UNIVERSITY OF LJUBLJANA

SCHOOL OF ECONOMICS AND BUSINESS

MASTER’S THESIS

DEMAND FORECASTING FROM SALES DATA IN THE PRESENCE OF STOCKOUTS

Ljubljana, February 2024

KERIM MAŠOVIĆ

AUTHORSHIP STATEMENT

The undersigned Kerim Mašović a student at the University of Ljubljana, School of Economics and Business, (hereafter: SEB LU), author of this written final work of studies with the title DEMAND FORECASTING FROM THE SALES DATA IN THE PRESENCE OF STOCKOUTS, prepared under supervision of Marko Jakšič

D E C L A R E

1. this written final work of studies to be based on the results of my own research;
2. the printed form of this written final work of studies to be identical to its electronic form;
3. the text of this written final work of studies to be language-edited and technically in adherence with the SEB LU's Technical Guidelines for Written Works, which means that I cited and / or quoted works and opinions of other authors in this written final work of studies in accordance with the SEB LU's Technical Guidelines for Written Works;
4. to be aware of the fact that plagiarism (in written or graphical form) is a criminal offence and can be prosecuted in accordance with the Criminal Code of the Republic of Slovenia;
5. to be aware of the consequences a proven plagiarism charge based on the this written final work could have for my status at the SEB LU in accordance with the relevant SEB LU Rules;
6. to have obtained all the necessary permits to use the data and works of other authors which are (in written or graphical form) referred to in this written final work of studies and to have clearly marked them;
7. to have acted in accordance with ethical principles during the preparation of this written final work of studies and to have, where necessary, obtained permission of the Ethics Committee;
8. my consent to use the electronic form of this written final work of studies for the detection of content similarity with other written works, using similarity detection software that is connected with the SEB LU Study Information System;
9. to transfer to the University of Ljubljana free of charge, non-exclusively, geographically and time-wise unlimited the right of saving this written final work of studies in the electronic form, the right of its reproduction, as well as the right of making this written final work of studies available to the public on the World Wide Web via the Repository of the University of Ljubljana;
10. my consent to publication of my personal data that are included in this written final work of studies and in this declaration, when this written final work of studies is published.
11. that I have verified the authenticity of the information derived from the records using artificial intelligence tools.

Ljubljana, ________________________ Author's signature:_________________________

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF APPENDICIES

LIST OF ABBREVIATIONS

AIC	Akaike Information Criterion
AR	Autoregressive
ARIMA	Autoregressive Integrated Moving Average
ARMA	Autoregressive Moving Average
DES	Double Exponential Smoothing
MA	Moving Average
MAD	Mean Absolute Deviation
MAPE	Mean Absolute Percentage Error
MSE	Mean Squared Error
RMSE	Root Mean Squared Error
SARIMA	Seasonal Autoregressive Integrated Moving Average
SES	Simple Exponential Smoothing

1 INTRODUCTION

Demand forecasting is a critical business process which specifically focuses on predicting the quantity of goods or services that customers are likely to purchase within a specific time frame. Demand forecasting is something that is present in every company today, and even people that are not in the supply chain, or analytics, have probably come across this term. Forecasting can be defined as accurately predicting the future, given all of the information available, including historical data, as well as knowledge and awareness about any event that might have impact on the forecast (Hyndman & Athanasopoulos, 2018).

Being able to accurately forecast demand is not only an added value activity, rather, it is absolutely necessary for survival. Accurate forecasting will result in fewer shortages, less useless inventory, and more sales. However, forecasting demand can be challenging, particularly in situations where stockouts occur. The stockouts truncate the demand distribution and, if uncorrected, wrong results are obtained (Bell, 2000).

As mentioned earlier, an accurate forecast can reduce costs. It reduces costs by allowing the company to successfully manage inventory and avoid overstocking or understocking. To successfully manage the inventory, companies rely on historical data, mainly past sales which are used to compute the demand forecast. If the specific item for which we want to predict future demand was always in stock and available for sale, then the company could do forecasting based on past sales data without any modification and adjustments. But unfortunately, that is not always the case. Occasionally a stockout for a specific item will occur causing the amount sold to be less than the amount demanded (Wecker, 1978).

Various other factors besides stockouts can influence the accuracy of the forecast. Some of those factors are hard to predict, while some can be spotted and considered prior to forecast. For example, some of the factors that can be taken into consideration are weather, holidays, big sports events, promotions, marketing campaigns, etc. (Ivanov et al., 2018). Fortunately, the majority of holidays or big sports events happen on the same date every year, so for some cases, it is enough to rely on historical data.

The problem that many companies are facing is that they do not distinguish between forecasting demand and forecasting sales. Demand is about what the customer wants and not about what was sold or shipped. Tracking sales only is alright as long as the company has enough stock (is not facing stockouts), and in that case demand and sales are equal. Usually, forecasting disregarding stockouts will give a picture that there is less demand for a product than there actually is. However, this is not always the case. When facing stockouts for one product, the effect of substitution comes in play. This effect will lead to selling more of a substitute product than there is original demand (Vandeput, 2023a).

As a master student of supply chain and logistics, and data engineer and data analyst by profession, I found this topic really interesting and suitable for me. My passion for this topic comes from my desire to combine theoretical knowledge obtained during the last two years of master's studies at University of Ljubljana, School of Economics and Business, and practical knowledge and experience in working with data. For the purpose of this thesis, I used R (programming software) to demonstrate how different data analytics techniques can be used to improve demand forecasting accuracy, how to handle stockouts, and how to measure accuracy of the forecast. R software contains many useful packages for time series forecasting, statistical modeling, as well as plotting the graphs, and that is why it is a popular choice for many data analysts and data scientists.

In this master thesis, I dove deeply into the topic of demand forecasting. I explained what the best practices are, how to forecast demand, what statistical methods to use, what to include in the forecast, as well as what to avoid, how to handle stockouts in the sales data, and how to compare results of different forecasting methods. This master's thesis consists of both theoretical research and practical examination. The theoretical section entails a descriptive explanation of demand forecasting and various forecasting methods using existing literature. To demonstrate this issue practically, I used data set generated with Python. Data set is described in more details later in the thesis.

In the analysis, I employed two forecasting methods. Based on the data set, representing the sales figures with random stockouts spanning four to seven days within 25 periods, I aim to forecast the demand. Forecasting methods used for this specific problem and data set are triple exponential smoothing (Holt-Winters) and Auto ARIMA method. The first method was adopted due to the datasets trend and seasonality attributes. While the data set trend and seasonality attributes suggested the use of SARIMA, I opted for Auto ARIMA for the second method due to its advantages over solely relying on SARIMA. This is because Auto ARIMA is a convenient function in the forecast package that automatically determines the best ARIMA parameters for a given time series, using a combination of unit root tests, minimization of the Akaike Information Criterion (AIC) and maximum likelihood estimation. Naturally, the Auto ARIMA function chose the SARIMA method, without requiring me to manually pick most optimal parameters.

Of course, prior to implementing these forecasting methods, I needed to fill in the gaps in the sales data using different approximation methods. These are primarily divided into average of last known periods method, two interpolation methods, and time series imputation algorithms using ImputTS package in R software. At the end results are compared and conclusion is made.

As a result of the work described above, I was able to answer the main research question:

- How to forecast demand from the sales data in the presence of stockouts?

Additionally, I was able to answer the following research sub-questions:
- How to approximate demand during stockout periods?
- How to measure the accuracy of the forecast using appropriate statistical measures and what is the improvement of the accuracy of the demand forecast when adjusting for stockouts?

The structure of the thesis follows a logical progression, beginning with the introduction and followed by chapters that delve into different aspects of demand forecasting, and later into the practical problem and applied solution. Initially, chapter 2 explores the role of demand forecasting, laying the foundation for understanding the significance of accurate forecasting. Within this chapter, an exploration forecasting principles and processes is presented, providing a theoretical framework. Subsequently, in chapter 3, an emphasis is put on demand forecasting methods, categorizing them into qualitative and quantitative methods. Even though every forecasting method is not going to be used during the practical part, it is important to get an overall picture of different methods. This can provide the reader with clarity on our selection process and the rationale behind choosing specific methods for this particular problem. In chapter 4, the focus is shifted to the thesis main research question: demand forecasting from sales data with the presence of stockouts. The chapter introduces the problem of stockouts in sales data, what problems can this cause, and how to deal with this issue. A thorough discussion of different methods for approximating demand during these situations is also presented in this chapter. This part is crucial for the thesis, since accurate forecast cannot be done without approximating stockouts data first.

The practical part of the thesis starts at chapter 5. Within this chapter, the dataset that is going to be used for problem demonstration is introduced. An explanation of how the dataset was generated and its attributes is presented. After that, different approximation methods discussed in chapter 4 are applied to the dataset. This section is beneficial in explaining the logic behind using specific approximation methods, their behavior, and their implementation in R software.

The focus of the first part of chapter 6 is the application of forecasting techniques on "clean" data with approximated stockouts. Two mentioned forecasting methods were applied on data adjusted by each approximation method. This was done using the for loop to make this forecast efficient and to make sure that everything undergoes the same process. Next part of chapter 6 involves the comparison of results generated by different forecasting and approximation methods. Here, the comparison between different methods and their accuracy is presented, together with visual representation of findings. This offers a comprehensive evaluation of the forecasting and approximation method used for this specific case.

Chapter 7 features a conclusion with summarized results and a final comment on how to handle this specific issue.

All code written in Python and R can also be found on my personal GitHub. Different codes with specific problems will be separated in individual files, so it is easier to find whatever is looked for. Codes are presented in the Appendices and can be accessed on GitHub through the link.

The entire data set can be found on my personal Kaggle profile on this link. Link to my profile, and hence the data set can also be found on GitHub under readme.md file.

2 DEMAND FORECASTING

Forecasting is crucial for business planning. If forecasts are too high, the company can overstock (overproduce), and if it is another way around, the company is missing sales (Ivanov et al., 2018). In both scenarios, the company is facing additional costs. Either it is wasting material, workforce, and disrupting entire logistics, or it is losing sales and potentially losing customers. Not to get the wrong impression, forecasting demand is employed by both the companies that are selling physical goods and the ones that are service oriented. However, to forecast demand in the service industry a company cannot turn to just forecasting sales, rather it has to be more focused on forecasting demand by tracking various metrics, such as number of customers in a restaurant during lunch, call volumes at a call center, number of orders, etc. (Dorne et al., 2008).

Forecasting demand now, and forecasting demand a few years back, are hardly the same things. The principle is the same, and the desired outcome is the same, but the tools and abilities that now exists are much better. The introduction of artificial intelligence and machine learning has changed to the practice of forecasting and are now sources of competitive advantage. These tools allow companies to do data science and extract meaningful conclusions from their data. Looking into the past, forecasting demand required forecasting software which used complex machines, complex algorithms, and not every company could afford it. However, things have changed today. With an increase in computing power, availability and know-how of data collection, better methods, free tools (and free packages that come within) one person can make a difference. People usually associate data science and data analytics with hard-core IT and coding. Of course, coding skills and good understanding of these concepts are necessary but what is even more important is to have the right scientific mindset (Vandeput, 2021). To truly be able to get everything out of data, one must have a tendency to observe, experiment, question the results, look for patterns, as well as have the ability to be creative and think outside of the box.

A common practice is to predict demand based on past sales. If the item for which demand is being forecasted was always in stock, and available in sales, any forecasting technique could be applied without any modification (Wecker, 1978). However, items that are demanded for sales are sometimes not in stock, and hence the stockout occurs. This can significantly influence the accuracy of the forecast since it gives the false picture of the demand. That is why companies should always make forecasts based on demand, not sales. Nevertheless, if the company did track only sales and not demand, it will be necessary to approximate demand. The issue of stockout is something that I will be specifically focusing on in this master thesis, where demand approximation of real-life demand will be applied to fill out these gaps in data, and different forecasting methods will be tested using new "corrected" data.

2.1 Role of demand forecasting

Forecasting can be defined as accurately predicting the future, given all of the information available, including historical data, as well as knowledge and awareness about any event that might have impact on the forecast (Hyndman & Athanasopoulos, 2018). Demand forecasting might be one of the most important concepts in every supply chain, and it is something that definitely controls and influences every other activity in the supply chain. It is used to determine the estimated demand for the future and prepare the organization to provide enough level of stock to match that demand (Adhikari et al., 2019).

Forecasting demand is a critical aspect of supply and operations management. The results of demand forecasting can have a direct impact on various aspects of the logistics process, including production, transportation, warehousing, and inventory planning. In addition, demand forecasting can also inform strategic decisions such as facility location planning, revenue management, and process design (Ivanov et al., 2018). Different functional roles within companies might use forecasts to support their business activities. As an example, logistic manager can use short-term forecast to decide what quantity of goods should he display in the shop, plant manager might use medium-range forecast results to see the need for hiring additional personnel or purchasing new equipment, or marketing strategist can realize the need for developing new products or entering new markets based on long-term forecast.

Forecasting is an activity that is required in many situations, from deciding how much bread a bakery needs to bake for a day to deciding whether to build an additional factory five years in future. Hence, in order to have an effective and efficient supply chain, forecasting is needed regardless of the forecasting horizon (Hyndman & G Athanasopoulos, 2018). It is important to note that, in general, the longer the forecasting horizon, the less accurate the forecasting will be (Adhikari et al., 2019). This happens because, in the long time period, there is a higher chance of various factor influencing the demand, in comparison with short-term forecasting. Accurate forecasting involves identifying patterns and relationships in historical data disregarding the events that will probably not happen again. It is important to distinguish between random fluctuations in past data and meaningful patterns that should be modeled and projected (Hyndman & G Athanasopoulos, 2018). For example, if the sales were off due to pandemic for two months, that does not mean that during the same period next year sales will be off again. However, if there are some events like festivals, that are occurring regularly every year at the same time, a good forecasting method should recognize that and include it in prediction. That is why the more data exists, that is, the more data goes back in time, the degree of forecasting accuracy will increase. On the other hand, if a company were to introduce a new product line, implement a new business policy, or launch a marketing campaign, relying solely on past data to forecast demand could be misleading. In such cases, it is important to consider the impact of these new developments and adjust forecasting methods accordingly.

What makes demand forecasting fairly complex is the fact that demand sometimes changes without prior warning, and there is hardly a method that can predict that. For example, pandemics like COVID-19, where almost every industry faced profit loss due to demand crunch or supply shock, while health industry faced shortage of protective equipment and medical supplies (Ivanov et al., 2018). Factors similar to these ones are something that makes usage of traditional statistical methods challenging.

2.2 Forecasting principles and processes

Regardless of what is being forecasted, there are some principles that always hold true. According to Sanders (2017), there are three main forecasting principles. The first one is that the forecasts are rarely perfect. This holds true mostly due to the fact that there are so many factors in a business environment that are hard to predict. That is why forecasting error exists. Forecasting error is the difference between forecasted value and the actual value, and ‘acceptable’ value of the error is defined by each company individual. The second principle states that forecasts are more accurate for groups than for individual items. This implies that while individual items within a group may be unstable, the overall data for the group can remain stable. Consequently, forecasting accuracy can be significantly improved by focusing on the group as a whole rather than individual items. The third principle implies that forecasts are more accurate for shorter than longer time-horizon. This principle holds, because in the short time period, there is less uncertainty. When forecasting for shorter time periods, such as one week, it is less likely that unforeseen events or disruptions will significantly impact the forecast. However, when forecasting for longer time horizons, such as a year in advance, there is a higher probability of encountering various disruptions that can influence the forecast.

