The fundamental goal is to turn inventory as fast as possible without losing customers (Gitmann and Zutter, 2012). Effective inventory management helps in reducing costs related to inventory holding, stock-out costs, as well as procurement costs by buying in the optimal quantity, keeping price movement in consideration (Basu and Wang, 2011).

However, according to Koumanakos (2008), inventories can be an asset or liability from an operations management perspective. Holding excess inventory consumes the physical space, creates a financial burden on the organisation and increases the risk of demand, defects and loss. Moreover, the risk of obsolescence of the parts also increases.

Spare parts availability is critical for maintaining the business operations. Therefore, the management of spare part inventory becomes an essential part of inventory management (Muller, 2011). According to Hu et al. (2018), the non-availability of spare parts on time can cause significant financial losses and dissatisfaction among customers, and thus, inventory management for spare parts plays a critical role in ensuring the desired parts availability.

According to Celebi et al. (2008), there are several differences between traditional inventory management and spare parts inventory management. While traditional inventory management focuses on managing the finished goods whereas spare parts inventory management focuses on managing the spare parts, which are used in different maintenance processes, such as repairing or replacing faulty parts in equipment. It requires a high level of accuracy and precision as the availability of spare parts is critical to maintain the service business, and the financial effects of a stockout can be multiple times the price of spare parts.

There are several key differences between the two approaches, which are demonstrated in Table 1.

Table 1. *Difference between Traditional and Spare Part Inventory Management (Celebi et al., 2008)*

	Traditional Inventory Management	**Spare part Inventory Management**
Focus	Mainly Finished Goods	Maintenance Parts
Demand	Smooth/Low variability	Lumpy/Intermittent
Demand Forecasting	Easier duo to smooth demand and low variability	Requires more accurate models' duo to high variable demand
Lead Times	Shorter duo to easier manufacturing	Usually Longer due to special set-ups in manufacturing.
Characteristics	Most Items are same kind	Most items are distinct
Stockout Cost	Less Critical	High Impact. Loss of production
Order Policies	Less Orders but high quantity	More orders and less quantity

As shown in Table 1 that both approaches have several differences, such as traditional inventory focuses on the finished goods, whereas spare part inventory focuses on goods which as used for the maintenance of machines or equipment. Secondly, traditional inventory has demand which is smooth and has low variability among orders; however, spare parts demand is usually intermittent and lumpy,

which is hard to forecast as compared to traditional inventory with smooth demand. One more important factor is lead time, where traditional inventory has a shorter lead time as compared to spare parts inventory, where spare parts inventory takes time to set up the manufacturing for small quantity orders. The impact of stockout on business is less critical with traditional inventory than with spare parts inventory, where a critical spare parts breakdown can stop production. Moreover, traditional inventory is easier to manage compared to spare part inventory because it has more quantities and less order frequency, whereas spare part inventory has more order frequency with fewer quantities.

3.2 Demand Planning and Forecasting

The common terms of demand refer to the amount of quantity needed to fulfil the customer's consumption for a specific time period. This phase of the sales and operations process focuses on meeting customer demand requirements with the use of demand forecasts and the risk associated with demand forecasts (Sodhi and Tyagi, 2011). The baseline forecasts are defined by the past sales, and it is provide the forecast quantity which is not constrained by the production capacity of the organisation (Grimson & Pyke,2007). This demand forecast used by the organisation to make supply planning, procurement, purchase and production planning accordingly; however, demand uncertainty makes this process difficult. The overstocking of the finished goods leads to an increase in inventory costs, and not meeting customer demand leads to lost sales (Sodhi and Tyagi, 2011). Therefore, according to Vereecke et al. (2018), demand forecasting is a crucial process in the whole sales and operations planning process, as the planning for the resources to customer satisfaction depends on it.

Demand planning and forecasting is the process of using data data-driven approach for making informed decisions to take infomed decisions in today’s competitive market, helping organisations to provide products to customers on time while minimising operational costs. According to Mentzer & Moon (2005), organisations typically use two distinct types of demand forecasting. First, organisations forecast total sales through diverse methods available in management's market perspective and organisational targets, typically have longer time horizons and focuses on strategic business objectives. Second, organisations require short-term demand predictions for specific periods, such as the upcoming week or month. These shorter-term forecasts are developed using historical data and quantitative approaches, referred to as statistical or unbiased forecasting.

The fundamental difference between these forecast types lies in their foundational basis and organisational purpose. Strategic forecasts come from management judgment and organisational ambitions, while operational forecasts come from historical patterns and empirical data. Stevenson (2021) states that accurate forecasts at both levels are crucial for capacity planning, inventory management and overall organisational planning. An unbiased forecast provides the projection of future outcomes by extrapolating past data patterns; however, determining its accuracy requires careful consideration of taking relevant historical data, which reflects the conditions likely to exist in the future. Both forecast types remain important for effective organisational planning, despite serving fundamentally different strategic purposes over different time horizons.

The demand forecasting process involves risk management as an essiantial part of the process is to predict the uncertain future. According to Armstrong & Green (2017), using experimental evidence, the systematic forecasting approach could reduce the prediction error by half compared to commonly used traditional methods. However, organisations cannot solve all the demand planning challenges by simply using new forecasting approaches. The demand forecasting methods provide the tool to enhance organisational perspective and inform management decisions. Armstrong (2001) emphasises that forecasting should only be undertaken when genuine uncertainty about future outcomes exists; forecasting adds no value when results are already predetermined or fully certain. Thus, organisations should employ forecasting strategically as one component within a broader planning framework that integrates expert judgment, operational constraints, and market intelligence.

Before investing in a forecasting organisation should consider the expected return of the investment. Forecasts provide information for decision-making, and the returns from them are therefore indirect (Stevenson 2014). Since the return on investment in the case of forecasting comes indirectly from, for example, reduced safety stocks organisation needs to be ready to follow the plan based on the improved forecast to get a return (Feigin 2014). These expected indirect returns should be identified at an early stage of investment planning since management needs to get benefits from the investment.

3.3 Moving Average

The moving average method (MA) is a straightforward approach used for demand forecasting, where future demand is predicted by averaging the demand observed in a certain number of preceding periods. By doing this, the moving average method smooths out fluctuations in the data, providing a more stable estimate for future demand. The forecast for the next period is essentially the average of actual demand over the previous number of periods.

$$P_{t+1} = \frac{D_t + D_{t-1} + D_{t-2} + \ldots + D_{t-N+1}}{N}$$

In which: P_{t+1} = forecasted demand for period t+1

D_t = Demand for period t

N = Number of periods

The method is applicable when there is no significant variation or seasonality in past demands. Consequently, a deviation is likely to occur between the forecasted and actual values when a trend emerges during the selected period. However, a major disadvantage lies in determining the periods to utilise in the forecast. A longer period tends to diminish the variation, whereas a shorter period tends to illustrate demand changes in recent times. A key limitation of the moving average method is that it gives equal importance to all the most recent observations within the chosen time frame, disregarding any prior data. This approach assumes that all past values are equally relevant, but it can be argued that more recent data points may offer greater insights into future trends compared to older observations (Wheelwright & Makridakis, 1973). Nonetheless, moving averages do offer advantages, such as ease of use, speed, and relative cost-effectiveness. Overall, this method yields satisfactory results when there is no trend or seasonality, and the forecasting is intended for a short period (Russell and Taylor, 2011). Moving average is still one of the most used methods in traditional businesses.

3.4 Exponential Smoothing

Exponential smoothing (ES) was developed to overcome the limitation of the moving average method, which is its inability to adequately reflect recent data changes. Due to its minimal historical data requirements, Exponential smoothing is widely adopted

in management practices and extensively used by managers to forecast accurate demand and make planning reliable with the forecasted data (Russell & Taylor, 2011). Exponential smoothing-based demand is calculated by a weighted average of forecasted and actual demand from previous periods, applying adjustment parameters based on the assigned weights to recent demand hikes. The formula for exponential smoothing is as follows.

$$P_{t+1} = P_t + \propto\left(D_t - P_t\right)$$

In which: P_{t+1} = forecast demand for period t+1

D_t = demand for period t

$\propto$ = weighting parameter smoothing constant, $0 < \propto < 1$

The smoothing constant (α) represents the impact of recent demand on the forecast, and it is typically calculated based on expert judgment and the specific characteristics of the production process (Russell & Taylor, 2011). Lower α values are suitable for steady-growth products, while higher values respond quickly to demand variations (Hazriani et al., 2022). Optimal α values can be determined through trial-and-error analysis to minimise forecast errors, such as mean absolute Deviation (MAD) or Mean Squared Error (MSE) (Paul, 2011; Sari & Winarno, 2023). The method used for adopting the higher value in demand variation, which means that keeping the α values higher would be perfect for minimising forecast error.