As it is already established, forecasts have a significant impact on a company's profitability and success, making it essential to create them effectively. This involves following a specific set of steps and processes. It is crucial to monitor and control the forecasting process from beginning to the phase of implementation of decisions. To achieve this, the good demand forecast ought to employ a specific framework. For example, such frameworks can be similar to the one presented in Figure 1 proposed by Vandeput (2023a), or to the one from Figure 2 proposed by Ivanov et al. (2018).

Figure 1: Five-step framework for demand planning

Source: Vandeput (2023a).

Let’s now take a closer look at the specifics of the individual components of the five-step framework for demand planning.

Objective/Define Objective: At the core of both frameworks is the need to emphasize the purpose of the forecast. Both Vandeput (2023a) and Ivanov et al. (2018) emphasize the importance of defining why a forecast is needed. For example, what decisions will this forecast support? Who are the primary stakeholders? How long is the forecast horizon? A well-defined objective provides clarity on the requirements and ensures the subsequent steps are aligned with the business needs.

Data/Data Acquisition: Both frameworks stress the importance of data. Vandeput (2023a) emphasizes the necessity of high-quality data, as the availability of data alone is not sufficient. This entails having data that accounts for variables such as stockouts or includes information about significant events that could affect demand. Ivanov et al. (2018) focuses more on different ways in which data can be acquired. Thus, he provides a practical example, suggesting that if the objective is to determine how many handouts should be printed for the next lecture, the data required would be the attendance list.

Metrics: Vandeput in his framework talks about importance of measuring the accuracy and effectiveness of a forecast. Accuracy is usually measured by splitting data into training and test set, applying forecasting method on training set and comparing forecast outcome with the actual results. Different metrics such as mean absolute percentage error, mean absolute error, and similar can be used to evaluate forecasting accuracy. Vandeput also mentions the value-weighted metrics, which emphasizes the differentiation between products based on their importance. This means that not every product is equally important, and that importance should be emphasized. In other words, the value of forecasting error should be weighted based on product unit cost. So, if the forecasting error is 10 units for a product worth 2€, value-weighted error is 20€. While, for more valuable products that have, for example, value of 50€ per unit, forecasting error of only one unit has a much higher value-weighted error. Ivanov et al. (2018) is not explicitly mentioning metrics but does hint at the importance of selecting the right method for forecasting, which basically means comparing different methods based on their accuracy and selecting an appropriate one.

Baseline Model/ Select the Method: Once the data has been acquired and initial analysis (data cleaning, pattern identification, etc.) has been conducted, the next step is to create or select the appropriate forecasting method. Generally, forecasting methodologies can be categorized into two main approaches. The first approach that is more mathematical, which relies on statistical which rely on data, and the less rigid approach that involves expert judgment (Dorne et al., 2008). Choosing the right forecasting method depends on various factors such as nature of the data, the time horizon of forecast, and the complexity of demand patterns (data with high variations in demand due to different factors).

Ivanov et al. (2018) also discusses the importance of specifying the time horizon when doing the forecast. Both Vandeput (2023a) and Ivanov et al. (2018) distinguish between three major time frames in forecasting. Short-term forecasts are often used to predict demand for the near future and are used for day-to-day operations. They tend to be most accurate, since there are the least possibilities of disruptions that might happen. Medium-range forecasts support management decisions for planning products, plants, and processes that require a more comprehensive analysis. These forecasts are critical for strategic planning and resource allocation decisions, such as hiring and training staff, purchasing raw materials, and similar. Long-term forecasts are usually used for a longer time frame and involve strategic decision-making related to new product planning, facility location planning, research and development, planning the budget for next year, etc.

Review Process/Control: This step comes as a final stage, after the first four stages are implemented. Successful implementation of forecasting method will create baseline forecast which is furtherly review by various teams. Baseline forecast refers to the initial or starting point forecast that is generated using a particular forecasting method. Different teams contribute by proposing improvements and modifications, leveraging their human expertise and insights to enrich the forecast. By examining the modifications suggested by each team, it's possible to determine which suggestions improved the forecast and which did not. This can help in a way that companies can realize which team should focus on which product, or part of the business. Ivanov et al. (2018) also points-out a frequency of reviews, like a weekly check, to ensure the forecast remains on track.

Figure 2: Forecasting process

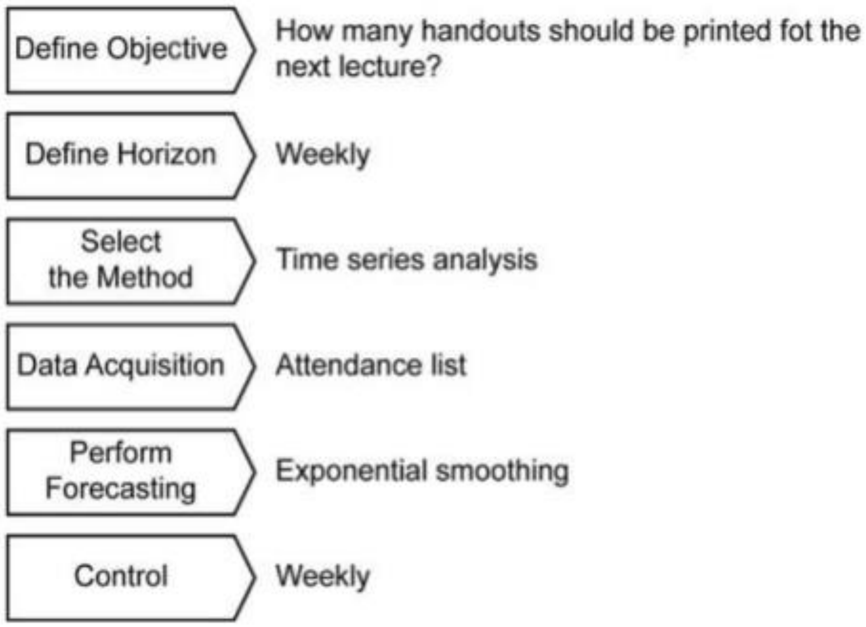

Source: Ivanov et al. (2018).

2.3 Demand forecasting methods

In general, we distinguish between two main types of forecasting methods, judgmental methods and statistical methods (Figure 3). Often used terms are quantitative method for statistical forecasting and qualitative method for judgmental forecasting.

The quantitative methods rely more on numerical data and statistical methods to forecast future events, while the qualitative methods are based on the individuals and their knowledge of the market and business environment, also can be based on market research, surveys, etc. For example, statistical forecasts will work better if a company has reliable data and a product with stable sales history (Adhikari et al., 2019). On the other hand, when a company needs to forecast demand for a product that has just been introduced, it relies on market research and industry experts' opinion, since there is no data to base the forecast on.

Figure 3: Forecasting methods

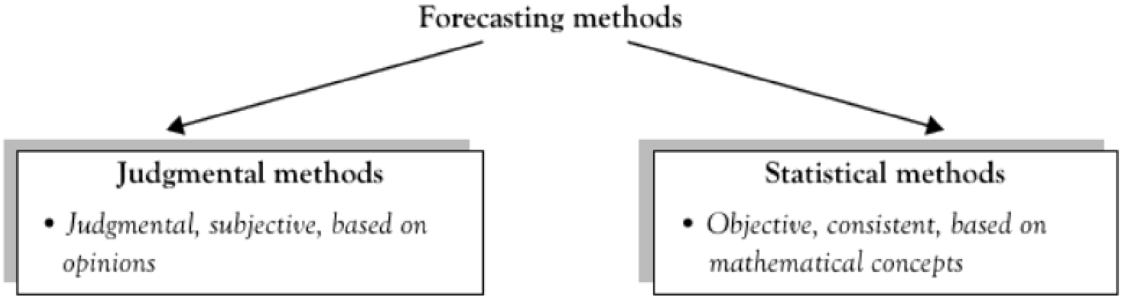

Source: Sanders (2017).

Both the scientific and judgmental methods are categorized into various sub-methods or sub-groups. Different quantitative methods exist because each method is designed in a way to be used in different situations, that is, for a specific purpose. For example, some forecasting methods are better used for long-term forecasting than short-term. Also, some methods might be better suitable for data influenced by trends or seasonality. Trend and seasonality are patterns in the data which are observed by time series methods, which are part of statistical forecasting.

Companies usually narrow down the choice on two or three different forecasting methods. This is because not all methods are appropriate for all forecasts. When selecting a forecasting method, several factors should be considered (Sanders, 2017). One factor is the amount and type of data available. Certain methods may require a larger amount of data, while others may be suitable for forecasting with limited data, depending on the specific forecast and product characteristics. Furtherly, the degree of accuracy should be considered when selecting the forecasting method. The company should consider how important is the accuracy of the forecast and how much money does it want to spend, since good forecasting methods might be costly to develop.

Some methods are better suited for long-term forecasting, and some are better for short-term forecasting, hence the time-horizon is an additional thing that should be considered in the selection process (Sanders, 2017). Additionally, it is always important to evaluate data type, and data characteristics that I described earlier, which includes trend patterns, seasonality, and cycles. Of course, selecting the method based on these criteria is not the final step.

After the selected method has been used, companies should always measure the forecast quality. Some of the most popular measures are mean absolute deviation (MAD), mean squared error (MSE), and mean absolute percentage error (MAPE) (Ivanov et al., 2018). Later, these methods are discussed in more detail.

Brandimarte (2012) explains that data plus knowledge equals the decision. So, it is important to emphasize the benefits of combining qualitative and quantitative methods, as they are not mutually exclusive. Companies should test the success of different combinations, and the combination that gives the best results should be used to conduct a forecast. According to Sanders (2017) combining the quantitative and qualitative forecasts can lead to error reduction by 14 percent on average, while that reduction can go even higher. Combining the forecasting methods is good for protection against biases, which is a big problem for qualitative forecasting methods.

2.3.1 Qualitative forecasting methods

Qualitative forecasting, also called judgmental forecasting, directly relies on human judgment, hence the name. Prior to development of the first statistical methods, judgmental forecasting was used to determine demand and accordingly prepare supply. However, today when complex and advanced statistical methods are developed, together with machine learning and artificial intelligence, and high availability of data, why is there even a need for judgmental forecasting? Well, even though data science has progressed so much that it is able to identify trends and patterns that humans simply cannot, we still rely on judgmental forecasting quite a lot.

Companies usually rely on qualitative methods when little data exists. This is typically the case for the new products that have just been introduced to the market, and in this situation, a company must rely on expert opinion, market research, and similar Ivanov et al. (2018). Of course, this is not the only case when judgmental forecast is used. For example, a company might know that the prices are going to increase next month, and if they also know that demand for that product is elastic, they should reduce their forecasts. Another example is when demand might be lower in the next period since a new competitor is entering the market, and the company should adjust the forecast as it might expect lower demand due to customers shifting. Also, when a sudden pandemic happens, as is the case with COVID-19, the company should know that the sales are going to drop significantly, while there is no way for statistical method to take that into consideration. In essence, judgmental forecasting is used mostly in the case where there are no historical data (e.g., new product forecasting) or there are disruptions (e.g., pandemic) (Sanders, 2017).

Probably the main issue with qualitative forecasting is that humans tend to make predictions about the future which might be affected by a range of decision making and cognitive biases (Dorne et al., 2008). Avoiding bias in forecasting is a challenging, and often impossible task due to its presence in both intentional and unintentional forms. Eroglu & Croxton (2010) differ between three types of cognitive biases: optimism biases, anchoring biases, and overreaction bias.

Optimism, anchoring, and overreaction bias are types of cognitive biases which occur because our judgement is influenced by different cognitive and psychological factors. The concept of optimism bias pertains to forecasters having a propensity to predominantly adjust forecasts in a positive, upward direction. Anchoring bias, on the other hand, involves forecasters being hesitant to deviate from a particular "anchor" value, which could be the initial statistical forecast for instance, even if they correctly identify the required adjustment direction. Lastly, overreaction bias occurs when forecasters make correct adjustments in judgment, but the magnitude of the adjustment is excessive, resulting in a larger error in the opposite direction (Eroglu & Croxton, 2010).

It is also important to mention intentional bias that is intentionally incorporated into the forecasting process to align with a specific, potentially driven by strategic motives, objective (Pennings et al., 2019). So, as the name suggests, different teams might intentionally bias the forecast. For example, person that works in customer service might intentionally lean to higher forecast numbers, so that every customer will be satisfied (no shortages), or marketing department that intentionally increase sales forecasts so that stakeholders would have a better felling of a new product that is being launched.

Biased forecasting process refers to making forecasts based on partial or incomplete information. One of the most common examples, which is directly related to the topic of my master thesis, is by forecasting demand by looking at historical sales rather demand. People can often face stockouts, which will lead to sales being less than the actual demand, so forecasting sales might be misleading.

There are many judgmental forecasting methods, as they are many different ways of how people make decisions. However, three most important and most common judgmental forecasting techniques can be emphasized. According to Sanders (2017), these methods, as evident in Table 1, are executive opinion, market research, and Delphi method. In Table 1, every method is explained shortly by its characteristics, strengths, and weaknesses.

Executive opinion, as shown in Table 1, is a qualitative forecasting method involving a group of managers getting together and making forecasts as a group. This method is characterized by being unstructured and relies on group discussions and opinions. It is commonly employed for forecasting sales, market trends, strategic forecasts, and new product predictions (Sanders, 2017). A potential drawback arises when a single individual's opinion, particularly if they hold a higher position, can exert significant influence over the entire forecast, even if their opinion proves to be incorrect.

Table 1: Qualitative (judgmental) forecasting methods

Method	**Characteristics**	**Strengths**	**Weaknesses**
Executive opinion	A group of managers meet and come up with a forecast	Good for strategic and new product testing	One person's opinion can dominate the forecast
Market research	Uses surveys and interviews to identify customer preferences	Good determinant of customer preferences	It can be difficult to develop a good questionnaire
Delphi method	Seek to develop a consensus among a group of experts	Excellent for forecasting long-term product demand, technological	Time consuming to develop

Source: Sanders (2017).

Another method is market research, which is a perfect way of understanding customers and the market. It is always good to do market research when trying to enter a new market or launch a new product. Conducting market research is not an easy task, and it should be carefully designed. Hence, it might be expensive since companies usually rely on outsourcing marketing experts for such research.

The Delphi method is a third qualitative forecasting technique. This method was developed back in 1950, and it is based on the fact that in general, a forecast that is coming from a group of people tends to be more accurate than those from individuals. The aim of this method is to construct a consensus forecast from a group of experts in a structured iterative manner (Hyndman & Athanasopoulos, 2018). It is based on the assumption that experts in the industry might not agree on everything, so things that they do agree on are very likely to happen (Sanders, 2017).