3.5 Croston Method

The exponential smoothing method is good when the recent data has a demand history with a significant quantity; however, according to Croston (1972), a key challenge for the Exponential smoothing method is to forecast the demand when the recent data has many zero demand and the method is not able to generate reliable forecast because the sudden spike in demand cannot be predicted by exponential smoothing when the weighted parameter is low because of recent data shows the zero demand. As a result, it is common to differentiate the demand size and the interval of demand occurrence. Subsequently, the average demand per period is calculated by dividing the average demand size by the mean transaction interval. Originally introduced by Croston (1972) and later refined by Rao (1973), the Croston method focuses on forecasting the mean demand per period by applying

exponential smoothing separately to demand intervals and demand sizes. As noted by Eaves and Kingsman (2004), the Croston method is similar to traditional exponential smoothing when demand is smooth with less consecutive zero demand. The key difference is that, in the Croston method, the period of forecasting is performed only when demand is positive. If demand is zero, the method increases the period count until demand becomes non-zero.

Initially, the Croston method defines the average interval by calculating the average interval between non-zero demand and the size of non-zero. Both are respectively denoted as demand periods (denoted P) and the average size of demand (denoted F), specifically,

When $D_t = 0$,

$$P_{t+1} = P_t$$

$$I_{t+1} = I_t$$

$$Q = Q + 1$$

Else,

$$P_{t+1} = P_t + \alpha(D_t - P_t)$$

$$I_{t+1} = I_t + \alpha(Q - P_t)$$

$$Q = 1$$

Subsequently, the forecasted average demand per period is calculated by dividing the average demand size of non-zero demand by the average interval of occurrence of non-zero demand that is,

$$C_{t+1} = \left(\frac{F_{t+1}}{P_{t+1}} \right)$$

In which:

D_t = demand for period t

F_{t+1} = forecast demand for period t+1

P_{t+1} = forecast interval between transactions for period t+1

Q = interval from the latest period with positive demand

C_{t+1} = forecast mean demand for period t+1

α = smoothing parameter for forecast intervals, 0 < α < 1

In a 1994 study, Willemain et al. stated that the Croston method has an advantage over exponential smoothing as it can handle the zero demand with many periods by separating the demand size of non-zero demand and demand interval; however, the method did not perform well in real-world case studies as it did in theory. The need for a method was still there for managers to forecast demand with many zero demands. In addition, Syntetos and Boylan (2001, 2005) showed that both exponential smoothing (ES) and the Croston method have biased demand forecasts, with exponential smoothing being particularly affected. Their research revealed that, even though the individual components mean demand size and mean demand interval are calculated accurately, the combined forecast of mean demand per period is often off. Specifically, the demand forecast using the Croston method tends to be either too high or too low compared to the actual demand, with the error growing larger as the smoothing parameter α approaches 1 (Syntetos & Boylan, 2001; Wallström & Segerstedt, 2010).

3.6 Syntetos-Boylan Approximation

The drawback of the Croston method is the high error when the smoothing parameter approaches 1. Syntetos and Boylan (2011) modified the Croston method by adding the new deflating factor. Syntetos and Boylan (2011) argue that the Croston method is recommended only for a low value of α. The bias becomes more significant when the α value goes above 0.15. Therefore, it is necessary to modify the CR method to ensure that the demand estimates per time period are unbiased.

The modified method uses the forecast demand and forecast interval between transactions as the original method, a modification to the Croston method is

$$C_{t+1} = (1 - \alpha) \frac{P_{t+1}}{I_{t+1}}$$

In which: P_{t+1} = forecast demand for period t+1

I_{t+1} = forecast interval between transactions for period t+1

C_{t+1} = forecast mean demand for period t+1

α = smoothing parameter for forecast intervals, 0 < α < 1

This adjustment ensures unbiased forecasts even with higher smoothing constants. The Syntetos-Boylan Approximation has been widely adopted as an improvement over the Croston method for intermittent-demand scenarios.

3.7 Teunter, Syntetos, and Babai

Intermittent demand forecasting became more challenging for the items which are at risk of being obsolete. Babai et al. (2014) put forward the argument that the Croston method and the modified Syntetos-Boylan Approximation method only update the demand size and demand interval when the positive demand occurs. To solve this problem, Teunter & Duncan (2009) developed a method called Teunter-Syntetos-Babai (TSB), which differs from previous methods in focusing on the demand interval to predict demand size, whereas the Teunter-Syntetos-Babai method utilises demand probability to estimate demand size. This estimation comes true, and this method gave better results than previous methods. Syntetos et al. (2015) stated that the difference between the Syntetos-Boylan Approximation and Teunter-Syntetos-Babai is that Teunter-Syntetos-Babai provide unbiased results, resulting in less forecast error.

The formula for Teunter-Syntetos-Babai is as described below,

If $Y_t = 0$ then $P'_t = P'_{t-1} + b(0 - P'_{t-1}), Z'_t = Z'_{t-1}, Y'_t = P'_t Z'_t$

If $Y_t > 0$ then $P'_t = P'_{t-1} + b(1 - P'_{t-1}),, Z'_t = Z'_{t-1} + a(, Z'_t - Z'_{t-1}), Y'_t = P'_t Z'_t$

Where the below,

Y_t = demand for an item in period t

Y'_t = estimate of mean demand per period made in period t for period t +1

Z_t=demand size in period t

$Z_t^{'}$=estimate of mean demand size in period t

$P_t^{'}$=estimate of the probability of demand occurrence in period t

a,b= smoothing parameter

This thesis will focus on Croston and variation methods, as most of the methods are designed to handle many zero demands, which the spare part demand usually has.

3.8 Forecasting Accuracy Measurement

The exploration of various forecasting models needs to be quantified to select the best-performing model for a particular dataset. Many measures of forecast accuracy have been discussed in the literature in the past, and several authors have discussed what methods should be used while comparing the different forecasting models which have univariate time series data (Hyndman & Koehler, 2006). Hyndman & Koehler (2006) provided the categorisation of these methods into four parts as below

1. Scale-dependent measures
2. Measures based on percentage error
3. Measures based on relative error
4. Relative measures

Firstly, there are several measures whose scale depends on the scale of the data. They are useful when comparing different methods applied to the same set of data. The forecast error e used in scale in scale-dependent measure is calculated by subtracting the actual demand quantity from the forecasted demand quantity (Vaskinen,2017).

$et = Y\ t - Ft$

The focus of the forecast has a long horizon, so a single value of the forecast error is not feasible; however, multiple periods have forecast errors, and averaging those gives feasible error results. The most used formula to calculate the forecast error based on scale is listed below.

$$Mean\ Absolute\ Error\ (MAE) = mean(|et|)$$

$$Geometric\ Mean\ Absolute\ Error\ (GMAE) = gmean(|et|)$$

$$Mean\ Square\ Error\ (MSE) = mean(et^2)$$

Syntetos & Boylan (2005) explain that Mean absolute error, Geometric mean absolute error and Mean square error are only useful when comparing a single series because of their scale dependency; they cannot be used in multi-series. If one series has scale in thousands and another in millions, the error will not truly show which one has better results.

Secondly, the percentage-based measure solves the problem faced in scale-based measures by becoming scale independent, which allows comparison of multiple data series to be compared. It is also easier to understand without needing extra information, such as if the error is 20%, it means that the forecast is off by 20% average for all parts. The most used percentage-based metric is mean absolute percentage error, which is defined by below formulas.

$$Percentage\ error\ (pt) = 100\ *et\ /\ Y \quad (14)$$

$$Mean\ absolute\ percentage\ error\ (MAPE) = mean(|pt|)$$

Percentage-based error metrics have some important challenges. They cannot be applied to time series that have many zero-demand periods because, as per the formula, it is required to divide by zero. Additionally, when actual demand values are very small, these metrics can distort the assessment of forecast accuracy, giving a misleadingly exaggerated error.

Thirdly, one of the alternatives to scale-independent methods is the relative error metric, which compares the forecast error obtained by the tested method to some benchmark method. After absolute errors have been calculated for both methods, relative errors can be calculated by below formula, where et* is the forecast error of the benchmarked method.

$$Relative\ error\ (rt) = et\ /\ et^*$$

The most used benchmark method is the naïve method, where the forecast is taken as the last actual demand forecast. If the last year's demand for a certain part is 90, the benchmark model will take 90 s this year's forecast demand and compare the results with the tested method. However, Hyndman and Koehler (2006) add that

when using different demand series, relative errors cannot be used. When errors are very small, as is commonly in the case of Intermittent demand and when using a naïve forecast as the benchmark method, the relative error becomes impossible because it would again require dividing by zero.

Therefore, mean absolute scaled error (MASE) has been introduced as an option which can be used universally to measure forecast accuracy for all demand types. If the scaled error is less than one, the forecast being measured is more accurate than the naïve forecast generated from previous demand and greater than one if the forecast is worse than said naïve forecast. As with other measurements error of one period is not that useful and the mean of multiple periods is better for performance comparison of different forecast methods. Mean absolute scaled error can be calculated with below formula (Hyndman & Koehler, 2006)

$$Mean\ absolute\ scaled\ error\ (MASE) = mean(|q_t|)$$

Overall, the forecast error literature provides many methods which are useful in different cases. This thesis will utilise the scale-based measure and percentage-based measures as they are easier to explain to non-technical stakeholders.

4 Framework

4.1 Introduction of the framework

This chapter presents the comprehensive analytical framework for demand forecasting in the spare parts operation of the spare parts of company. The framework consists of five interconnected processes, as shown in Figure 5.

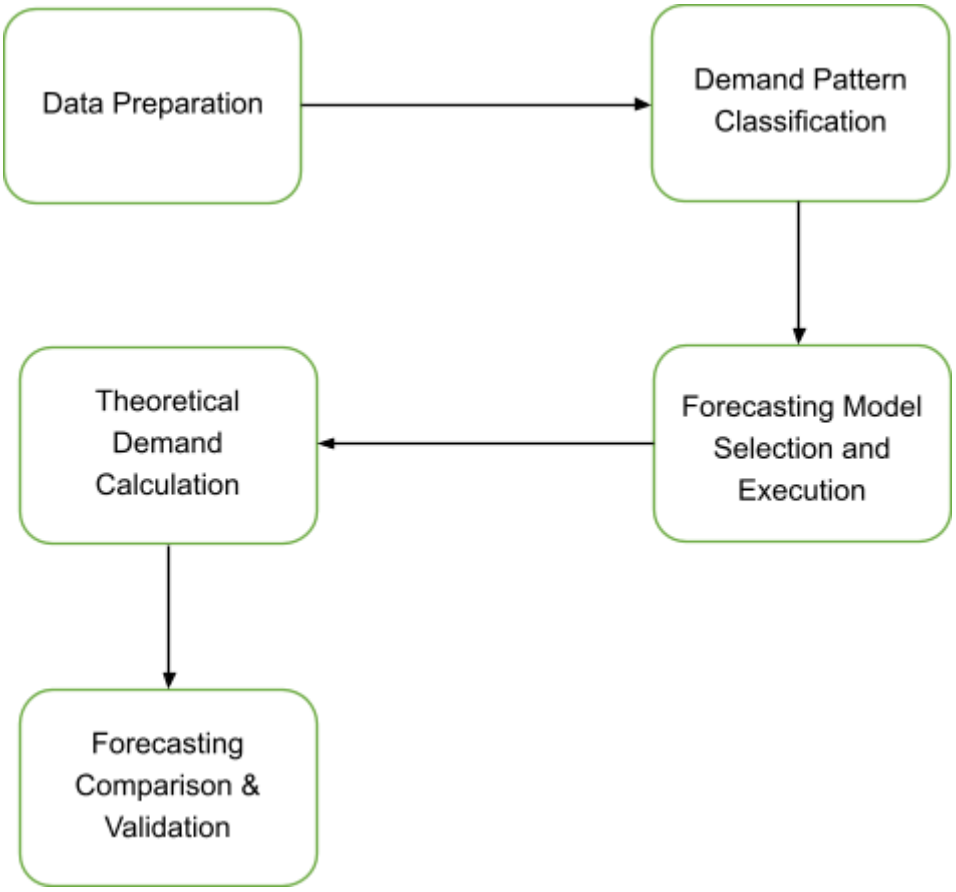

***Figure 5**. Integrated Forecasting Framework - Five-Step Workflow*

As shown in the Figure above, these five processes offer a systematic framework for finding an accurate demand forecast while also allowing comparison with theoretical demand to validate the demand expectation.

The framework is designed to address the practical problems of forecasting thousands of spare parts items with different demand characteristics. It integrates the data-driven classification of spare parts, comparison-based forecast method

selection across different demand patterns, and theoretical validation to take informed inventory decisions and dealer management.

The development of this framework recognises gaps in practical forecasting implementation within spare parts management contexts. While academic literature establishes theoretical principles for demand pattern classification and method selection, limited literature exists for operationalising these rules across large, diverse spare parts items in industrial settings. This framework bridges that gap by providing a structured, systematic, and scalable methodology that maintains academic rules while addressing practical challenges, including data availability, computational capabilities, and organisational capabilities.

As a similar methodology discussed by Syntetos et al. (2005), the framework takes a two-dimensional classification system that classifies the spare parts based on their demand intermittence and variability characteristics. This approach enables management of different demand patterns based on different forecasting methods and inventory decisions across all the items.

4.2 Data Preparation and Quality Assurance

Before classification, the historical data is gone into preparation and filtered to find which parts are suitable for quantitative forecasting and practically possible to apply different methods. Historical data has been extracted from enterprise resource planning systems for the past few years. The timeframe provides sufficient observations to identify different demand patterns. The raw transactional data includes multiple fields which describe the factors affecting the demand. The raw transactional data has many null values, which are a challenge during the analysis. A critical part of data preparation is to handle null values. It is key to identify what a null value is and what a zero demand. Missing data suggests a lack of information which needs to be excluded from analysis. The handling of missing values also depended on what the column represents and does it impact the analysis.

Handling the product code and items is important to ensure. Many product codes in the data represent different items. The items which are not spare parts need to be excluded from the analysis. Identifying the items which are not spare parts needs to be communicated to the case company.

Before the analysis, the data undergoes data validation steps which rule out quality issues which might affect the classification, as well as be irrelevant to the forecast.

The validation ensures that the data quality is enough to classify the item and apply the forecasting model to find the pattern. There are two criteria to ensure that the analysis focuses on relevant products.

1. **Minimum Transaction Frequency Criterion:** This criterion ensures that each item has a minimum amount of transaction frequency, which is needed to apply the statistical forecasting model. With fewer observations, forecasting accuracy becomes highly uncertain and unreliable. The coefficient of variation calculation shows high sampling error with small samples, potentially resulting in unstable classifications that change drastically with each new observation.

2. **Recent Activity Criterion:** This criterion ensures that items which are included in the analysis have a recent sales history. Parts that sold historically but show recent sales may represent obsolete products which need focus if they represent high volume in inventory. Including such items leads to unnecessary data points and misleading results, which affects the forecasting accuracy. Parts failing this Criterion are excluded from analysis.

After applying the filtering criterion, the remaining parts can be considered as active. This framework makes sure that focus will be on items which has recent sales and enough history to classify and forecast the demand. Overall, this phase makes sure that the data for analysis has high quality and focuses is on what is a relevant spare part for the company and has the capacity to generate good results. This phase excludes many spare parts as the majority do not have enough data, which makes it challenging to forecast the spare parts.

4.3 Demand Pattern Classification

For the items which are meeting active criteria, the framework calculates two quantitative measures that characterise fundamental aspects of their demand patterns. These measures operationalise the two-dimensional conceptual framework, providing values that enable classification.

1. **Average Demand Interval (ADI):** The average demand interval quantifies the total number of periods to the number of periods with non-zero demand. The higher the value of ADI, the harder it is to predict the demand.

2. **Squared Coefficient of Variation (CV^2):** The squared coefficient of variation quantifies demand magnitude variability relative to the average demand level. A higher CV^2 value indicates higher difficulty in predicting demand magnitude. A lower CV^2 value demand pattern is easy to predict with simple models.

The calculated ADI and CV^2 values position each spare part within a two-dimensional space defined by these measures. The classification framework puts this into four regions using established threshold values: ADI = 1.32 and CV^2 = 0.49. These threshold values, derived from research by Syntetos, Boylan and Croston (2005), represent points where forecasting method performance characteristics change. They are classified into four categories, namely as below.

1. Smooth Demand (ADI < 1.32, CV^2 < 0.49)
2. Erratic Demand (ADI < 1.32, $CV^2 \geq 0.49$)
3. Intermittent Demand (ADI ≥ 1.32, CV^2 < 0.49)
4. Lumpy Demand (ADI ≥ 1.32, $CV^2 \geq 0.49$)

All the parts are classified into these four demand patterns, and then forecasting methods are implemented for each demand pattern. According to Syntetos, Boylan and Croston (2005), the majority of the spare parts fall into intermittent demand, which is harder to deal with as it has many periods with zero demand, and then it has demand. The framework adopts this classification as it is easier to adopt in a practical case, and it has no bias as it is based on mathematical classification compared to other methods, which are based on expert judgement with their own bias in classification.

4.4 Forecasting Method Selection and Implementation Framework

4.4.1 Selection Of Forecasting Models

The forecasting method selection framework is built on a demand pattern classification framework to assign the right forecasting approach to different categories. The aim is to apply selected statistical methods to different categories and measure their forecast accuracy to select which one performs better in which categories. This framework ensures that the chosen forecasting approach is not only theoretically consistent but also empirically validated using real data.

After the completion of the two phases, the next phase focuses on exploring the different forecasting models discussed in the literature. There are many forecasting models in the literature, from basic statistical models to advanced machine learning models, which include many factors in the model. However, this framework focuses on a basic statistical model, which only uses historical demand data. Given the various demand categories, the framework employs multiple forecasting methods, which are discussed in the literature. These include:

1. **Moving Average**: A simple but effective technique suitable for smooth and moderately variable demand.
2. **Exponential Smoothing (ES)**: A responsive model that adjusts to demand level changes through weighted averages. It gives more weight to recent observations.
3. **Croston Method:** Specifically developed for intermittent demand, separating the estimation of demand size and demand interval.
4. **Syntetos–Boylan Approximation (SBA):** A bias-corrected version of the Croston method, improving accuracy for highly lumpy patterns.
5. **Teunter-Syntetos-Babai (TSB) Method:** An intermittent-demand method that models the probability of demand occurrence and the demand size separately using exponential smoothing for both components.