2.3.2 Quantitative forecasting methods

Quantitative or statistical methods, as the name states, are based on mathematics and statistics. Statistical forecasting methods have advantages over qualitative methods due to their capability to efficiently handle large volumes of data and provide objective forecasts without bias. It might be that the usage of statistical methods is to minimize the work of forecasters while providing the organization with valuable information (Adhikari et al., 2019).

As mentioned earlier, statistical methods might not be suitable for every product, but they work if a product has some sales history. Statistical methods can be generally divided into two major groups: The ones that are utilizing explanatory variables, also called predictive methods, and those that use time series data (Dorne et al., 2008). They differ in the way that time series methods are based on analyzing patterns from historical data, while predictive methods are modeled based on relationships between variables.

One of the most important things prior to statistical forecast of any kind is to make sure that clean data is being used. This means that data is accurate, complete, and consistent. One thing that caused me a lot of difficulties in my career as data analyst and data engineer is the inconsistency of the data. This can be caused by various reasons such as simple human errors like data entry mistakes, or when data is being generated through various sources, etc. It is important that the data type is consistent, otherwise it will lead to inaccurate information.

Earlier, when the forecasting process was discussed, the importance of evaluating the forecasting object and the data itself before proceeding with method selection was highlighted. When we are not sure which method to select, or even when we think we are, it is always a good idea to try different methods and compare the results. If there is a method that has the best accuracy by far, it is reasonable to only use that method. However, if several methods are giving similar results, that is, the close accuracy, it is smart to average their results into a single final forecast.

Quantitative forecasting methods can be divided into time series methods, casual (regression) methods, machine learning methods, and explanatory methods. As the name suggests, time series methods are based on time series data. Essentially, time series data is collected whenever it is necessary to observe and analyze something over a period of time (Hyndman et al., 2008). Usually, time series are associated with a graph that shows certain patterns such as trends or seasonality. Time series methods are discussed in more detail in the next chapter, along with practical examples, as I will use these extensively in the practical part of the thesis.

Regression methods are also called casual or predictive methods. They are based on the relationship in variables (Cowpertwait & Metcalfe, 2009). We can distinguish between simple linear regression, multiple linear regression, and non-linear regression. In linear regression, explanatory (independent) variable(s) x and dependent variable y are used. A dependent variable is the one that we are interested in, and we use explanatory variable to explain the dependent one (Bingham et al., 2010).

Non-linear methods are more complicated, but usually more realistic (Ivanov et al., 2018). Non-linear methods provide the most accurate explanation for various business-related measures such as cost curves, production functions, earnings functions, measures of profitability, revenues, and sales (Ritz & Streibig, 2008). These methods capture the complex relationships and patterns inherent in business data that cannot be adequately described by linear methods (Webster, 2013).

Machine learning methods use algorithms to automatically learn and adapt to patterns in data. Neural networks, random forests and support vector machines are examples of machine learning methods applied to forecasting. A notable advantage of machine learning methods is their ability to handle large and high-dimensional data sets and extract complicated patterns that can be challenging for traditional statistical methods. Additionally, these methods can adapt to changing patterns over time, making them particularly effective. However, machine learning models often require a significant amount of data for training and can be perceived as “black boxes” due to their complexity, making it difficult to interpret the reasoning behind certain predictions.

Explanatory forecasting methods incorporate external factors and underlying theories to make predictions. Econometric methods and simulation models fall into this category. Econometric methods work on a principle of incorporating economic theories and relationships to make forecasts. On the other hand, simulations methods use mathematical methods to simulate real world processes and predict outcomes. These methods allow forecasters to incorporate domain-specific knowledge and assumptions into the modeling process, making them suitable for situations where the causal relationships are well understood. However, explanatory models can have difficulties when dealing with highly complex or uncertain systems, as they are based on predefined theories and assumptions that may not fully capture the intricacies of the real world.

3 TIME SERIES FORECASTING METHODS

Time series data can be found in various contexts, and in different time-frames, such as tracking minute-by-minute stock prices, recording hourly number of customers at shop, monitoring daily arrivals at a subway, analyzing weekly sales of a product, examining monthly rise of water level, evaluating yearly sales of specific auto company. In essence, time series data is collected whenever there is a need to observe and analyze something over a period of time (Hyndman et al., 2008). Usually, people associate time series with a graph that shows specific patterns such as trend or seasonality. By looking at Figure 4, we can see some of the typical time series data from the business environment, and as we can see, there is typically a patter that can be observed just by looking at the plot, that is why the graphs are essential part of the time series. Besides regular line charts, companies often tend to use boxplots, since they are a great way to see outliers, and other statistical values such as mean, first and third quantile, and how dispersed the data is.

Figure 4: Typical time series graphs of business and economic data

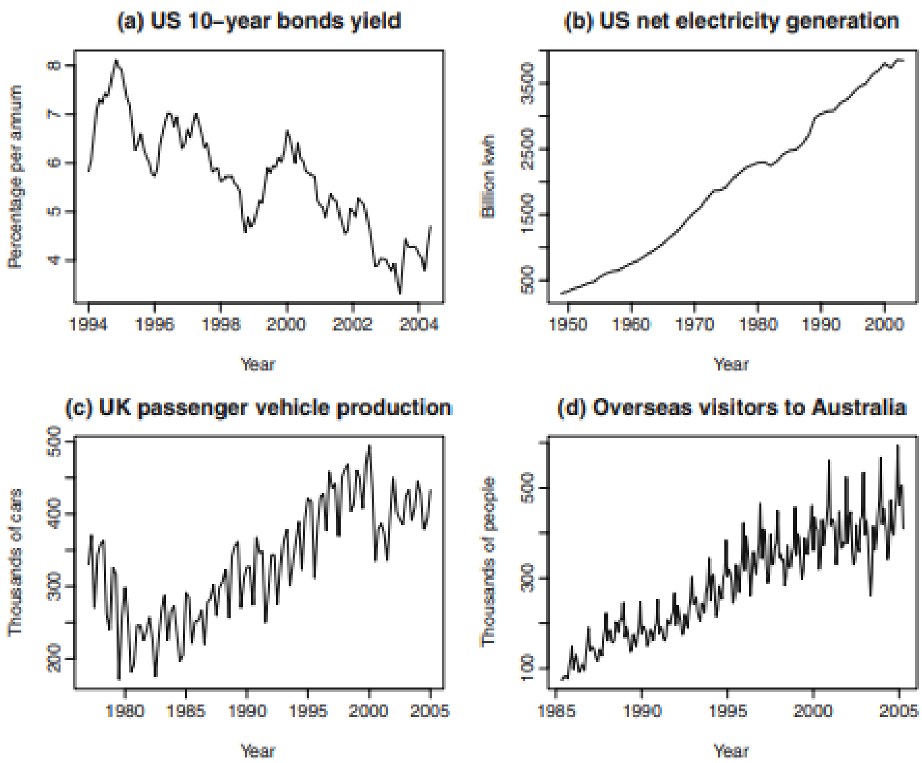

Source: Hyndman (2010).

Time series forecasting is being used to predict the future by looking at historical data and observing patterns such as trend, seasonality, and cyclicality. As can be observed in Figure 4, all of the graphs are using at least 8 years of historical data. This is the method that will be used in the practical part of this mater thesis, since after all, sales (historical) data are the subject of analysis.

Quantity of data available is often a concern when conducting the time series forecast. Dealing with this concern is not straightforward. It usually depends on the data itself, and the statistical method that is being used. Often, the more data we have, the better it is. However, some would argue that there is no point using data from more than two or three years ago because demand is changing too quickly for older data to be relevant (Hyndman & Kostenko, 2007).

According to Moroff et al. (2021), the time series is composed of four components (trend, seasonality, cyclicality, and residual component), and Vandeput (2023a) mentions level as an additional component. In Figure 5 visual difference between trend and seasonality is presented. The combination of these components can be additive or multiplicative, and combination of these components will create different demand patterns. The term additive effect describes circumstances in which the influence of one variable on the anticipated result is unrelated to the values of the other variables. On the other hand, a multiplicative effect means that the different components of a time series are not independent but are multiplied together to get the overall time series.

Figure 5: Time series components

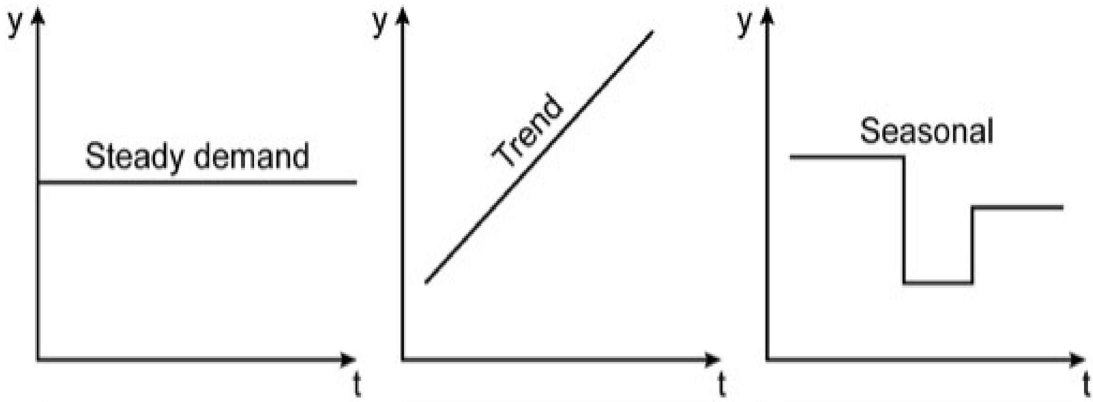

Source: Ivanov et al. (2018).

Level is the average value around which the demand varies over time (Vandeput, 2023a). It captures the long-term behavior of the variable. Usually, we use levels to establish reference points in the data, so we can understand trends and fluctuations.

The trend is a demand component which refers to consistent change in data from one period to another, as shown in Figure 5.

Seasonality is a common thing in the supply chain. It refers to seasonal products which have their high sales peaks in certain seasons and low peaks in another season, as we can see in Figure 5. Many factors, such as holidays can cause seasonality. So, Santa's suits could be an example of seasonal products that only sell around Christmas. They are regular and predictable patterns in a time series.

Seasonality, same as trend, can be additive and multiplicative. If it is additive, it means that seasonality pattern has constant amplitude through time series, while in case of multiplicativity the amplitude is proportional to the level of time series.

In Figure 6 by Vandeput (2023a) we can see the graph representing the sales, level, and trend of BMW car sales in Norway. On the x axis we observe the time-horizon of 10 years, while on the primary y axis we can see sales and level variables ranging from 200 to 1600, and on the secondary y-axis we observe the value of trend ranging from -40 to 40 units. This trend component is showing centered moving average over ten-year period, which can highlight whether the data is above or below the trend.

Figure 6 can help us better understand the demand components by looking at them visually using a real-life example. For example, we can see that the level often looks like a smoothed version of demand. We can also see how the trend line can be used to smooth out short-term fluctuations (noise) in the data, giving us a clearer understanding of the trend in sales.

Figure 6: Sales, level, and trend of BMW car sales in Norway

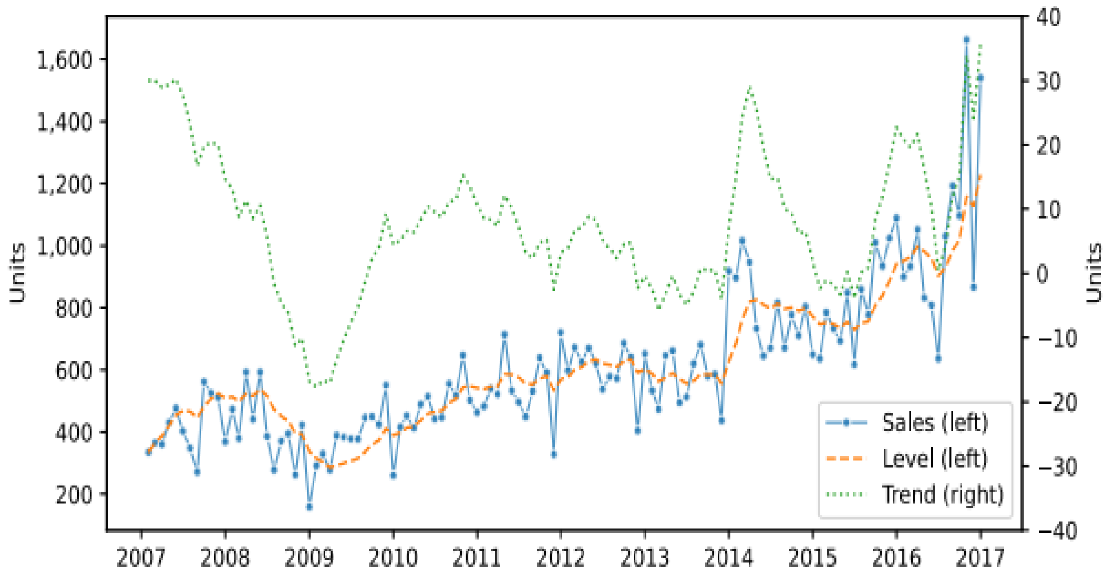

Source: Vandeput (2023a).

3.1 ARIMA methods

3.1.1 Moving average method

Moving average is probably the simplest time series forecasting method. In this method, we assume that the future demand is equal to the average demand of the period that we observed. However, it does not mean that it is necessary the mean of the entire data set, meaning that it can easily compute a mean of the part of the data (Sanders, 2017).

The name moving average comes from the fact that with every period the average value change, thus *moves* around.

The formula is:

$$F_{t+1} = \frac{D_{t-N+1} + D_{t-N+2} + \cdots + D_t}{N} \tag{1}$$

In equation (1),

F_{t+1} is forecast for next period

D_t is demand for current period

D_{t-N+1} to D_t are demands for previous N periods

N is the number of periods that we take average of

The selection of N, representing the number of periods included in the moving average calculation is a decision that balances responsiveness against stability in the forecast model. A lower N emphasizes sensitivity to recent changes, making the forecast more responsive to short-term variations. On the other hand, a higher N smooths out random fluctuations, which can be advantageous when the focus is on long-term trends.

Table 2: Moving average method with different N

t (week)	Demand	F (N_1 = 2)	F (N_2 = 3)	F (N_3 = 4)
1	27	-	-	-
2	30	-	-	-
3	33	28.5	-	-
4	31	31.5	30	-
5	29	32	31.33	30.25
6	35	30	31	30.75
7	26	32	31.67	32
8	30	30.5	30	30.25
9	33	28	30.33	30

Source: Own work.

For example, if we look at Table 2, the moving average method with various values of parameter N is tested to determine the impact of the parameter on forecasting accuracy measured by the mean squared error (MSE). In the case that we go with the two-week moving average method with N_1, the demand for t_3 (week 3) is calculated by taking the demand of t_1 and t_2 (first two weeks) and dividing them by 2: (27+30)/2. The same logic applies for N_2 (3 periods) and for N_3 (4 periods). This choice is a strategic decision based on the balance between capturing short-term changes and emphasizing stable, long-term trends in the demand forecasting process.

Figure 7 plots the 11 years long sales data. The plot contains four different moving averages where the number in front of the MA indicates the value of parameter *N* denoting the number of past years included in the moving average calculation: 3-MA, 5-MA, 7-MA, 9-MA.

Figure 7: Different moving averages applied to data

Source: Hyndman & Athanasopoulos (2018).

As mentioned, different values of *N* parameter will influence the model responsiveness. Figure 7 shows the effect of changing the order of the moving average for the residential electricity sales data. Typically, simple moving averages like these are chosen to be of an odd order (e.g., 3, 5, 7, etc.) to maintain symmetry. Symmetry ensures that the middle observation, along with an equal number of observations on both sides, contributes to the average. If an even order were selected, it would disrupt this symmetry, potentially leading to less accurate representation of the data as it might not appropriately consider both sides of the central point. (Hyndman & Athanasopoulos, 2018). Since this is quite a simple method, it is expected to have some limitations such as not being able to observe trend or seasonality.

3.1.2 Autoregressive method

In an autoregression method, the variable of interest is predicted by combining past values of the variable in a linear manner. The term "autoregression" signifies that the regression is performed using the variable's own historical data. (Rossi, 2021). Meaning, that the method works by assuming that historical values have impact on the current values. The following formula is proposed by Cowpertwait & Metcalfe (2009):

$$\boldsymbol{x_t = \alpha_0 + \alpha_1 x_{t-1} + \alpha_2 x_{t-2} + \ldots + \alpha_p x_{t-p} + \omega_t} \tag{2}$$

In equation (2) $\boldsymbol{\omega_t}$ is the white noise(error), and the α are the method parameters for p which is the order of the autoregressive part. Autoregressive methods are known to be widely flexible and to handle a wide range of different time series patterns (Hyndman & Athanasopoulos, 2018). In an autoregressive AR(p) method, p represents the order of autoregressive model, it signifies the number of past observations considered when predicting the current value. For instance, an AR(1) method ($p = 1$) employs only the immediate previous value for prediction, while an AR(2) method ($p = 2$) takes into account the two most recent values.

To decide on how many past predictions to employ, we use PACF statistical tool, to determine a correlation between the time series variable and its own lagged values. For example, determine the correlation between quantity sold in the current month versus quantity sold 5 months ago. If they are not statistically correlated, that value of 5 months ago is excluded from the method.

3.1.3 ARMA method

ARMA stands for autoregressive moving average, where p represents the order of the autoregressive part, and q represents the number of moving average terms. As you can conclude, it is combination of AR and MA method that were just discussed. Both AR and MA methods requires data to be stationary. Stationarity in data basically means that there is always the same mean during time series, variance is always the same, and there is no seasonality. Same applies for ARMA method. In ARMA method p and q parameters are chosen using criteria like AIC (Woodward et al., 2017). ARMA method has the followings formula:

$$\boldsymbol{l_t = \beta_0 + \beta_1 l_{t-1} + \varphi_1 \varepsilon_{t-1} + \varepsilon_t} \tag{3}$$

In equation (3) l_t represent the forecasted value at time t. In the formula β is the coefficient of the autoregressive term, and it is used to determine the influence of past values in predicting current/future value. β_0 is constant term that provides baseline level for the series. $\beta_1 l_{t-1}$ is the autoregressive part of the model, where β_1 is the coefficient that multiplies the value at the previous time point l_{t-1}. $\varphi_1 \varepsilon_{t-1}$ is the component from MA, with φ_1 being the coefficient of the forecast error from the previous time point ε_{t-1} which is the difference between the actual value and the forecasted value from the previous period and ε_t is error from the current time period. So, for example, we could use ARMA(2, 1) method to predict future sales based on the two previous values, that is, last two time points (AR component with $p = 2$), and one lagged forecast error, just from previous time point (MA component with $q = 1$).

3.1.4 Non-seasonal ARIMA method

ARIMA stands for Autoregressive Integrated Moving Average. As it can be concluded from the name, it is combination of autoregressive method, moving average method, and different components to capture the patterns and dynamics of a time series. In contrast to the ARMA method, ARIMA method can be used for non-stationary data where the mean and/or variance change over time.

Time series follows ARIMA(p,d,q) process where parameters inside the bracket (p,d,q) represent the orders of the autoregressive, degree of first differencing involved (I), and order of moving average part components, respectively (Hyndman & Athanasopoulos, 2018).

More specifically, I in ARIMA stans for “integrated”, signifying the process of differencing. This means that instead of forecasting demand from time-series itself, ARIMA predicts the differences in demand between different timestamps, making the data stationary. By doing this, we are getting rid of the non-stationary, and the data becomes stationary, more will be explained in the formula below.

If we set our ARIMA method to be simply (1,1,1) this will give following formula:

$$\boldsymbol{z_t = \varphi_1 z_{t-1} + \theta_1 \varepsilon_{t-1} + \varepsilon_t} \quad (4)$$

In equation (4) $\varphi_1 z_{t-1}$ is the AR part, while $\theta_1 \varepsilon_{t-1}$ is MA part, and ε_t is error in current time period. The integrated (I) part is implicitly handled by the differencing operation. The differenced series z_t is calculated as the difference between demand or sales in consecutive time points. More specifically, it is calculated as $z_t = a_{t+1} - a_t$, where a_t represents the demand or sales at period t. In the context of demand forecasting, z_t captures the change in demand from one period to the next, allowing us to model and predict these variations.

As mentioned, by transforming the series into the differences between consecutive time points, we solve the problem of the mean not being stationary, and now it has the same constant mean over time and the same variance (Rossi, 2021).

3.1.5 Seasonal ARIMA method (SARIMA)

In contrast to the ARIMA method, we can use the SARIMA method when modeling seasonal data. In order to adapt the ARIMA method to seasonality, it is important to include the seasonal term in the method: ARIMA(p, d, q) $(P, D, Q)m$.

To incorporate seasonality into the ARIMA framework, a seasonal term $(P, D, Q)m$ is added to the model where the set of parameters (P, D, Q) represents the seasonal part, and (p, d, q) is the non-seasonal part, as discussed earlier (Hyndman & Athanasopoulos, 2018). p, d, q parameters, introduced earlier, are used both in the non-seasonal and seasonal components. In the seasonal part they are combined with m**,** which stands for the number of observations per year.

Let's break down the entire SARIMA method into simple words. Autoregressive (AR) part is used to predict the value of the time-series today, based on the values from the past, integrated (I) part is used to get rid of the non-stationary (trend, variance) by using differencing to make data stationary, moving average (MA) use the error from the previous period to get more accurate results today, and finally there is a new seasonality (S) part.

So, for example, our new (S)ARIMA method can hypothetically look like $ARIMA(1,1,1)(1,1,1)4$. Seasonal part of the model $(1,1,1)4$ can be interpreted as following: One lagged seasonal value is used in autoregressive model part (P), the series is seasonally differentiated once to achieve stationarity (D), and one lagged seasonal forecast error is used in the moving average part (Q). The seasonal cycle (m) is considered to be length of 4.

To assess the need to use SARIMA, we need to determine whether the data set has seasonal characteristics. A common approach for this is to graphically represent the dataset and gain a visual understanding of its characteristics.

3.2 Exponential smoothing

Exponential smoothing encompasses weighted averages of past data, as the weight (importance) decreases exponentially as the observation (data point) gets older, so that new observations have more weight (Hyndman & Athanasopoulos, 2018). As it is described later, we differ between three types of exponential smoothing, based on the company's needs and what kind of data we have.

3.2.1 Simple exponential smoothing

The simple exponential smoothing is suitable for forecasting data that has no clear trend or seasonality. Level (average value) is the only pattern that this method is capable of learning. Up until this point, the simple exponential smoothing might sound the same as weighted moving average. While both methods involve averaging previous data points, they differ in the way the weights are assigned. When it comes to weighted moving average, we must specify the weight that is assigned to each time point, while in the case of simple exponential smoothing, the weight that is assigned to every time point exponentially decreases over time so that most recent observation has the highest weight.

The logic behind exponential smoothing is that a method always learns from the most recent observations, and using it to update the last forecast (Vandeput, 2021), which is captured in the following mathematical formula:

$$\boldsymbol{f_t = \alpha d_{t-1} + (1-\alpha) f_{t-1}} \tag{5}$$

In equation (5) α is the smoothing parameter that determines the importance that has been allocated to the most recent data point, where α values are between 0 and 1.

The higher the α, the more importance will be allocated to the latest data point, that is, latest observation. This means that the method will learn fast, and it will be reactive to the changes. On the other hand, it will be more sensitive to outliers and noises, and it will not be as smooth. While lower α will give less sensitive, smoother method, which will also be less responsive to the changes (Vandeput, 2021).

In practice, α is between 0.05 and 0.5, depending on the forecast itself, and of course human judgement. When α is higher than 0.5, it suggests that the method is giving low importance to the historical data, making the method responsive, but sensitive to noise and outliers.

Besides the ability to exponentially assign weight to the time periods, simple exponential smoothing is same as the moving average method. That means that that they share the same limitations. Limitations of this method are that it does not project or observe trends, or any seasonal patterns.

3.2.2 Double exponential smoothing

The double exponential smoothing method is often called Holt's trend method, since he is the one that first introduced the trend adjustment to exponential smoothing in 1957. The simple exponential smoothing gave us a great starting point to further develop more complex and accurate methods.

The need for the development of a double exponential smoothing method arises from one of the major issues that simple exponential smoothing has, and that is the inability to identify and project trend. So, this method is used when we believe that trend already exists (Ivanov et al., 2018). However, this method does not recognize any seasonality.

Before, I mentioned the α that determines the weight, which determines the responsiveness of the method, in this method, the β is being introduced which determine how quickly the method picks up the trend in history (Adhikari et al., 2019). So, our α and β are two smoothing factors, and they are both always between 0 and 1. To create demand forecast using this method, we need to firstly evaluate the level and the trend, since the forecast is simply equal to the level plus trend, that is a_t in period t plus b_t in period t, as we can see in the formula below. As it will be evident, the formulas are similar to those we have for simple exponential smoothing, except that now a trend is also included. Level a_t in period t can be estimated as:

$$\boldsymbol{a_t = \alpha d_t + (1 - \alpha)(a_{t-1} + b_{t-1})} \tag{6}$$

$a_{t-1} + b_{t-1}$ in (6) is the previous level estimation increased by the trend, and $(1 - \alpha)$ will give the importance (weight) to this estimation.

Trend b_t can be estimated using the same logic:

$$\boldsymbol{b_t = \beta(a_t - a_{t-1}) + (1 - \beta)b_{t-1}} \tag{7}$$

Now, the forecast for the next period can be simply shown as:

$$\boldsymbol{f_{t+1} = a_t + b_t} \tag{8}$$

Another thing important to mention is the double smoothing with the damped trend. Here the idea is to fix the assumption that trend goes on forever. So, the new factor is added to the method: damping factor phi(φ). Using this factor, the trend will be exponentially decreased over time. It is particularly useful for long-term forecasts because it mitigates the assumption that the current trend will continue indefinitely.

3.2.3 Triple exponential smoothing (Holt-Winters)

The double exponential smoothing was updated by Holt(1957) and Winters in (1960) to capture seasonality, so this methods is often called Holt-Winters method (Hyndman & Athanasopoulos, 2018). So, as the name suggests, this method has three smoothers. We already worked with the first two, one for level (alpha) and another for trend (beta), the new one is the seasonal component which will be marked with gamma (γ). The seasonal factor in this method can be additive or multiplicative. If it is multiplicative, following formula proposed by Vandeput (2021) can be used:

$$\boldsymbol{f_{t+1} = (a_t + \varphi b_t) s_{t+1-p}} \tag{9}$$

In simper words, this method can be presented as:

Next period forecast = (Level + Trend) * Season.

As seasonality is included in mode, demand factors need to be updated in the next way:

$$\boldsymbol{a_t = \alpha \frac{d_t}{s_{t-p}} + (1-\alpha)(a_{t-1} + \varphi b_{t-1})} \tag{10}$$

$$\boldsymbol{b_t = \beta(a_t - a_{t-1}) + (1-\beta)\, \varphi b_{t-1}} \tag{11}$$

$$\boldsymbol{s_t = \gamma \frac{d_t}{a_t} + (1-\gamma) s_{t-p}} \tag{12}$$

Same as α and β smoothing parameters, γ is there to determine the importance (weight) that has is given to every observation. Gamma is also between 0 and 1, and in practice γ should be lower than 0.3, since the seasonality should not change too much over the years, it is usually quite constant. p in s_{t-p} is used to represent the frequency of seasonality. For example, in quarterly data we should use $p = 4$ and for monthly data $p = 12$ (Hyndman & Athanasopoulos, 2018).

One of the problems with multiplicative seasonality is that the method is prone to result in mathematical errors when the demand for the certain product is very low, in this case the method will tend to overreact. This will also happen if the seasonality factors are too close to 0. Hence, multiplicative seasonality is not suitable for every product since it will give false results. Multiplicative seasonality is appropriate when the seasonality effect increases or decreases proportionally with the overall trend.

In the situation where demand for a product has periods with very low volumes it is good idea to use additive seasonality. In this case formula will be adjusted in that matter so it is:

$$\boldsymbol{f_{t+1} = a_t + \varphi b_t + s_{t+1-p}} \tag{13}$$

In (13), next period forecast = Level + Trend + Season.

Multiplicative seasonality is expressed in percentages, while additive is expressed in numbers. Seasonality should not impact the amount forecasted. This means that if we were to forecast demand using double exponential smoothing and got a certain amount, same forecast should be done using triple exponential smoothing, however the distribution of amounts should be different. In simpler words, if we sell 1000 units more in July than usual, we should sell 1000 units less during the rest of the year.

3.3 Statistical measures of forecasting accuracy

Different forecasting methods will give different results, and this is actually a good thing, since we can compare those results and see which forecasting method is the best one to use, for a specific example. There are several different methods that can be used to estimate forecasting accuracy (quality). We should start with computing the forecast error:

$$\boldsymbol{e_t = d_t - f_t} \tag{14}$$

This formula (14) simply shows how forecast error for the period (e_t) is a difference between actual value (d_t) and forecasted value (f_t). So, if the forecast is higher than it should be, the error will be positive, and vice versa.

Before going into different ways on how to measure the forecast accuracy, we should first understand what is the "fit" of the method (Sanders, 2017). Fit of the method basically means that we are taking the historical data and testing the different methods on those data. This is an important step, because, as discussed, not all historical data are the same, some have trends, some seasonality, some both, and similar. However, the problem is that the “fit” of the method only shows us how well it performed in history, and it is not giving the same guarantee for the future. Hence, it is recommended to test the method on the different data from the ones that we used in method creation, as well as to keep monitoring the performance of the method that is being used (Sanders, 2017).

Here, we came across terms, test set and training set. On the training set, we train our method based on historical data. Then we use the test set to evaluate the performance of the method. This is usually done using the same data set, which we separate into two parts. Part of the data on which we test the method is usually around 20% of the entire data set, and the rest goes on training set (Hyndman & Athanasopoulos, 2018).