There are five statistical models included in the analysis. These models provide a balanced set of approaches to generate highly accurate forecasts for each demand category.

4.4.2 Parameter Optimisation and Model Execution

After the selection of each method, the different parameters used in different method are selected. The parameter then has different values based on each model. Parameters need to be optimised to get optimal results from each forecasting method. The table below shows preferred methods and literature, and the parameters optimisation.

Table 2. *Different models for different demand categories (Syntetos & Boylan, 2005)*

Demand Category	Preffered Methods	Parameter Notes

Smooth	MA, ES	Short smoothing Windows (3-6 month)
Erratic	ES, MA	Moderate Smoothing (6 month)
Intermittant	Croston, Syntetos-Boylan Approximation, Teunter-Syntetos-Babai	Long smoothing windows (12 month)
Lumpy	Syntetos-Boylan Approximation, Teunter-Syntetos-Babai	Long smoothing windows (12 month)

As shown in Table 2, each demand category has different models according to the literature, which gives a better forecast. Parameter tuning ensures that the models remain responsive to genuine demand changes while filtering noise.

Smoothing Parameter (α): In exponential smoothing-based models (ES, Croston, Syntetos-Boylan Approximation, and Teunter-Syntetos-Babai), the *smoothing parameter α (alpha)* determines how much weight is given to the most recent observation compared to older observations.

- A higher α (closer to 1) makes the model more responsive to recent changes, suitable when demand patterns shift quickly.
- A lower α (closer to 0) gives smoother forecasts, suitable when demand patterns are stable but require filtering of short-term noise.

In the Teunter-Syntetos-Babai method, two smoothing parameters are used:

- α_1 for smoothing demand size, and
- α_2 for smoothing the demand occurrence probability.

These parameters are optimised to minimise forecast error across each time series. For moving average models, the smoothing window(n) represents the number of

past periods used to calculate the average. A short window captures short-term changes effectively but may introduce volatility. A long window produces stable forecasts but may lag behind actual changes.

Both smoothing parameters and windows are fine-tuned through cross-validation over historical data to achieve an optimal balance between stability and responsiveness.

4.4.3 Model Accuracy Evaluation and Best Model Selection

After every model is deployed in each demand pattern, their performance is evaluated using different statistical accuracy measures. This includes:

- **Mean Absolute Percentage Error (MAPE):** Measures relative forecast accuracy across items with different scales.
- **Mean Absolute Deviation (MAD):** Indicates the average absolute forecast error in demand units.
- **Root Mean Squared Error (RMSE):** Penalises larger errors, highlighting models that occasionally make large deviations.

For each spare part demand, the model with the lowest error value is selected as the best-performing model. Also, it is worth mentioning that forecasting results also vary based on the forecast horizon.

4.5 Theoretical Demand Validation Framework

Theoretical demand provides one additional layer that complements statistical forecasting. It estimates expected spare part demand based on physical and operational demand drivers rather than historical sales. This comparison helps identify a gap in strategy, coordination with dealers, as well as inventory strategies.

The theoretical demand is calculated using the installed machine base, expected maintenance intervals, and assumed annual operating hours. The following relationship is applied:

$$D = N * \frac{H_{annual}}{H_{maint}}$$

Where:

N = Total installed base of machines in the region

H_{annual}= Average annual operating hours per machine

H_{maint}= Recommended maintenance interval in hours for the spare part

This formula estimates how many times each spare part is expected to be replaced annually, providing a baseline demand level. It also indicates that the aftersales market loses sales if the theoretical demand is more than actual sales.

5 Case Company

This project has been conducted in collaboration with a Finnish company operating in the forestry machinery industry. The case company specialises in the design, manufacturing, and sales of forestry equipment, including harvesters and forwarders, primarily used in logging operations. In addition to machinery, the case company provides after-sales services, including technical support, maintenance, and spare parts supply.

The company operates globally, with its main market being Europe. While most international sales are handled through independent dealers, Finland remains an exception, where the company sells and services machines directly to end customers. The spare parts business is integral to after-sales operations, and spare parts are primarily delivered from the main warehouse in Finland. The dealer network of the company handles spare part distribution in other regions. The figure below shows the sales distribution in each country.

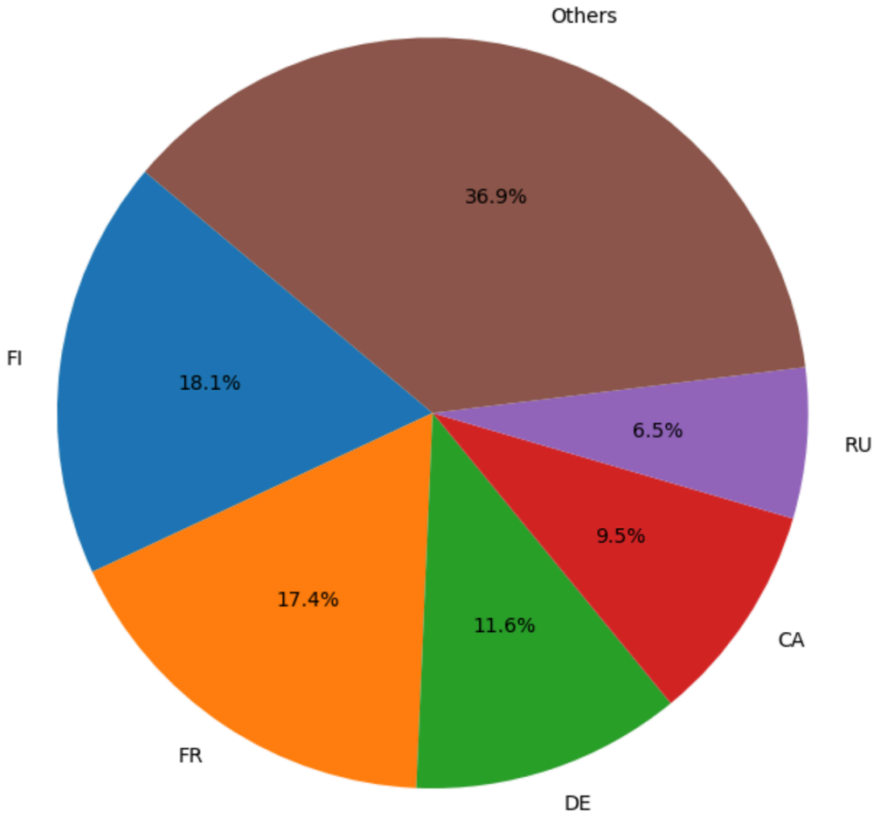

Figure 6. Biggest markets based on sales value

As shown in Figure 6, the majority of the sales are coming from Europe, with Germany and France being the leading countries. The case company has suspended all operations in Russia after 2022 and the company does not have any plans to operate in the future however this dataset takes in count of sales before 2022. The case company operates on its main markets with dealers who manage the customers and inventory in their stores. The company has a vision to improve its sales of the machine and provide better after-sales services to customers to improve the satisfaction rate, as well as revenue generation. The industry standard for after-sales revenue is 20% however case company is lagging in this. Several steps are being taken to improve this by digitalisation, better dealer management and marketing in key areas.

Currently, the management of spare parts at the case company lacks formal classification and relies on professional judgment. There is no current process of determining reorder points, safety stock levels, or inventory control mechanisms. Most procurement and stock control decisions are based on historical consumption data, supported by the experience of staff. As a result, there are concerns about inventory inefficiencies, such as holding obsolete items or missing parts critical to minimising customer machine downtime. The company also relies on dealers and their judgment for spare part demand, which concerns the lost sales to other competitors.

The current process of spare parts starts with a customer requirement for any spare part item in any region where a dealer of the case company is located.Dealers maintain the inventory of commonly sold spare parts which they provide to the customer based on demand. Dealers inventory is maintained based on expert judgment rather than forecasting models. In exceptional cases where customers want some parts which are not available with dealers then dealers send the purchase orders for specific spare parts with information on how urgent it is for the customer. After getting the purchase order, the spare parts team of the case company delivers the products from spare part inventory. If the product is not available in inventory then the spare part team sends the purchase order to the supplier to get the spare parts. Afterwards, it was sent from Finland to the region where it is needed.

The company has access to a substantial amount of historical sales data dating from 2019 onward. However, this data includes irrelevant entries, such as freight and labour costs, which require cleaning before analysis. Furthermore, there is no

integration between spare parts sales data and machine operating data, limiting the ability to build predictive models based on usage patterns.

At present, the case company does not use any forecasting models for spare parts demand. Inventory levels are largely driven by past consumption trends, without deeper analysis into customer behaviour, machine population by region, or the criticality of parts. While this approach has ensured service continuity to an extent, it may lead to excess inventory or unavailability of certain items.

This thesis aims to address these challenges by applying literature-supported methods for spare parts classification and demand forecasting, with a focus on intermittent demand. By exploring a structured approach to data-driven demand planning, the study seeks to provide the case company with practical tools to improve forecasting accuracy, optimise stock levels, and enhance overall service performance.

6 Testing the Framework

6.1 Introduction

Following the proposed framework in Chapter 4, this chapter represents the application phase, demonstrating the implementation of the proposed methodology with real data from the case company, evaluating the results from the implementation and validation using both statistical and business-driven metrics.

The initial dataset received from the case company had 9456 unique spare part product codes with 120,000 records from 2019 to 2025. The analysis covers more than 9000 SKUs across multiple regions, with a focused case study in the French region, chosen for its leading market with one dealer presence. The objective of the chapter is to demonstrate the implementation of the proposed methodology with real data and processes.

6.2 Data Cleaning & Preparation

Before proceeding with the analysis and forecasting, the raw dataset goes through several steps of data cleaning and pre-processing. These steps ensure accuracy, consistency and completeness of the information used in the upcoming stages of the framework. This step is critical, and the information revealed several missing, inconsistent and irrelevant entries that required systematic treatment. Table 3 below shows the missing values in each column of the initial dataset.

Table 3. *Initial Missing Values*

Column	Missing Value
Invoice Date	0
Customer	11
Product Code	213
Title	332
Supplier	17,019
Quantity	145
Region	11
Currency	146
Sales	1

As shown in the above table, most columns contained only minor gaps, but the supplier field had a significant number of missing records. After initial talks with the case company about it. It revealed that many unrelated entries exist in the data which do not have a supplier, so category missing values were filled with representative placeholders such as “UNKNOWN CUSTOMER” and “UNKNOWN SUPPLIER”. For numerical missing values, they were replaced with zero after verifying that they represented a true absence of the transactions rather than system errors. One exception in handling the missing values was the product code column. The application of the forecasting framework is based on the product code and its quantity. The rows which were missing product code were being removed from the analysis to maintain data completeness for the analysis.

To improve the data reliability and make the analysis for relevant products, filtering steps were performed on the raw data. During the data exploration, the key

information came out that many sales column values are negative. As they represent the credit notes or return cases rather than the true demand. These rows were removed from the analysis. In total, 4500 refund entries were excluded. The initial unique codes were over 9000, of which there were many irrelevant codes. Many of them represent such costs as freight cost, labour cost, and merchandise. The rows that contained them were removed from analysis, which were around 9000 records. After these steps, the dataset contained only genuine spare parts transactions that represent customer demand patterns.

The next step involved standardisation of the data type of each column, ensuring that during the forecasting stage, it would be easier to manage the data and apply the models. All date-related entries were converted to DATETIME format to enable aggregation based on days, months and years and time-series modelling. Other columns were cast as strings and numeric based on their character tics and what they represent. The column name was changed to industry best practice to do data analysis.

The framework introduced filtering criteria in Chapter 4. The product code which has fewer than five sales records, was excluded from analysis as there are not enough for the forecasting model to have a pattern. As well as the product codes which were not sold even one time in the last three years were excluded from the analysis as they were in the obsolete category. The total unique parts were around 9456 and Figure 7 shows the excluded parts from the analysis.

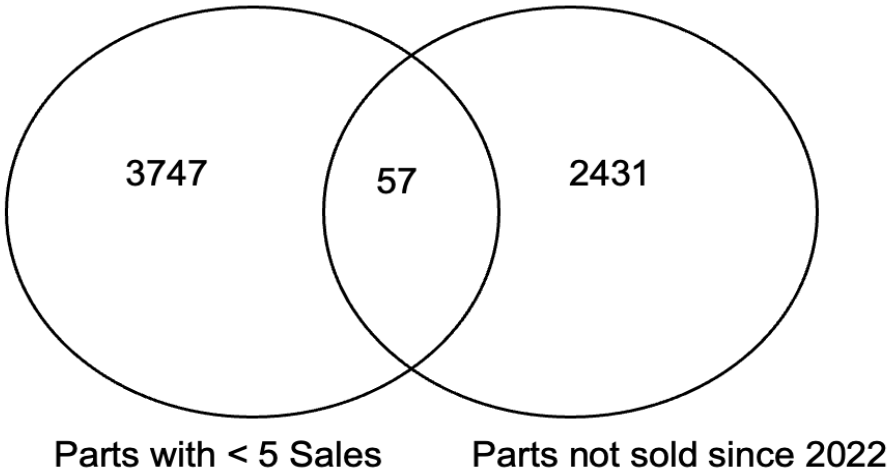

Figure 7. *Filtering criterion excluded parts*

As shown in Figure 7, most of the unique spare parts are removed from the analysis as they could not meet the filtering criteria. The total of 6235 unique parts was excluded from the analysis, where 3747 parts had fewer than five sales observations, and 2431 parts were not sold at all since 2022. There are 57 parts had fewer than five sales records; at the same time, they have not been sold since 2022. The dataset now contains over 3000 unique products, which can be analysed and used for testing forecasting methods.

Following these steps, the dataset was free of inconsistencies and ready for analysis. These changes ensured that the remaining dataset represents accurate spare part consumption behaviour across the customer base. The cleaned data now provides a good foundation for demand pattern classification and the subsequent application of forecasting models.

6.3 Demand Classification

The next step in the framework is demanding classification of the remaining unique spare part products. The objective of this step is to find out what kind of demand pattern each spare part has, and utilising this information, different forecasting models could be explored to find the optimal forecasting results for each spare part. The classification of the spare part based on the proposed methodology. Figure 7 shows the classification based on average demand interval and the squared coefficient of variation, which was calculated using cleaned monthly data.

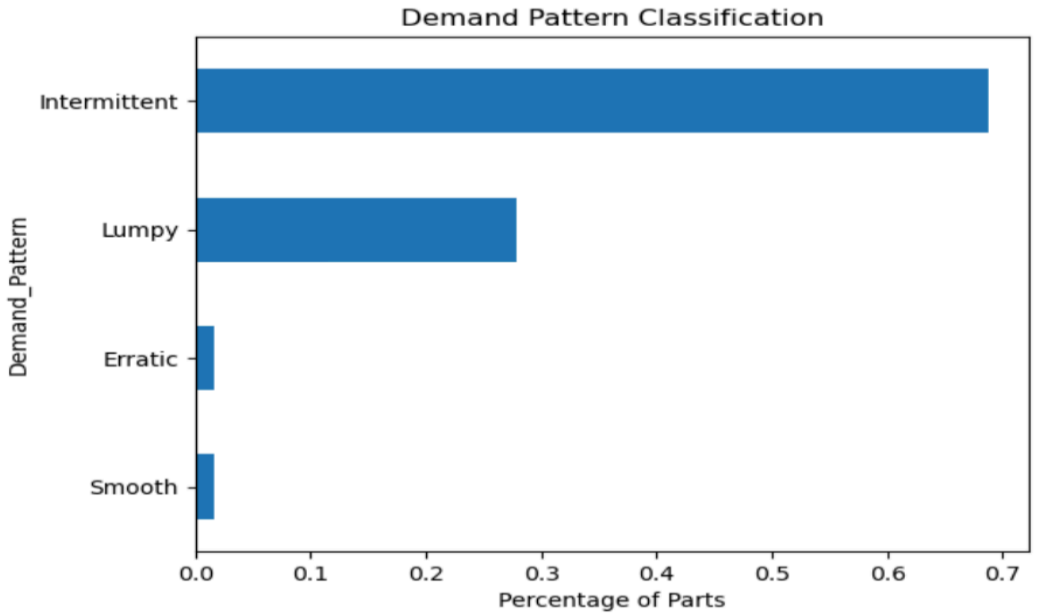

Figure 8. *Demand Patterns classification*

As shown in Figure 8 above, the Intermittent demand pattern dominates the spare part portfolio with almost 70% parts belonging to that category. The second most demanded pattern is lumpy, which has almost 25% spare parts belonging to that category. It means that most of the parts show irregular and infrequent demand, which is hard to forecast. The other remaining part has a smooth and erratic demand pattern.

6.4 Forecasting Model Selection

After the classification of spare parts based on their demand pattern characteristics, each category gives different results with different models. The forecasting models which were explored in the literature were the moving average, exponential smoothing, Croston method and variation. Each model was evaluated with different forecasting accuracy metrics such as mean absolute deviation, bias, and mean absolute percentage error.

Firstly, all the demand patterns were given with different window sizes so that each of them can have a smoother pattern which making be easier to capture the demand. The smooth demand pattern was given a 3-month window, erratic demand was given 6 months, and intermittent and lumpy demand were given a 12-month window size.

The set of figures below represents the accuracy for each model using the different accuracy metrics. The comparisons of results based on each forecasting error measurement technique were discussed earlier in the literature.

Figure 9 below illustrates the mean absolute deviation metrics to compare the accuracy of each method with each category. The key focus is on Intermittent demand, which represents most of the spare parts, and lumpy demand, which represents a quarter of the spare parts.