We can also differentiate between standard error measurement and relative error measurement (Sanders, 2017). Standard error measurements are scale dependent and are expressed in numbers, meaning that results for two different objects are hardly comparable with each other, since for example, something is sold in hundreds and something in thousands pieces. On the other hand, relative error measurements are expressed in percentage and are not scale dependent, so we can compare different products using same percentages.

3.3.1 Mean Absolute Error (MAE)

As the name implies, the MAE is a means of absolute errors. According to Hyndman et al. (2008) the MAE can be formulated easily as:

$$MAE = \frac{1}{n}\sum_{n} |e_t| \tag{15}$$

e_t in (15) is the same error term that was mentioned above. Results are expressed as absolute value, meaning that if the MAE is equal to 20, it is not known if it is 20 more or 20 less. The main point to understand about these errors is that they are directly linked to the scale of the data. Therefore, MAE is referred to as scale dependent forecasting error. If, for instance, forecasted/actual value represents sales volume in kilograms, then the error term, e_t, will also be measured in kilograms. This scale-dependency poses a limitation for accuracy measures based solely on e_t, making it unsuitable for comparing different series that operate on diverse scales (Hyndman et al., 2008). This hold for both MAE and MSE that is mentioned below.

3.3.2 Mean Squared Error (MSE) and Root Mean Squared Error (RMSE)

MSE has similar approach as MAE, the difference is that MSE uses average of squared errors, not absolute errors (Wang & Bovik, 2009). This can be observed in the formula below:

$$MSE = \frac{1}{n}\sum_{n} e_t^2 \tag{16}$$

The squared differences can magnify the impact of larger errors, making it more challenging to intuitively grasp the significance of the metric. For instance, the difference between 3 squared and 5 squared (9 vs. 25) appears substantial, and this can make the overall MSE less intuitive for interpretation (Sanders, 2017). However, the bright side is that is gives more weight on the larger errors than on the smaller errors.

To address interpretability, people will often use root mean squared error (RMSE) which is basically squared root of MSE. The benefit of using RMSE over MSE lies in its ability to provide a more interpretable measure of the standard deviation of errors, as it is presented in the same scale as the original data.

$$RMSE = \sqrt{MSE} \tag{17}$$

3.3.3 Mean Absolute Percentage Error (MAPE)

This is one of the most commonly used methods to measure forecasting accuracy. Percentage errors offer the benefit of being unaffected by scale, making them a popular choice for comparing forecast performance across various data sets. So, difference between MAE and MAPE is that MAPE use percentage errors, so it can compare performance regardless of the scale. According to Hyndman et al. (2008) MAPE can be computed as:

$$MAPE = mean(|p_t|) \tag{18}$$

In formula (18) the percentage error is given by $p_t = 100\frac{e_t}{y_t}$.

Percentage errors suffer from drawbacks when applied in certain scenarios. They can become infinite or undefined when any observation in the test set (y_t) equals zero, and they yield extreme values when yt approaches zero. Furthermore, an often-neglected issue with percentage errors is their assumption of a quantity-based scale, and the assumption that the unit of measurement possesses a meaningful zero point.

For example, assessing the accuracy of temperature forecasts using percentage errors is inappropriate since temperature has an arbitrary (not meaningful) zero point, whether in Fahrenheit or Celsius scales. (Hyndman et al., 2014). A low forecast, such as 0, can result in a percentage error of 100%, while excessively high forecasts are not limited to a specific percentage error. So, if you had demand of 100 units and forecast of 1000, the error will be 900%. General rule is that MAPE should not be used if the data is not cleaned properly, since it can be misleading for outliers, intermittent zeros, and low-demand (Sanders, 2017).

The answer to a question which error measurement to pick, there is no straight answer. If we have a case where forecast one has a lower MAE and RMSE than the forecast two, than we should, of course, pick the forecast one. But often we have a scenario when we have conflicting results from different measures (Sanders, 2017).

Some researchers would argue against the use of RMSE and MSE in forecast accuracy evaluation, precisely because of their sensitivity to outliers. Outliers are extreme data points that deviate significantly from the rest of the data set, and they can disproportionately impact the performance assessment of RMSE and MSE, leading to potentially misleading results. In contrast, MAE is more robust in the presence of outliers, as it does not square the errors. Hence, it makes it less affected by extreme values, providing a more balanced and reliable assessment of forecast accuracy, particularly in datasets where outliers are common (Hyndman & Koehler, 2006).

3.4 Comparing the different forecasting methods and their accuracy

In this chapter, the same data set will be used to test previously introduced time series forecasting methods. The demand forecast generated using these method will be evaluated using statistical measures of forecasting accuracy that were introduced in the previous chapter. This is done so we can gain better understanding of how different forecasting methods perform, as well as behavior of different statistical measures of forecasting accuracy.

The provided data set in Figure 8 captures a monthly sales data covering a period of three years. The x-axis represents the date, with time progressing in monthly increments. The y-axis represents the sales values for the month in question. The line in the graph shows the sales trend and seasonal fluctuations. The chosen granularity of the data set is monthly to provide a focused view of sales dynamics over time.

Figure 8: Sample sales data with seasonality and trend

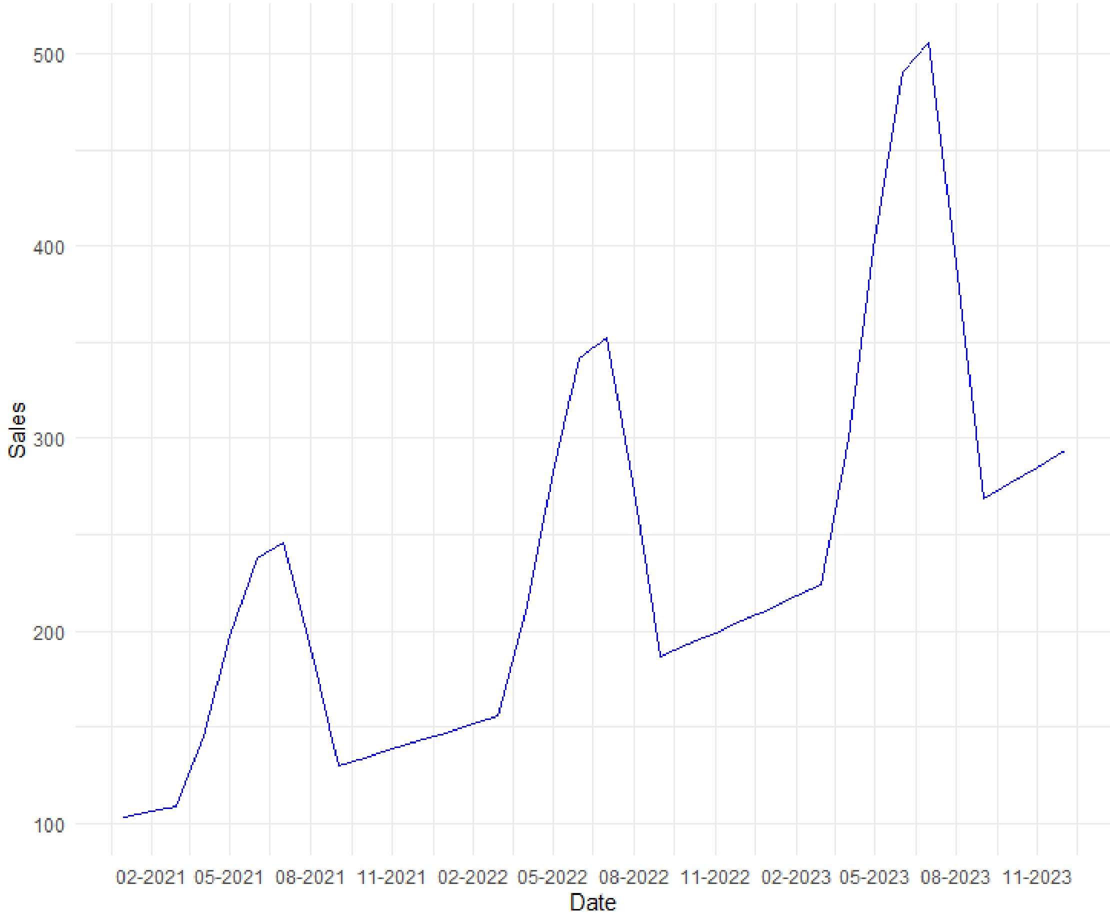

Source: Own work.

Now that we understand the data, various time series forecasts can be performed on the data set. Each time series forecasting method, which can be seen in Figure 9, is based on the most appropriate parameters for that method. To achieve this, a 'for' loop was implemented in R software that iterates through different sets of parameters and returns the one that gives the lowest MAE value. In this way, I made sure that each time series forecasting method is used in the best possible way.

Figure 9 illustrates the monthly forecast for the year 2024 using different forecasting methods applied to the dataset from the previous three years. The x-axis represents a one-year date period in month-year format (mm-yy). Each colored line corresponds to a specific forecasting method, with the dots marking the forecasted values for each period. A red and black dashed line represents actual sales and serves as a benchmark for assessing forecast accuracy. In this figure, SES stands for single exponential smoothing, DES for double exponential smoothing and TES for triple exponential smoothing, and the other acronyms are known to us.

Figure 9: Comparison of different forecasting methods

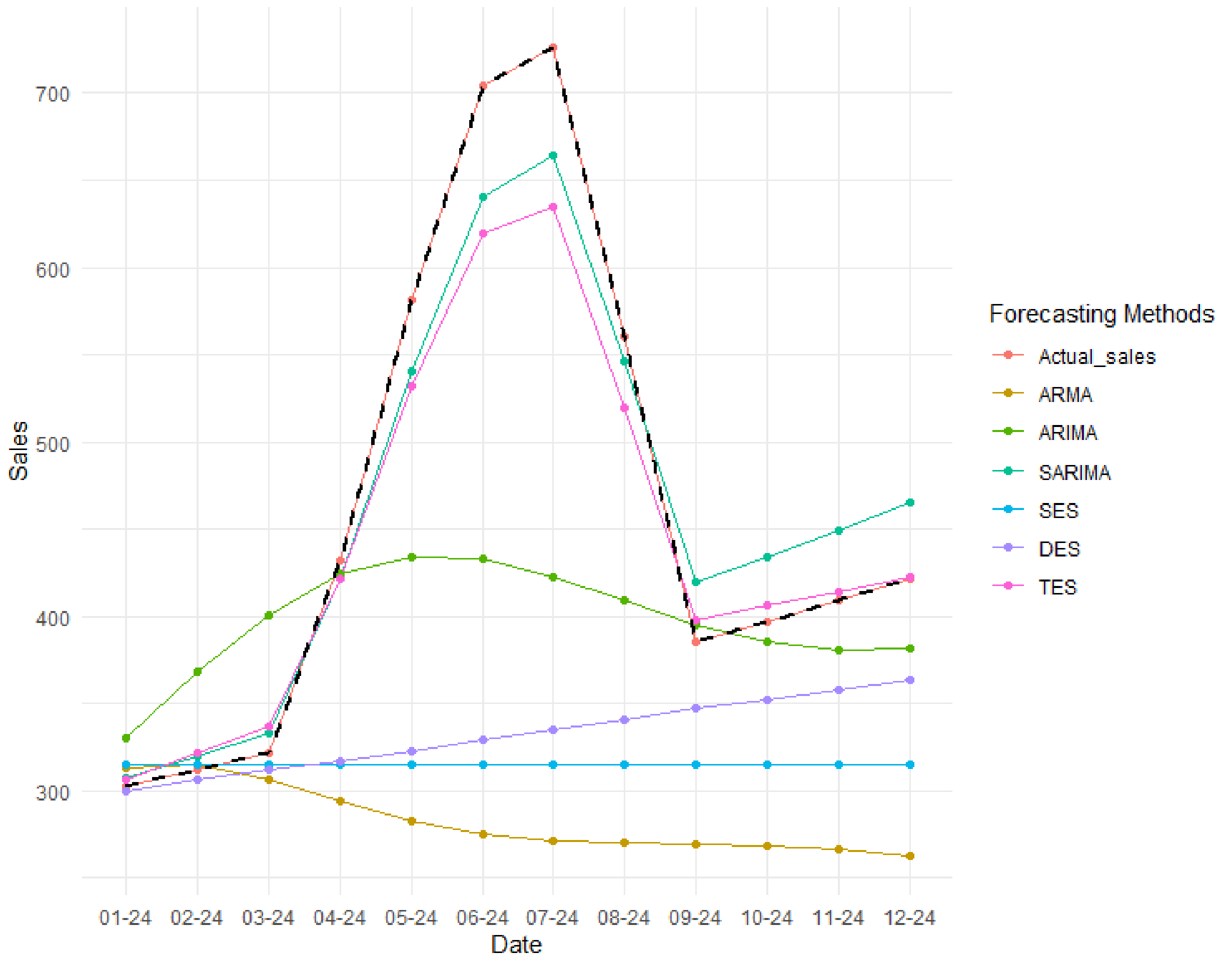

Source: Own work.

Reflecting on what we have learned in Chapter 3, we can conclude that the behavior of forecasting methods from Figure 9 is more or less expected with the respect to the data set shown in Figure 8. Since the data set that is subject to both trend and seasonality was used, forecasting methods that can handle both are expected to perform best. For this reason, SARIMA and TES provided the most accurate results in this example. They are the only time series forecasting methods that can capture both trend and seasonality.

ARIMA models are more flexible compared to ARMA models and that is why it show much better results than ARMA. ARMA performed poorly with any combination of parameters, since it is suitable only for stationary time series data, without trend and/or seasonality. SES, when applied with any set of parameters gave a straight line. This is because its only attribute is to exponentially decrease the impact of past observation as they get older. DES accounts for trend which could explain why it follows the actual sales trend more closely than SES, but it does not account for seasonality, which is causing inaccurate forecast.

Now let's evaluate various forecast accuracy measurements and how they behave by looking at Table 3. We can see which demand forecasting method performed best by looking at Table 3 as well as by looking at Figure 9. However, Table 3 shows us how different measures of forecast accuracy behave. For example, if we look at the MAE value for any given forecasting method, it is not easy to interpret it and draw insights from it because we do not see the magnitude of the data. However, if we had information about forecast versus actual sales for each period, the MAE would be more useful.

MAPE is a better metric if we have no knowledge of the data. MSE do not tell us much either. Both MSE and RMSE are more useful when comparing models to each other rather than using them as standalone metrics. When evaluating a model's performance, lower MSE and RMSE values indicate a better fit to the data, assuming all other factors are equal.

Table 3: Different error measurements

Period	MAE	MSE	RMSE	MAPE
ARMA	182	54975	234	33,58%
ARIMA	94	18565	136	17,30%
SARIMA	31	1343	37	6,25%
SES	150	41406	203	26,88%
DES	130	35940	190	22,57%
TES	27	1682	41	4,86%

Source: Own work.