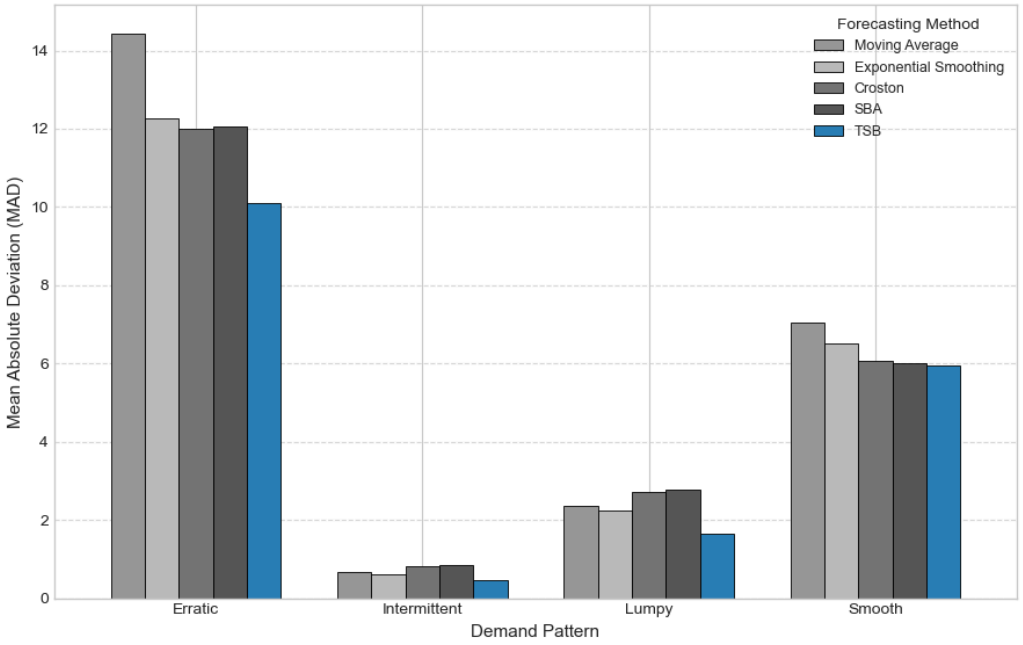

Figure 9. *Mean Absolute Error (MAD) by Forecasting Method and Demand Pattern*

The above figure shows the results obtained from applying each forecasting model to each demand pattern. In the erratic demand pattern, Teunter-Syntetos-Babai shows the lowest error, which is around 10. The smooth demand pattern shows the lowest error, 6, which was obtained in the Teunter-Syntetos-Babai method. And the most important demand patterns, Intermittent and Lumpy, obtained overall lower error in each method, but the Teunter-Syntetos-Babai performed exceptionally, with only 0.46 error in the intermittent demand pattern. Looking at the results from mean absolute deviation, Teunter-Syntetos-Babai is clearly the best model for the majority of spare parts with less than 1 unit deviation from actual demand.

The second error metric discussed in the literature was mean absolute percentage error. One of the downsides of this error is that when the demand has many zeros, it increases the demand percentage. The figure below shows the comparisons of error obtained from each model with respect to each demand pattern.

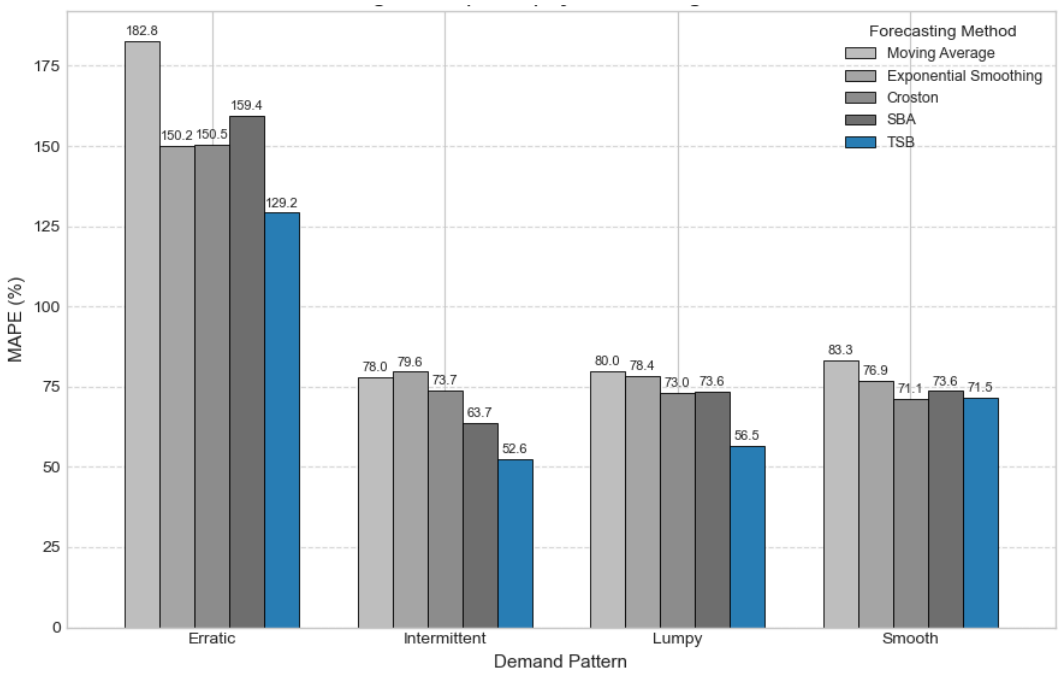

Figure 10. *Mean Absolute Percentage Error (MAPE) by Forecasting Method and Demand Pattern*

As shown in Figure 10, the same trend is followed for MAD metrics. Teunter-Syntetos-Babai performed exceptionally with the lowest error among all other forecasting methods in every demand pattern. The intermittent demand forecast was 10% better than the next best forecasting method, which saw the superiority of the Teunter-Syntetos-Babai method among all. The lumpy demand pattern was predicted with less error with Teunter-Syntetos-Babai compared to other forecasting methods, which had an over 70% error percentage.

The third error metric explored was bias. It shows how much on average method predicted over the forecast or under the forecast. Each method was applied to each demand pattern, and the obtained results are shown in Figure 11.

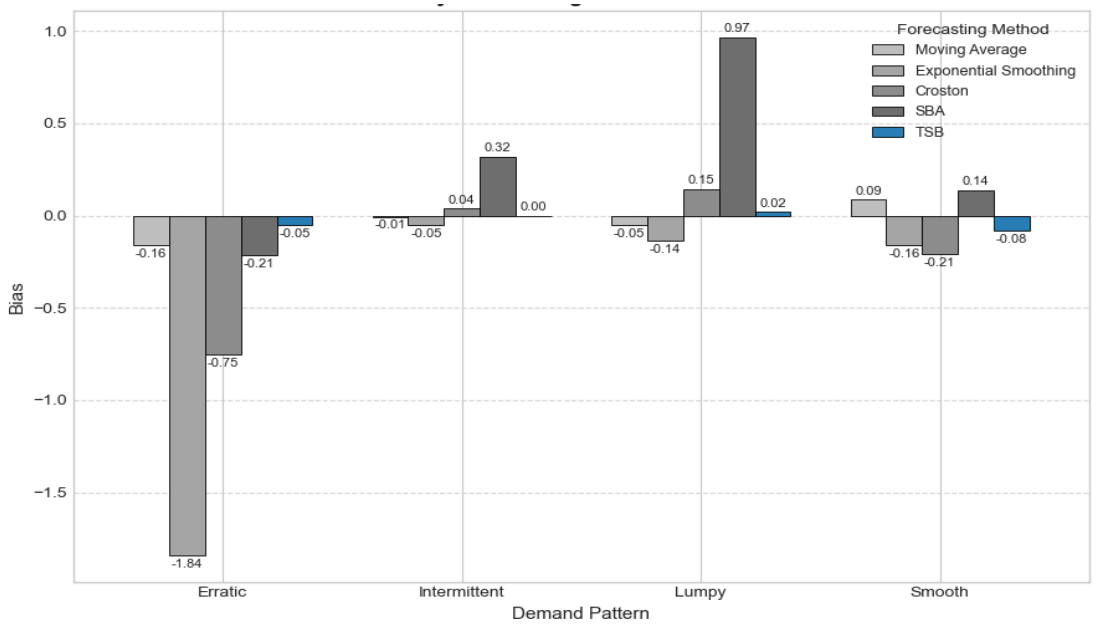

Figure 11. *Forecast Bias by Forecasting Method and Demand Pattern*

As shown in the above figure, the forecast bias was recorded near zero with the Teunter-Syntetos-Babai model in each demand pattern category. The Teunter-Syntetos-Babai model predicted neither an over- nor under forecast. Whereas Exponential smoothing and Croston were among the highest bias models in this implementation.

After the comparisons of each method on different performance evaluation metrics, Teunter-Syntetos-Babai saw the best results among the chosen forecasting models. The chosen method would be implemented in the case company for generating a forecast with a horizon of one year. The implementation of the chosen method gave a 30% weighted average improvement across all the parts included in the analysis. Meanwhile, the largest impact was obtained among over 2200 parts with intermittent demand, where it gave 34% improvement compared to the simple moving average method.