4 COLLECTION AND APPROXIMATION OF DEMAND FROM THE SALES DATA WITH STOCKOUTS

This chapter serves as the theoretical basis for the practical part in chapters 5 and 6. In this chapter I describe the relevance of the appropriate data preparation for forecasting from sales data with stockouts. By the end of this chapter, the reader will be able to understand the magnitude of this problem and issues that comes with it. I believe that the problem of stock-outs in data is still very much neglected as it is rarely mentioned in forecasting books, articles or journals dealing with forecasting issues.

Before discussing the quantitative models used to approximate demand during stockouts, I first discuss how one can avoid these data gaps in the first place. Here, different methods for capturing demand during shortages are discussed. Finally, I evaluate different quantitative methods for estimating demand during shortages. Various methods are discussed here, such as the average of the last known periods, interpolation and imputation of time series. Here, I only discuss their advantages and challenges, but not yet their implementation.

Later, in Chapter 5.2, these methods are implemented, and in Chapter 6.2 the results are evaluated and compared.

4.1 Problem of stockouts in sales data

At the beginning of this chapter, the difference between sales and demand should be emphasized. Sales are equal to how much the company has sold, and the demand equals to how much the company is supposed to sell. Demand can be higher than sales, but it cannot be lower. Demand has an unconstrained point of view as it defines what, when, and how much of a specific product the customer wants. (Vandeput, 2020).

New let's understand stockouts. Stockouts can be defined as a situation where the inventory is depleted to zero, or in other words there are no more units of that product available for sales. Consequently, we cannot just know whether one or more additional units were demanded. Relying solely on sales data without considering stockouts could lead to a misleading forecast. We need to understand that actual demand might have been higher than sales, especially during stockout events (Bell, 2000).

A common practice in the company is to forecast demand based on the historical sales data. If that item was always in stock, that means that all of the demand was satisfied, and that the sales matched the actual demand (Wecker, 1978). If this is always the case, there would not be a need to address this issue since the historical sales data would provide perfect information about the demand. However, quite often companies experience stockouts. This leads to demand being higher than the actual sales. That is why in this case, sales are not the right indicator of demand. So, if we were to take those sales data, without any adjustments, we would get a false forecast.

Stockouts, besides having a negative influence on the demand and thus companies' profit, also have an influence on customers. In the research paper by Kim et al. (2019), it is discussed how stockouts affect the customers, and conclusion is that the majority of them would either switch brands, go to another store, or turn to the product substitute.

An additional aspect that comes into play when stockouts are faced is the effect of substitution and cannibalization. Substitution can be a situation when, running out of stock for one product, might result in higher sales for another product. Cannibalization relates to situations where we stimulate sales for one product (price, advertising, etc.) which result in lower sales for other products. This effect also influences sales data in two ways: it can lead to lost customers and thus a reduction in potential sales, or customers may postpone their purchases, which can shift demand patterns. Both aspects create an issue of distinguishing between genuine demand changes and those influenced by substitution or cannibalization, which can disrupt accurate demand forecasts for individual products (Vandeput, 2023b).

Figure 10: Forecast based on demand vs. sales data

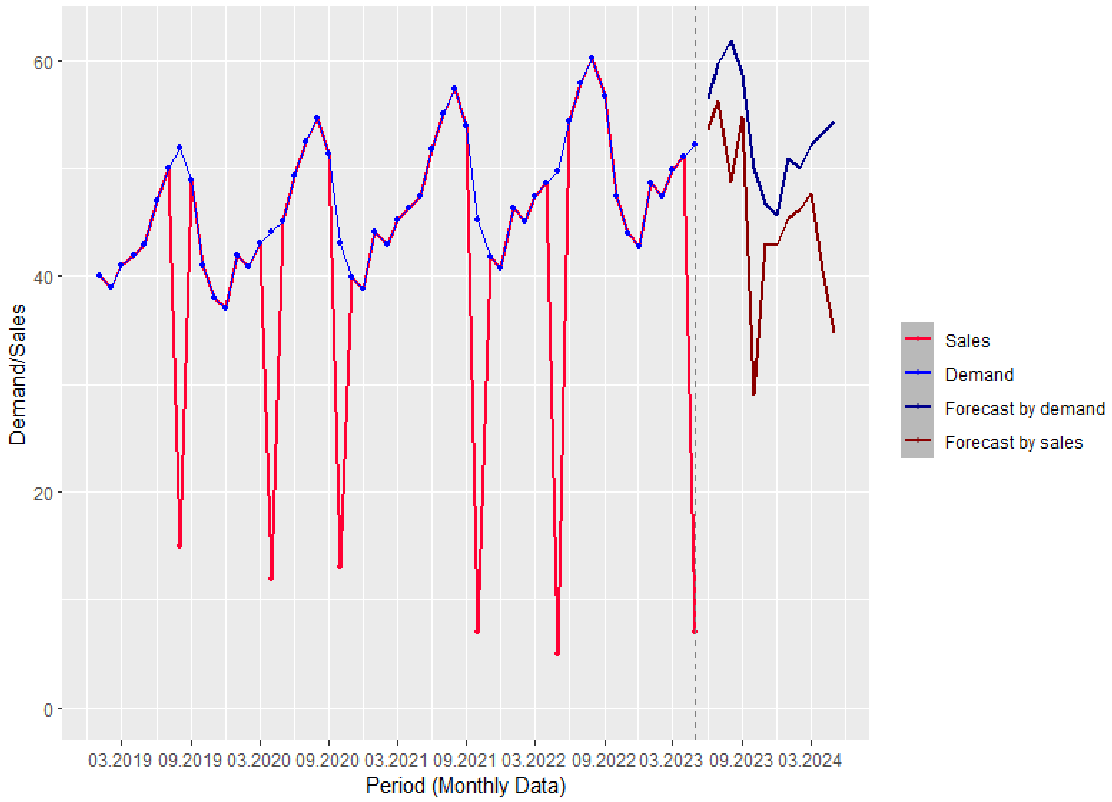

Source: Own work.

In Figure 10 an example is illustrated, showing monthly sales data with several periods of stockouts over the course of 4 years. Periods where sales are well below normal levels are indicators that there were stockouts on some days of these months. These large swings represent large deviations from the characteristic demand pattern in which positive trend and in particular seasonality in demand is observed. A possible approach to generate the demand data for these monthly outliers is to fit the demand data to the observed pattern.

To show how the unadjusted sales data affects the forecasting accuracy, I have added the forecast for the following 1 year (period after vertical gray line). We see that the forecast based on the demand data follows the characteristic demand pattern, while forecasting based on unadjusted data leads to underestimation of the demand and erratic swings. Hence, if we were to do forecast based on sales data without any adjustments, we would most likely face stockouts again in the future. This can lead to a series of issues such as profit loss, customer loss, bad word-of-mouth, etc. Therefore, this is an issue that every company needs to take seriously.

4.2 Methods of demand collection

Now, we can understand why is important to adjust the sales data when stockouts are present. Prior to discussing the mathematical methods to approximate demand during the periods of stockouts, let's introduce several methods used to collect the demand in such occurrences.

To avoid the problem of missing a realistic demand data, Vandeput (2023a) suggests how to collect the demand in case of shortages, where he discusses several methods and points out that 100% realistic collection of demand data is almost impossible. I can point out to the following methods that can be used to improve our demand data: order collection, shortage-censoring, and customer collaboration.

The first method is order collection. The purpose of this method is to trace orders even if we ran out of stock. In case of shortages, order management system can be used to handle incoming orders and client requests. In this way, collection of demand will continue even if the outbound shipments drop. This method focuses on tracking different types of orders affected by stockouts, such as open, duplicate, canceled, substitution, and uncollected orders. Here we can face challenge such as potential online customers leaving due to unavailable stock, since this is particularly hard to collect and quantify. DeLamater et al. (2018) discuss similar solution called point of sale systems and inventory management which offer features that allow businesses to track lost sales due to stockouts. Whenever a customer asks for an item that is not in stock, employees can record it in the system, providing a clearer picture of unmet demand.

The second method is shortage-censoring. The purpose of this method is to modify forecasting models to account for stockouts. This method is based on tracking inventory levels and sales data over time, so we know in which periods we had enough supply to meet demand. Then, we have to teach the method that the sales are down due to no stock, not due to low demand. One challenge is that many modern forecasting engines neglect inventory data, even when available, so it might be necessary to create a model that will account for this.

Third method is customer collaboration. The purpose of this method is try capturing demand by direct customer collaboration. This can be achieved through pre-orders or reservations for specific products, enabling the company to predict the required stock. Also, the company can provide a waitlist and notification system that allows customers to express their interest. Asur & Huberman (2010) also discuss considering data from external sources such as social media mentions, online forums, or web traffic analytics, online communication channels can be beneficial to get insight of the actual demand for the period. Also, directly asking customers about products they could not find or would have purchased can provide insights into demand during stockouts (Parasuraman et al., 1988).

As mentioned earlier, these methods may not be able to capture the entirety of the unfulfilled demand, but they should be able to capture a significant portion. An additional problem that one can face is assessing the accuracy of the collected data as there is no direct way to assess to what extent and how accurately we have captured the true demand. Indirectly, the accuracy can be assessed by using the collected demand data in forecasting demand. If adjusted demand data leads to more accurate future forecasts, then it is reasonable to assume the collection was appropriate. If the forecasting results are unsatisfactory, it could indicate a need for improvement in our forecasting methodology or the possibility that the methods that we are using to collect demand are not successful.

4.3 Methods to approximate demand

I have just discussed methods of collecting demand when facing stockouts. If we were to use these methods to collect the demand when stockouts occur, it would be easy to insert into sales data those captured demand values instead of actual sales, which were zero due to stockouts. By applying these methods, companies can capture the majority of demand, but mostly likely never 100%.

Unfortunately, many companies still do not use these methods to collect demand, and are left with sales data which have gaps inside, and cannot be used to accurately forecast the demand. In this case, we should use some mathematical methods to approximate the demand in the period of stockout.

An important practice for the company to implement involves monitoring inventory levels and incorporating them into sales data analysis and forecasting. This practice holds significant importance because relying on manual detection of shortages is not a good practice. Of course, some instances of stockouts can be identified by just looking at the data, but we can easily miss small differences between sales and the actual demand, particularly if both are highly variable.

4.3.1 Optimized average of last known periods

Calculating the average of the last known periods involves using the mean of the data points prior to a gap to estimate missing values. This method is often used as a quick-fix, and it is suitable for simple data set that are not subject to trend and/or seasonality. To enhance this method for datasets with such characteristics, a deeper analysis of demand patterns is essential. For time series with no random fluctuations and no evident trend or seasonality, a simple averaging approach may be sufficient. However, for data with seasonal trends, such as a spike in product sales during summer, an optimization of method would be required. This optimization might involve averaging sales from seasonal months.

For example, if we work with the data set that is subject to seasonality during summer months, to estimate a stockout in June, one should not simply take average of months prior to June, since those might be not influenced by seasonality. Furthermore, when trend is evident in the data set, it is smart to take average of only few recent periods around stockouts, giving less weight to past observations. This approach mitigates the issue of relying on outdated patterns that may no longer be relevant.

4.3.2 Interpolation

Interpolation is a mathematical technique that utilizes a given set of known independent and dependent values to estimate the values of the dependent variable at other points within the range of the given data. It involves constructing a continuous curve or polynomial that passes through the known data points. However, it’s important to note that if the technique is employed to estimate values beyond the range of the known data points, it is referred to as extrapolation (Murphey, 2022). Ali et al. (2018) defined interpolation as a mathematical function that predicts the values at positions according to the observed values in other locations. The word interpolation comes from Latin inter (between) and polare (to polish) which means to compute new values that lie between certain given values (Salomon, 2011).

Interpolation is a commonly used tool to approximate missing records in the time-series. It is usually used in cases where we have gap of one or more observations, which are stockouts in our example. These gaps should be followed by a long unbroken sequence of known values right before and after the gap.

Besides approximation of missing values, interpolation is also commonly used when it is necessary to replace outliers because they may introduce errors in forecasts if left unaddressed. For instance, when atypical events such as extreme weather conditions or strikes occur, they may disrupt the usual demand patterns for a product. In these scenarios, interpolation is a valuable tool to approximate the demand for the affected periods, thereby producing a more accurate and usable data set for forecasting needs (Brubacher & Wilson, 1976).

We can distinguish between several interpolation methods, each varying in complexity and capability. Here I will explain the basics of nearest neighbor, linear interpolation, spline interpolation, and polynomial interpolation methods.

Probably the simplest method is the nearest neighbor, also known as proximal interpolation. This is a straightforward method which fills in missing values by adopting the value of the closest known data point. The data interpolated using this method tends to be discontinuous, or in other words they can jump suddenly from one data point to another without gradual transitions in between. It will usually not give good results if, for example, we have stockouts occurring multiple days in a row (Lepot et al., 2017).

Another method that is frequently used is linear interpolation. Linear interpolation is a method of first order that draws a straight line between every two successive data points in the input signal (Miklos, 2004). It uses linear polynomials to create new data points within the set of existing data points. The fundamental assumption of this method is that the change (relationship) between two data points is linear, which follows a straight-line path.

Next very popular interpolation technique is cubic spline interpolation. Splines can be defined as set of polynomials of degree k that are smoothly connected at certain data points. At each data point, two polynomials connect, and their first derivatives (tangent vectors) have the same values (Salomon, 2011). The most commonly used one is cubic interpolation, which uses cubic polynomials (degree of three). Through cubic spline interpolation, a set of distinct cubic polynomials are tailored to fit between each pair of data points, ensuring that the resulting curve is not only continuous but also exhibits a smooth appearance. These cubic splines are then employed to calculate the rates of change and the overall change across a given interval (McKinley & Levine, 1998).

This method usually yields good results, since it is the most flexible, avoids oscillation and overfitting, and can be used for large data sets. The flexibility inherent in cubic spline interpolation ensures the resulting curve remains smooth throughout its length. Not only is it smooth, but its rate of change is also smooth, meaning there are no sudden jumps or sharp bends. This gives us a nicely flowing curve that fits our data points well (Faires & Burden, 2016). This method has the flexibility to fit a wide variety of data shapes. They can capture nonlinear trends more effectively than linear interpolation.

Another popular interpolation technique is polynomial interpolation. It involves creating a curve that passes through a given set of points. By calculating the coefficients of the polynomial equation, we can represent the points as a smooth curve (Salomon, 2011). Using this polynomial for approximation within the interval of know datapoints is called polynomial interpolation (Faires & Burden, 2016). This method fits single polynomial function that goes through data set. This means that classic polynomial interpolation is "global" fit. Because of this, it might be less flexible, prone to overfitting, and it might not be able to capture variations. However, we can adjust this by having a higher degree of polynomial, which will be discussed more in chapter 5.2 where I will apply the method.

4.3.3 Time series imputation

Time series imputation is a statistical technique used to estimate missing or incomplete values (such as stockouts) within a time series. Unlike imputation in cross-sectional data, time series imputation takes into account the temporal dependence structure of the data. The objective is to fill in missing data points in a way that is consistent with the time-dependent nature of the data set, which often includes considerations of trends, seasonality, and cyclic behavior (Hyndman & Athanasopoulos, 2018).

The challenge in imputing time series arises from the need to consider the temporal dependencies rather than the correlations between attributes that are usually considered by standard imputation algorithms. Temporal dependencies require an algorithm to predict or reconstruct missing data points based on the values before and after the gap in a way that matches the identified patterns in the data (Moritz & Bartz-Beielstein, 2017).

Essentially, imputation is about identifying the missing values, understanding the characteristics of the data set and then selecting and applying a suitable imputation method. For example, in a data set with multiple stockouts over three years, each missing period would be treated individually. The imputation for the first stockout would be based on the clean data prior to the event. Successive stockouts would be based on the previously imputed values to maintain the consistency and coherence of the entire time series.

Various statistical methods can be used for imputation. These methods often rely on the principles of forecasting, as they project the observed data of the past into the missing periods, assuming that the patterns of the past persist. For example, let's say we have sales data for the last three years and there have been four stock-outs in that period. Demand in the first stockout period would be estimated by applying certain forecasting methods to the data points prior to that stockout. Then the demand in the second stockout period would be estimated by taking all the data prior to that period, including the estimated value for the first stockout period, and we would do this for all four stockout periods in the data set.

State space models are an important component of the imputation of time series. They are mathematical constructs used in time series analysis to represent complex data with underlying trends, seasonality and cyclical behavior (Shumway & Stoffer, 2017). They work by decomposing a time series into its constituent parts, which are then analyzed in a unified framework.

Kalman smoothing is a method that complements state space models by improving the estimate of the state of the time series at each point. It utilizes all available data and adjusts previous estimates with new information as it becomes available to improve accuracy (Barratt & Stephen, 2020). Together, they provide a powerful approach to time series imputation, especially for non-stationary data where simple methods fail. Kalman smoothing can refine the imputation of missing values through its retrospective analysis, taking into account the full extent of the temporal structure of the data (Young, 1988). This leads to a more accurate and reliable data set, which is crucial for subsequent analysis and forecasting. Specific time series imputation methods are discussed later in chapter 5.2, where they are implemented.

5 APPLICATION OF DEMAND APPROXIMATION METHODS

This chapter is devoted to the application of the demand approximation methods to the generated data set. First, I present the data set generated by using a Python programming software. The data set includes stockouts and has trend and seasonal attributes. Once we understand the data, only than we can discuss which forecasting and approximation methods are appropriate to use.

In the second part of this chapter, I introduce methods to approximate the demand during the period of stockouts. This is the prerequisite for being able to apply the forecasting methods in chapter 6. Additionally, topic of selecting the appropriate approximation method is discussed, implementation of the method in R software is explained, followed by evaluation and representation of results.

5.1 Generation and presentation of the data set

For the purposes of this master's thesis, the ideal data set should have frequent stockout periods as a necessary characteristic. Additionally, it should have a span of at least three years in order to capture long-term sales trends and seasonal cycles. Data set should have clear seasonal patterns that reflect periodic sales increases or decreases within each year. In addition, an overall trend, that is, an increase or decrease in sales, is needed to simulate realistic market behavior over time. Since it is hard to find the data set that would match all of those characteristics, I realized that the best option is to create my own data set with such specifications that can be used for this particular problem demonstration.

With the criteria above being established, I constructed a simulated data set using Python programming software. For this purpose, I used libraries such as *pandas* for data manipulation, *numpy* for numerical computations, and *faker* library to generate simulated data. The process begins with generating a date range for a specific time period, with daily frequency. This is followed by a linearly increasing trend is added which exponentially boost sales. After that, the sales are generated as combination of trend and "starting sales" variable.

After that, I introduced the stockout periods that are 6 to 10 days long, occurring 25 times. The stockouts are intentionally excluded from the last year of the data set, since I will use this year only as a test set to compare accuracy of forecasting and approximation methods.

In addition, a seasonal component is included in the data set. This is achieved by increasing sales in the seasonal, pre-seasonal and post-seasonal months. Here, sales in the seasonal months are increased by 1.6, while sales in the pre-seasonal and post-seasonal months are increased by 1.25.

Finally, the data set is structured in a data frame and saved in an Excel format for easier further use. Each step of the dataset creation is shown in the Python code in Figure 11. These steps are followed by comments that provide clarity and guidance, allowing the reader to navigate through the code and understand the methodology used in the creation of this dataset.

An important thing that I want to mention is that, in this data set, only the sales with value of 0 indicate stockouts, which can be seen in the code. If sales are any higher than 0, that can only indicate low sales for that day, and not stockout. This is done for benefit of easier problem demonstrations. However, in real world scenarios this might not be the case, and companies should be able to flag stockouts as they occur, as discussed earlier in Chapter 4.2.

Figure 11: Python code used to generate the data set

```
# import libraries
from faker import Faker
import pandas as pd
import numpy as np

fake = Faker()

# Set random seed for reproducibility
np.random.seed(0)

# Generate the date range
start_date = pd.to_datetime('2019-01-01')
end_date = pd.to_datetime('2023-12-31')
dates = pd.date_range(start=start_date, end=end_date, freq='D')

# Create an increasing trend
trend = np.linspace(0, 1, len(dates))
starting_sales = 30

# Calculate sales as a combination of trend and starting sales
sales = (trend * 80 + starting_sales).round(0)

# Introduce gaps of 6 to 10 days with zero sales for 25 periods (stockouts)
gap_periods = 25
gap_length = np.random.randint(4, 7, gap_periods)
gap_start_indices = np.random.choice(np.where(dates.year != 2023)[0],
gap_periods, replace=False)

for i in range(gap_periods):
    gap_start = gap_start_indices[i]
    gap_end = gap_start + gap_length[i]
    sales[gap_start:gap_end] = 0

# Multiply sales by a scaling factor for higher sales during summer months
summer_months = (dates.month >= 6) & (dates.month <= 8)
sales[summer_months] *= 1.6  # Seasonality sales boost

# Multiply sales by a scaling factor for higher sales during almost-summer
months
almost_summer_months = (dates.month == 5) | (dates.month == 9)
sales[almost_summer_months] *= 1.25  # Pre-seasonality sales boost

# Create a DataFrame with generated data
df = pd.DataFrame({'Date': dates, 'Sales': sales.round(0), 'Customer Name':
[fake.name() for _ in range(len(dates))]})
```

Source: Own work.

To summarize, the code from Figure 11 generates the data set created from the daily sales from the beginning of 2019 to end of 2023. The data set shows daily increasing sales figures as the algorithm contains an exponentially growing trend. The data set also contains months where sales are higher due to the influence of seasonality.

Figure 12 is based on data from the generated data set I just described. Although the data set is based on daily sales, the figure is created based on a monthly aggregation of sales to eliminate the noise of daily fluctuations and provide a clearer picture and focus on broader patterns. The solid red line represents the total sales for each month. Each point on this line corresponds to the total sales for a particular month, making month to month fluctuations easily spotted. Note that there will be at least some sales every month, as stockouts never last a whole month. For this reason, the red y-axis, which represents actual sales, will not reach the value 0. However, there are some days in the month that have zero sales, and these stockouts are represented by the blue line on the secondary y-axis on the right. We can see a clear correlation between the increase in stockouts (blue line) and the decrease in actual sales (red line). Sudden drops in sales can be clearly explained by spikes in stockouts.

Figure 12: Preview of the data set

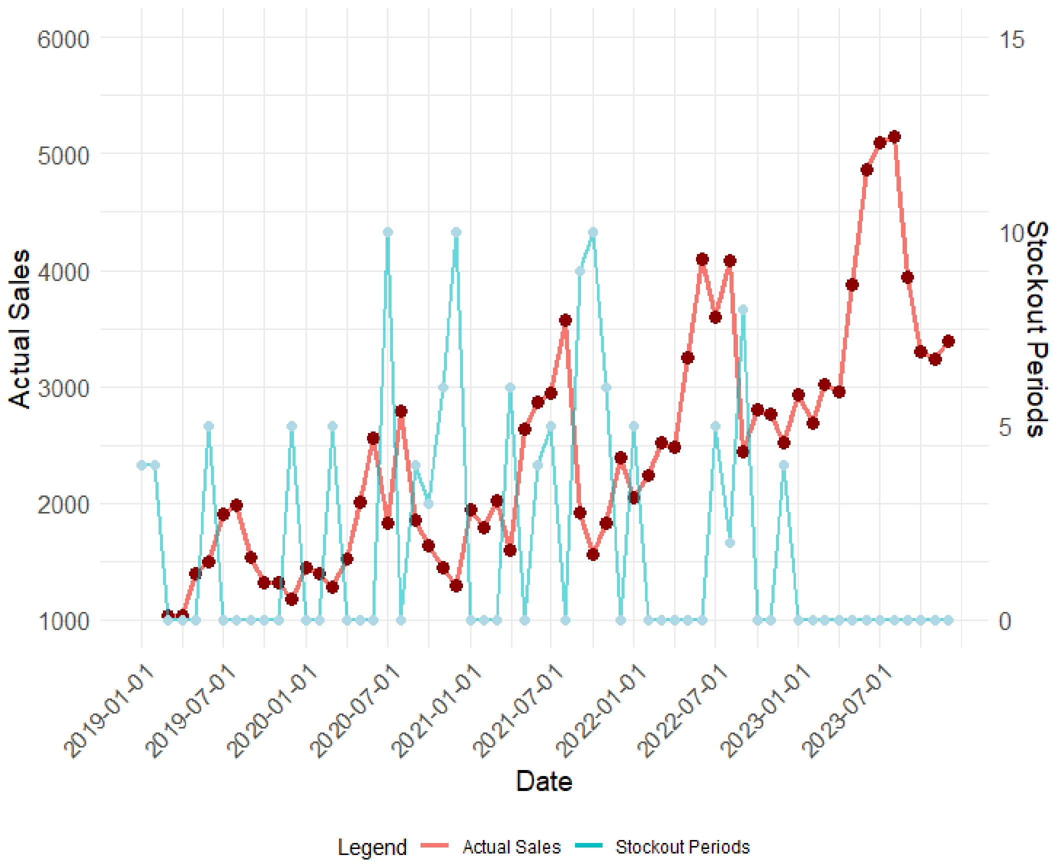

Source: Own work.

Let us now discuss the key observations from Figure 12 and better understand the data set. We have sales data from beginning of 2019 until the end of 2023, aggregated monthly. There is a clear upward trend in sales over the years. Starting from around 1000 units sold in January 2019, sales rise to over 5000 units in July 2023. We can also see a clear seasonality in the summer months. The peak of seasonality starts in May and ends in

September. And finally, we can identify some anomalies that indicate the stockout periods in this month. Once again, as can be observed, there are no stockouts in the last year, from the beginning until the end of 2023. This is done because later, when demand forecast methods are done, this year will serve as a test set against which accuracy of different forecasting and approximation methods will be compared.

5.2 Implementation of methods to approximate demand in R software

5.2.1 Optimized average of last known periods

Given our familiarity with the data set and the method in question, we recognize that a basic version of this method would be inadequate. Therefore, I will begin by optimizing the method to ensure its efficacy for the data set presented. Due to the consistently high growth trend, I will tailor the stockout period approximations by taking an average from the sales data only from the 15 days preceding a stockout. This approach assigns less weight to older data points. The problem also arises in evident seasonality, so certain adjustments to the code needed to be made. For example, when approximating sales at the start of May, the method considers the 15 days prior, which is end of April. This approach could lead to inaccuracies as April is the month that is not affected by seasonality and therefore sales are lower than in May. Same goes for June where sales are higher than compared to May, which is pre-season. Accordingly, I modified the method to account for such seasonal shifts. On the other hand, approximation of stockouts for first days of September must be decreased since it will take average from August, which is month affected by seasonality. Also, approximation for beginning of October must be decreased since it takes average from pre-season month September.

Figure 13 shows how the optimized average of last known periods performed. Once again, for purpose of better visualization, I am aggregating our daily sales data into monthly sales figures. Last year is excluded from the visualization, since it does not have any stockouts, and hence there is nothing to approximate for that year. By examining this figure, we can compare sales without approximated demand during stockout periods, indicated by the red line, against sales with approximated demand using this method, indicated by the blue line. The red line represents the actual sales, which correspond to those in the previous figure. This confirms that the sudden sharp drops in sales represent specific stockout periods within that month. But even without that knowledge, we can easily conclude that every discrepancy of red line against the blue one indicates the stockouts in that period. It is evident how the method filled in data during the stockout periods, and as anticipated, no anomalies are currently observed. Observing the blue line, we can now clearly spot seasonal patterns, which were difficult to discern previously.

Figure 13: Sales and approximated demand data using optimized average of last known periods method

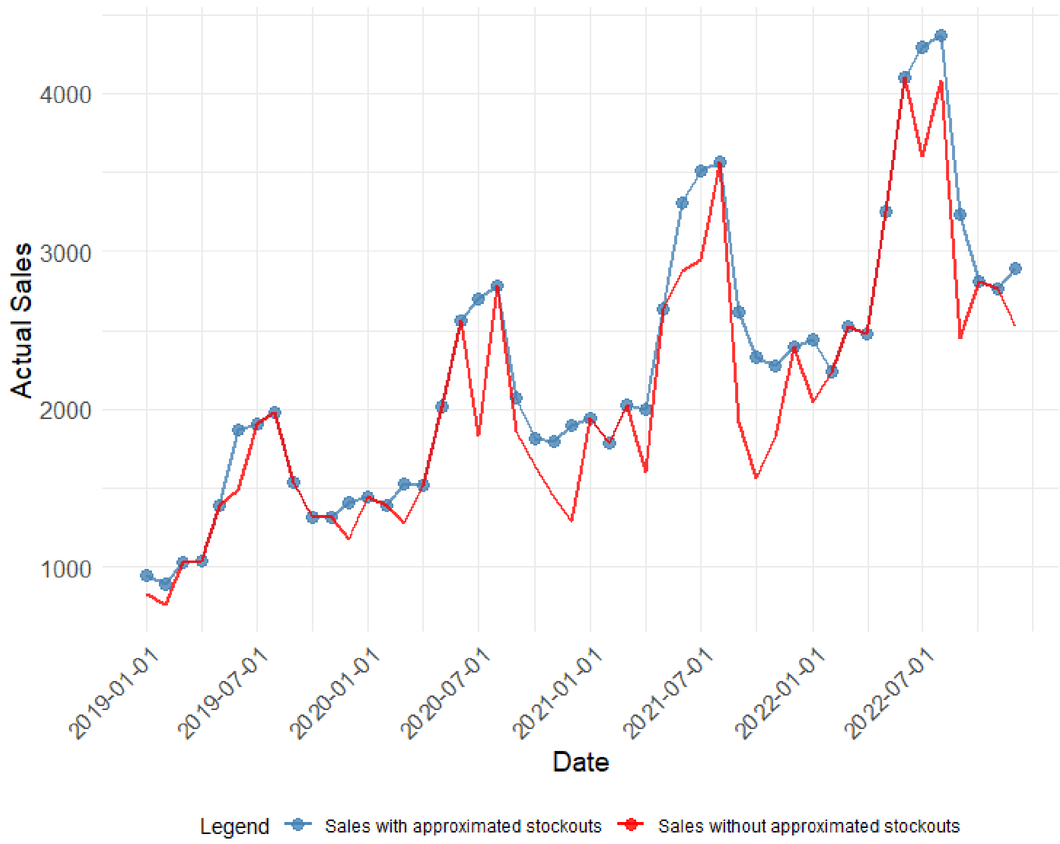

Source: Own work.