6.5 Theoretical Demand Calculation

After the selection of the best performing model, the case company wanted to explore the theoretical demand present in the specific region to improve their strategies and potential demand. The calculation of theoretical demand started with choosing the specific region for this framework. With the case company, one of the leading markets was chosen where only one dealer was present there from the company side.

The study started with assessing the number of machines sold during the last 10 years in that region, which was around 80. After that, maintenance documents were analysed, and the goal was to find the information which shows the minimum hours a customer needs to change certain parts. The assumption was taken based on the case company professional about the average annual machine our in the French region, which was assumed as 2000 hours. After getting the information, an Excel sheet was developed showing the theoretical demand for each spare part, which had maintenance data.

The following sample figure shows the comparisons of theoretical demand in the French region to actual demand, recorded and forecasted demand using the Teunter-Syntetos-Babai method.

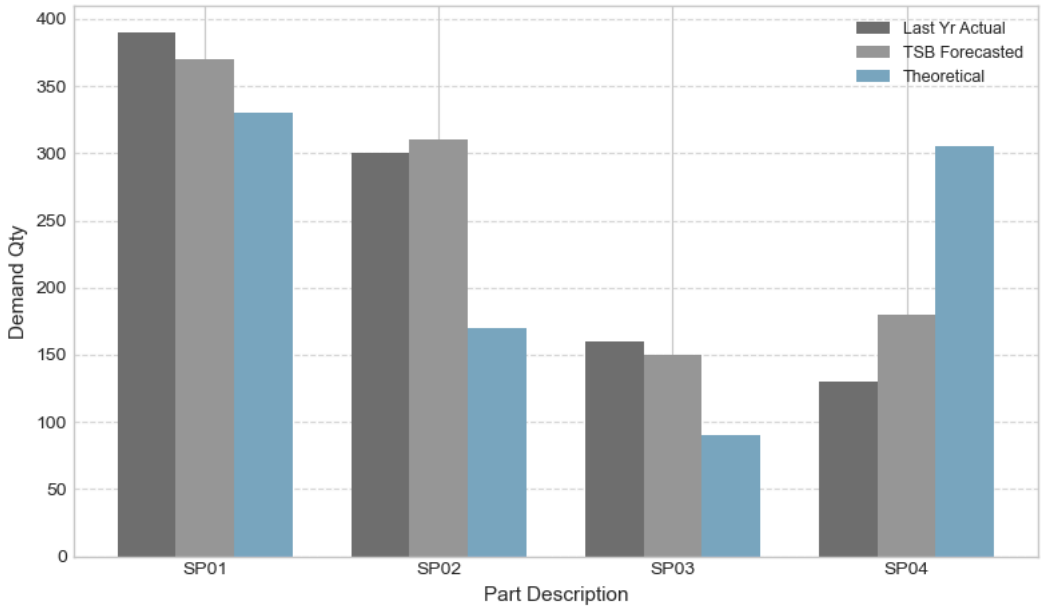

Figure 12. *Three-way comparisons in French region*

As shown in the above figure, a comparison sheet was developed showing last year demand, forecasted demand and theoretical demand. The results show that many parts had less theoretical demand than actual demand. As per the company manager, it could happen because of machines which were sold 10 years still in use, and they have higher wear and tear because of their age. This gave the case company a good insight into how they can collaborate with the dealer and have more inventory with the dealer on specific parts.

Another conclusion data gives are that they are cross-selling of spare parts with competitors, and how they can spend more efforts into marketing and price point to have more cross-selling demand. The overall framework gave the case company

with best performing model as well as the comparison sheet where they can make informed decisions based on the data available.

7 Conclusion

7.1 Key Findings

The objective of this thesis was to develop a method for forecasting spare part demand by examining various statistical models from the literature and identifying the best-performing model for the case company. The results will be used to inform decisions related to inventory, dealer management, and marketing strategies. It was found that the moving average and exponential smoothing models were ineffective due to the nature of spare part demand. Hence, Croston and variation were explored to reach the objective. The target was to find the sales quantity on a yearly basis to make informed decisions.

The best model based on prediction accuracy is the Teunter-Syntetos-Babai model. Using the Teunter-Syntetos-Babai model, 0.46 accuracy could be achieved in intermittent demand, which is 0.22 higher than the simple moving average model. With lumpy demand, the prediction accuracy of 1.66 is achieved, where the Teunter-Syntetos-Babai method gives is 0.70 higher prediction than the simple moving average model. Moreover, it is well known that spare part demand forecasts are uncertain regardless of the forecast method used, so it is not realistic to expect full accuracy. Consequently, we suggest that the case company apply the selected Teunter-Syntetos-Babai model for forecasting spare part demand for inventory and dealer-related decisions.

However, it is surprising how much difference the Croston and variation methods have among themselves. Croston and Syntetos-Boylan Approximation intermittently had less accuracy than basic models, even though the literature mentions Croston and Syntetos-Boylan Approximation as models specifically for intermittent demand in spare parts. However, Teunter-Syntetos-Babai proved to have way higher accuracy in all metrics compared to other Croston variation methods. The case company will have more data over time, and using more observations might lead to different accuracy among all the models, or it might improve accuracy because of more data points to train the models.

The case company can use theoretical demand to find more answers related to demand in a certain region and for certain products. They can utilise it to set the target sales as well as use it for marketing and dealer lead initiatives. The surprising

results were that many spare parts had less theoretical demand, but actual sales were higher. The manager utilised the data and understood that we did not include machines sold 10 years ago because of data constraints, which require more spare parts. Moreover, customers of the competitor company could have bought the spare part from the dealer, which could increase the sales for certain parts.

In all, this study can be considered as successful since the five tested models were able to select the best performing model for spare parts demand forecasting with a favourable accuracy, and the method is chosen especially for the needs of the purchasing function in the case company. As well as the theoretical demand framework for the business development function, which could be helpful for future decisions. Thus, the objective was fully achieved. The implications of the findings are discussed in more detail in the next chapter.

7.2 Theoretical and Practical Implications

This research makes a good theoretical contribution to the domain of spare part demand forecasting by developing and validating the proposed framework that integrates demand classification, forecasting model selection, and developing theoretical demand and comparison of theoretical, actual and forecasted demand. While prior literature has talked about these elements independently, very few studies have combined them into a single and systematic approach. The proposed framework contributes to bridging this conceptual gap by demonstrating how theoretical demand understanding can be operationalised into practical improvements in after-market business.

From the theoretical point of view, the study reinforces and extends the conceptual linkage between demand pattern classification and forecasting model selection. Earlier works, such as Syntetos and Boylan (2008) and Croston (1971), have established statistical properties of intermittent demand to influence the demand forecasting accuracy. This study builds upon that foundation by showing that not only the selection of a forecasting model, but theoretical demand can give more insights for the spare part business.

Moreover, this research supports the multi-level framework in forecasting theory. By situating the model assessment in statistical and operational perspectives, the study contributes to growing future research in a hybrid approach to forecasting. It serves as foundational knowledge about implementing the multi-layer framework in the spare parts demand forecasting case.

In addition to a theoretical contribution, this case study provides the framework which really has practical relevance for spare parts managers, planners, and decision-makers in after-sales operations. The integrated approach provided a clear, data-driven process for understanding demand behaviour, selecting suitable forecasting models, and validating performance against both actual outcomes and theoretical expectations.

Firstly, the demand classification steps help practitioners to differentiate between different demand behaviours, such as smooth, intermittent, lumpy, and erratic, based on objective criteria. This classification supports more informed decisions in upcoming steps, such as model selection, which links forecasting approaches to demand patterns with higher accuracy. For instance, while moving average methods may perform well for smooth demand, Croston-based models or Teunter-Syntetos-Babai approaches are more appropriate for intermittent and lumpy patterns. The framework thus acts as a diagnostic tool that helps practitioners assign forecasting models strategically rather than relying on trial and error.

Secondly, by incorporating multiple forecasting methods such as Moving Average, Exponential Smoothing, Croston, Syntetos-Boylan Approximation, and Teunter-Syntetos-Babai and evaluating them through performance metrics like MAPE and Bias, the framework provides a good performance comparison. This comparative analysis helps practitioners identify the method that consistently delivers the lowest error and minimal bias across different part categories. For organisations managing thousands of spare parts, such an evidence-based model selection process can significantly improve forecast reliability and reduce the workload associated with manual forecast adjustments.

Third, the inclusion of theoretical demand provides a practical control mechanism for forecast validation. By comparing forecasted values and actual consumption against theoretical demand expectations, planners can make informed decisions on the inventory as well as more strategic decisions. For instance, the case company found out that some of the products in one region are having less demand compared to theoretical demand, they started to find the reason behind it, which led them to find that customers are buying from other dealers because of less lead time or price difference.

From a broader operational perspective, implementing this framework can enhance inventory management and cost efficiency. Accurate forecasts reduce excess stock levels, minimize obsolescence, and prevent stockouts, all of which

contribute to improved service levels and customer satisfaction. The framework also facilitates communication between data analysts, maintenance planners, and procurement managers by providing a common, structured understanding of demand dynamics. In practice, this can improve collaboration and decision transparency across supply chain functions.

Finally, the modular design of the framework makes it scalable and transferable. Organisations can adopt it incrementally, starting with demand classification and model evaluation, and later integrating theoretical demand benchmarking as data maturity increases. This scalability makes the framework suitable for both large enterprises and smaller firms seeking to formalise their forecasting process from scratch.

In summary, the practical implications of this research lie in its potential to transform spare parts forecasting from a reactive, model-driven activity into a structured, context-aware, and theory-informed decision process. By linking demand classification, model selection, and theoretical demand evaluation, the framework not only improves forecasting accuracy but also strengthens organisational learning and operational resilience.

7.3 Limitations and Future Research

A few limitations of this study need to be acknowledged. First, the statistical models were explored because the case company had no prior experience with data-driven demand forecasting. The company was relying on professional judgment for demand management. The case company did not have different classes based on different properties of the spare part, which is a challenge when applying the model to all the parts together.

Lastly, there are a couple of concerns regarding the data used. First, there was a limited amount of data available in the case company. Many unique parts corresponded to the amount of data the sales point dataset had. Even though six years of data were collected still it was still limited to the SKU level, which can lead to less accuracy in the forecast. While considering including theoretical demand, one of the limitations was that maintenance documents did not have data for all the parts. There were a few parts that had the data to calculate the theoretical demand, which limited the study's focus to only a few parts for comparison.

When the limitations are kept in mind, the research and results could be considered as good quality for the case company to make decisions based on that. The study was designed in a way that it contributes validity and reliability. For example, the accuracy metrics used validate the reliability of the generated forecast. This study reveals potential directions for future research. The combination of theoretical demand with forecasting was a good addition to the framework. However, the limitations of the data make it challenging to explore more advanced models, such as machine learning models. This could be a good topic to explore with more advanced models, and better maintenance data can lead to better decisions on large products.

The development of the key product indicator could be a good future research direction with this framework. The development of new KPIs based on theoretical demand and actual demand would be a great direction in the improvement of data-driven tracking and decision-making. Also, the inventory policies and dealer management decisions can be taken based on this framework.

Overall, the future direction based on this framework could be towards more complex models, development in key performance indicators, as well as new inventory policies and dealer management policies based on demand forecasting data.

8 REFRENCES

[1] A.H. C. Eaves, B. G. K. (2004). Forecasting for the Ordering and Stock-Holding of Spare Parts on JSTOR. Journal of the Operational Research Society (2004) 55, 431-43, 431–437.

[2] Ahlsell, L. , Jalal, D., Khajavi, S. H., Jonsson, P. ;, & Holmström, J. (2023). Additive Manufacturing of Slow-Moving Automotive Spare Parts: A Supply Chain Cost Assessment. Journal of Manufacturing and Materials Processing, 7(1).

[3] Armstrong, J. S., & Green, K. C. (2017). Demand Forecasting II: Evidence-Based Methods and Checklists. Working Paper 89-clean, The Wharton School, University of Pennsylvania.

[4] Armstrong, J. S. (2001). Principles of Forecasting: A Handbook for Researchers and Practitioners. Springer Science & Business Media.

[5] A.S, B., & Ramanathan, U. (2021). The role of digital technologies in supply chain resilience for emerging markets' automotive sector. Supply Chain Management, 26(6), 654–671.

[6] Bacchetti, A., & Saccani, N. (2012). Spare parts classification and demand forecasting for stock control: Investigating the gap between research and practice. Omega, 40(6), 722–737.

[7] Bayraktar, D., Selcen, D., & Aykaç, O. (n.d.). MULTI CRITERIA CLASSIFICATION FOR SPARE PARTS INVENTORY.

[8] Boone, C. A., Craighead, C. W., & Hanna, J. B. (2008). Critical challenges of inventory management in service parts supply: A Delphi study. Operations Management Research, 1(1), 31–39.

[9] Boylan, J. E., & Syntetos, A. A. (2008). Forecasting for Inventory Management of Service Parts. Springer Series in Reliability Engineering, 8, 479–506.

[10] Cavalieri, S., Garetti, M., MacChi, M., & Pinto, R. (2008). A decision-making framework for managing maintenance spare parts. Production Planning & Control, 19(4), 379–396.

[11] Christopher, M., & Lee, H. (2004). Mitigating supply chain risk through improved confidence. International Journal of Physical Distribution & Logistics Management, 34(5), 388–396.

[12] Fortuin, L., & Martin, H. (1999). Control of service parts. International Journal of Operations and Production Management, 19(9), 950–971.

[13] Gajpal, P. P., Ganesh, L. S., & Rajendran, C. (1994). Criticality analysis of spare parts using the analytic hierarchy process. International Journal of Production Economics, 35(1–3), 293–297.

[14] Ghewari Asita Ashok. (2020). (PDF) An Inventory Control using ABC Analysis and FSN Analysis. International Journal of Engineering, Business and Enterprise Applications (IJEBEA).

[15] Hasni, M., Aguir, M. S., Babai, M. Z., & Jemai, Z. (2019). Spare parts demand forecasting: a review on bootstrapping methods. International Journal of Production Research, 57(15–16), 4791–4804.

[16] Hatefi, S. M., Torabi, S. A., & Bagheri, P. (2014). Multi-criteria ABC inventory classification with mixed quantitative and qualitative criteria. International Journal of Production Research, 52(3), 776–786.

[17] Hmida, J. Ben, Parekh, S., & Lee, J. (2014). Integrated inventory ranking system for oilfield equipment industry. Journal of Industrial Engineering and Management, 7(1), 115–136.

[18] Hua, Z. S., Zhang, B., Yang, J., & Tan, D. S. (2007). A new approach of forecasting intermittent demand for spare parts inventories in the process industries. Journal of the Operational Research Society, 58(1), 52–61.

[19] Huiskonen, J. (2001). Maintenance spare parts logistics: Special characteristics and strategic choices. International Journal of Production Economics, 71(1–3), 125–133.

[20] Huynh, N. (n.d.). LUT School of Business and Management Master's Programme in Supply Management Demand forecasting of spare parts-case study from automotive industry.

[21] Hu, Q., Boylan, J. E., Chen, H., & Labib, A. (2018). OR in spare parts management: A review. European Journal of Operational Research, 266(2), 395–414.

[22] Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. International Journal of Forecasting, 22(4), 679–688.

[23] Mentzer, J. T., & Moon, M. A. (2005). Sales Forecasting Management: A Demand Management Approach (2nd ed.). SAGE Publications.

[24] Molenaers, A., Baets, H., Pintelon, L., & Waeyenbergh, G. (2012). Criticality classification of spare parts: A case study. International Journal of Production Economics, 140(2), 570–578.

[25] Mukhopadhyay, S.K., Pathak, K. and Guddu, K. (2003),... - Google Scholar. (n.d.). Retrieved April 17, 2025,

[26] Nurcahyo, R., Malik, F. M., & Farizal. (2017). Aircraft spare parts inventory management using multi-criteria classification with AHP approach. 4th IEEE International Conference on Engineering Technologies and Applied Sciences, ICETAS 2017, 2018-January, 1–5.

[27] Patton, J. D. ., & Steele, R. J. . (2003). Service parts handbook. 556.

[28] Pinçe, Ç., Turrini, L., & Meissner, J. (2023). Intermittent demand forecasting for spare parts: A critical review. Quinlan School of Business, Loyola University Chicago.

[29] Ramanathan, R. (2006). ABC inventory classification with multiple criteria using weighted linear optimization. Computers & Operations Research, 33(3), 695–700.

[30] Roda, I., Macchi, M., Fumagalli, L., & Viveros, P. (n.d.). A review of multi-criteria classification of spare parts: From literature analysis to industrial evidence.

[31] Stevenson, W. J. (2021). Operations Management (14th ed.). McGraw-Hill Education

[32] Syntetos, A. A., Boylan, J. E., & Croston, J. D. (2005). On the categorization of demand patterns. Journal of the Operational Research Society, 56(5), 495–503.

[33] Teixeira, C., Lopes, I., & Figueiredo, M. (2017). Multi-criteria Classification for Spare Parts Management: A Case Study. Procedia Manufacturing, 11, 1560–1567.

[34] Teixeira, C., Lopes, I., & Figueiredo, M. (2018). Classification methodology for spare parts management combining maintenance and logistics perspectives. Journal of Management Analytics, 5(2), 116–135.

[35] Van der Auweraer S, Boute RN, Syntetos AA. Forecasting spare part demand using service maintenance information. International Journal of Production Economics.

[36] Williams, T. M. (1984). Stock control with sporadic and slow-moving demand. Journal of the Operational Research Society, 35(10), 939–948.

[37] Zheng, L., Terpenny, J., & Sandborn, P. (2015). Design refresh planning models for managing obsolescence. IIE Transactions, 47(12), 1407–1423. https://doi.org/10.1080/0740817X.2014.999898