5.2.2 Interpolation

In chapter 4.2.2 several interpolation methods were discussed. Now, when the data set is introduced, I can discuss which of those interpolation methods are going to be use and which not.

Let's discuss the "nearest neighbor", which I introduced as our first interpolation method. I previously highlighted its straightforward nature, where missing values are substituted with the data from the closest known point. One of the limitations of the method is that it can ignore important patterns such as trends and seasonality, which are crucial for interpreting our daily sales data, characterized by clear trends and seasonal fluctuations. Secondly, the method proves inadequate when dealing with stockouts lasting multiple days, as it assigns identical values to each day, ignoring the growth trend. This could be particularly misleading if the nearest data point also happens to be a day with zero sales, as it would suggest a continuation of no sales. Hence, I have decided against the usage of the nearest neighbor method for this specific case.

The second introduced interpolation method was linear interpolation. This model assumes linear relationships between data points. In our example, that is not always the case. Linear interpolation simply draws a straight line between two known data points to estimate missing values (Lepot et al., 2017). On one hand, we observe a clear and steady trend, which theoretically could be captured by linear interpolation. However, since our data set consists of daily data points, and attributes such as seasonality and stockouts, this method will have some challenges. For instance, if a stockout occurs over 4 days and we assume that the sales data point just before the stockout days and the one just after accurately represent demand, linear interpolation for the points in between could indeed be considered. However, the challenge arises when we account for seasonality, as linear interpolation might oversimplify the patterns present in daily data. We can take an example of stockout occurring at the last days of May. Linear interpolation assumes a constant rate of change between data points. Hence, it would estimate values between the last available data point before the stockout and the first data point after it, where sales might be much higher due to seasonal effect, and linear interpolation may fail to account for that change in demand. That is why I will avoid this method in this case.

The first interpolation method that I am going to use for approximation of stockouts is polynomial interpolation. I will use this method with high-degree polynomials, which can give more accurate results.

The second interpolation method that I am going to use is cubic spline interpolation. I will use it because of its ability to accommodate a diverse range of data shapes, and effectively capturing nonlinear trends more robustly than linear interpolation. Unlike polynomial interpolation, cubic spline interpolation uses multiple segments of third-degree polynomials to connect the points, which ensures a smoother curve and avoids the large fluctuations that high degree polynomials can cause (Burden & Faires, 2016).

Both the cubic spline and the polynomial interpolation methods were implemented in R software. When implementing the cubic spline interpolation, the data set was loaded into the environment and the sales column was converted into a time series using the ts() function to ensure temporal coherence. In the next step, the na.spline() function from the "zoo" package was used, which performs cubic spline interpolation by default. Essentially, the function identifies the NAs (flagged missing values) in the data set. It then fits a cubic spline to the non-missing values, ensuring a smooth and continuous curve that passes through these known data points. Once the spline is fitted, it is used to estimate the sales figures for the days with missing values, filling the gaps in the data with estimates based on the observed sales pattern.

The application of polynomial interpolation was not as simple as with spline interpolation. The main point that required some extra work was to find the degree of the polynomial that was best suited for our data set. As mentioned earlier, the general rule is that a higher degree polynomial fits the data better but can lead to overfitting and fluctuations between points (Hawkins, 2004).

The polynomial degree was chosen by creating a function in R software, as we can see in the code in Figure 14. First, I initialized variables to store the results. Then I chose a for loop that iterates from 1 to 27, since 27 is the highest degree that can be applied to this data set. Each degree matches the model, and AIC score is calculated for each model. At the end, the if statement updates the variables best_degree and best_aic if the condition is met. In our example, the best degree of the polynomial is 26, which is reasonable since our data emphasizes trend and seasonality, so our polynomial must be able to capture this.

Figure 14: Finding best degree of polynomial interpolation

```r
# Initialize variables to store the results
best_degree <- 0
best_aic <- Inf

# Try different degrees of the polynomial
for (degree in 1:27) {
  # Fit the model
  model <- lm(Sales ~ poly(Days, degree), data = sales_data_polynomial,
na.action = na.exclude)

  # Calculate the AIC
  aic <- AIC(model)

  # Update the best degree and AIC
  if (aic < best_aic) {
    best_degree <- degree
    best_aic <- aic
  }
}
print(paste("Best Degree:", best_degree))
print(paste("Best AIC:", best_aic))
#Best polynomial for our example is 26
```

Source: Own work.

When the best degree of the polynomial is chosen, we can proceed to the next step. This step involves fitting the polynomial model to the sales data with respect to the date using the lm() function. Within lm(), the poly() function is used to create polynomials of the specified degree (26 in our example) for the date variable. The fitted model is then used to predict the demand for the time of inventory.

Figure 15: Polynomial vs. Spline interpolation

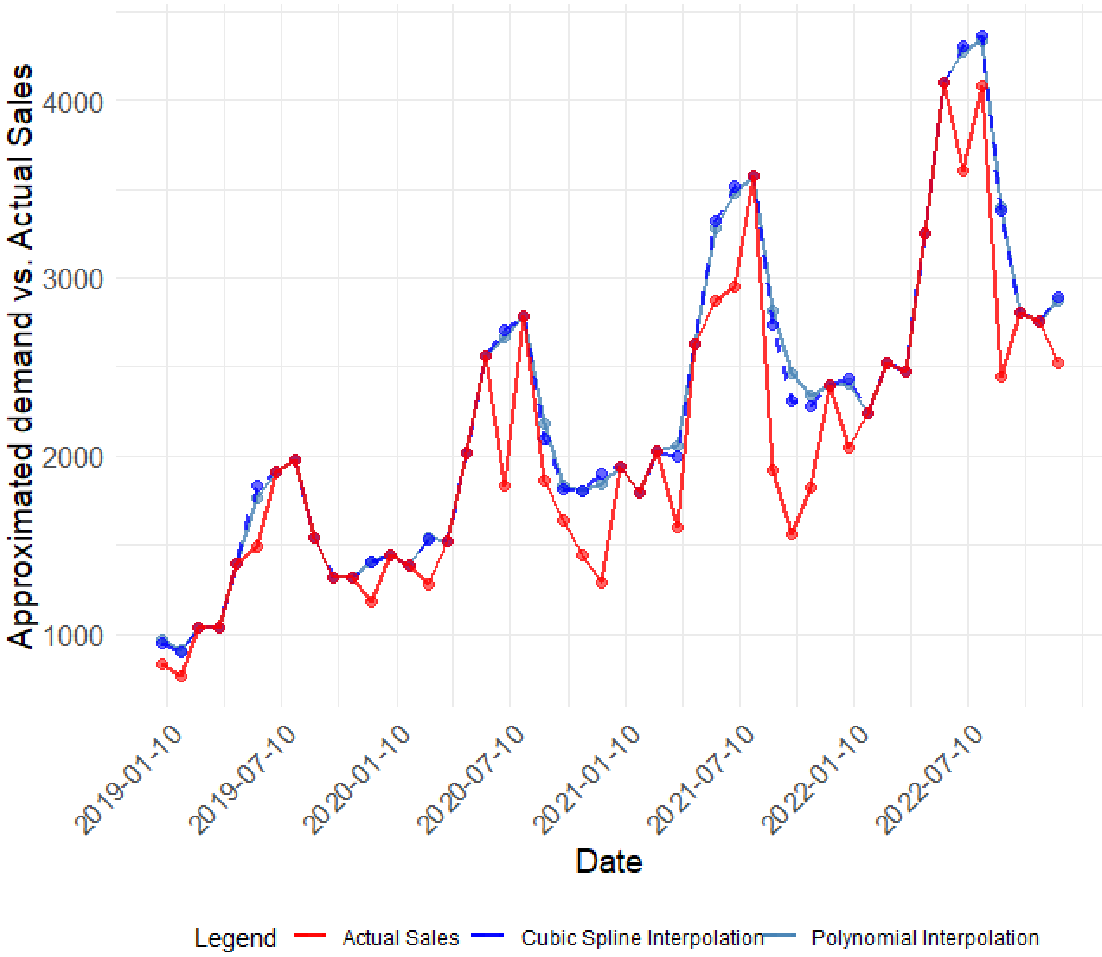

Source: Own work.

In Figure 15, we observe the values approximated using both polynomial and cubic spline interpolation. Once again, for purpose of better visualization, I am aggregating our daily sales data into monthly sales figures. The exclusion of the previous year from the visualization is based on the absence of stockouts, eliminating the need for any approximations during that timeframe.

As can be observed, both cubic spline and polynomial interpolation gave comparable demand approximations during stockout periods, exhibiting slight variations in specific intervals. Both methods seem to effectively mirror the data set pattern, adhering to its seasonal and trend characteristics. It can be observed how both interpolation methods filled in data during the stockout periods, and as anticipated, no anomalies are currently present.

5.2.3 Time series imputation using ImputeTS in R

Previously, in chapter 4.2.3, theoretical underpinnings of time series imputation were explored. There it was mentioned that many imputation methods rely on inter-attribute correlations, while time series imputation instead needs to employ time dependencies. However, in R software, there is a package that is mainly designed for time series. The package is called ImputeTS, which is short for "time series missing value imputation". This package is a collection of methods and tools used for approximation (imputation) of missing/incomplete values in time series, and it offers several different imputation algorithms. More information about this package can be found on R-project.org. (2019).

Of course, not every method is suitable for our data set. Certain methods, such as "na_locf" (Last Observation Carried Forward), which borrows the nearest neighbor's value, and "na_mean", which fills missing/incomplete data by calculating the mean, median, or mode, will be excluded from our application, since they will fail to capture non-stationarity of the data set.

In this section, I will point out which methods from this package are going to be used and why. Additionally, I will shortly explain each algorithm and compare them.

The first missing value imputation method I have used is the weighted moving average, for which I have used the "na_ma" code. I have explained in chapter 3.2.1 that the weighted moving average method gives good results due to its adaptive nature, where the window size can increase in case of long stockout periods (NAs), allows it to capture and adjust to trends in the data. The method offers different types of weighting such as simple, linear, and exponential weighting. For our example, exponential weighted moving average will be used because of its ability to adjust quickly to recent changes, making it especially relevant for approximating stockouts based on recent trends. With exponential weighting, if we were to predict value at position E based on A, B, C, and D, the values for those position will be given a weight of 1/16, 1/8, ¼, and ½ respectively.

An additional parameter is the value of k, which is integer width of the moving average window. However, with exponential weighting, the impact of observations diminishes exponentially as they move away from the missing value, making adjustments to the k parameter irrelevant. Hence, I will leave default value for this parameter.

The second missing value imputation method is the Kalman smoothing method, denoted as "na_kalman". This method can use two different methods (algorithms): Kalman Smoothing on structural time series models ("StructTS"), and state space representation of an ARIMA model ("auto.arima"). Structural time series, as well as Kalman Smoothing were introduced earlier in chapter 4.2.3. We can specify which algorithm this method will use, by stating them as a parameter when calling the "na_kalman".

"StructTS" is a method in the R software designed for fitting structural time series models using maximum likelihood estimation. It creates model based on the decomposition of the series into a number of components such as trend, seasonal, and irregular components. The resulting model is then utilized by "na_kalman" to impute missing values through Kalman smoothing.

In "na_kalman", "auto.arima" will use state space representation of an ARIMA model. The auto.arima in R software automates the process of selecting the best ARIMA model for a given time series based on the AIC score. The state space representation of the chosen ARIMA model is then utilized by "na_kalman" to impute missing values using Kalman smoothing.

Figure 16: Different time series imputation methods

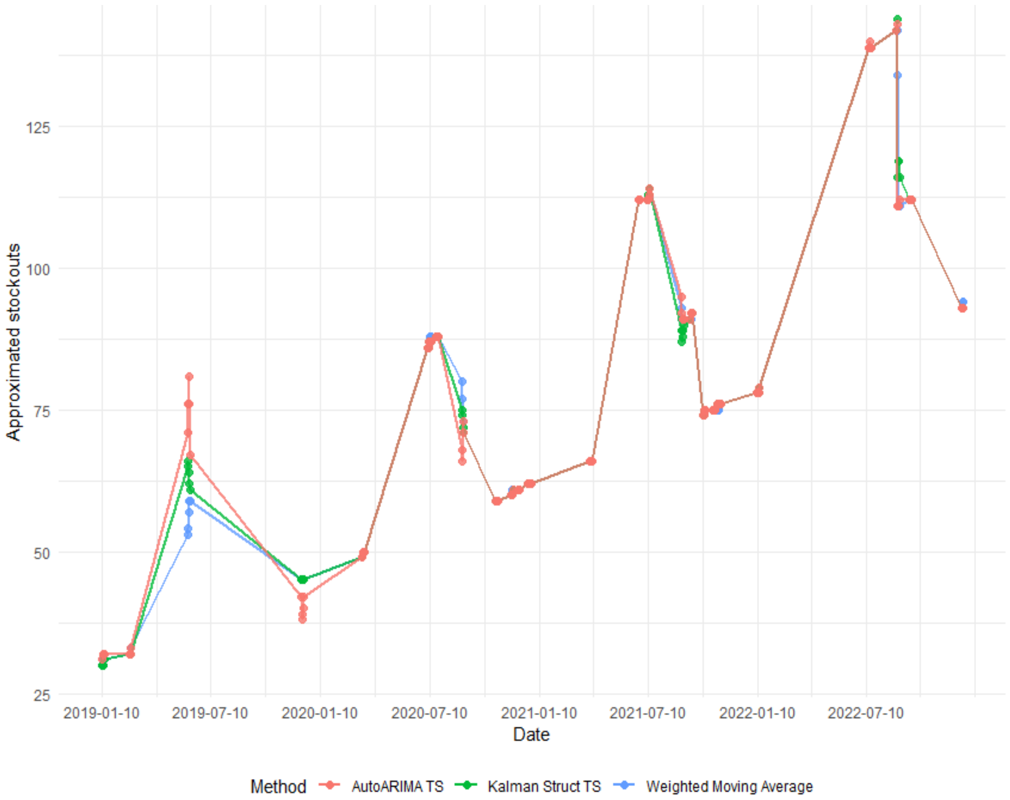

Source: Own work.

Both methods have their benefits and challenges, and both are suitable for our specific data set. Figure 16 illustrates the outcomes of applying these time series imputation methods to address stockout data points. It is important to note that only the data points subjected to imputation (the periods of stockouts) are plotted to avoid cluttering the graph with the extensive actual sales data, which would make the differences in imputed values less observable. Each method is represented by a distinct line color on the graph.

It can be observed that every time series imputation method gave similar results. Careful selection of imputation techniques as well as their imputation process might justify this. Despite the overall similarity, we can identify the periods where the differences between the imputation methods are more significant. At the beginning of 2019, for example, there are clear differences between the methods, as we can see in Figure 16. This is because there is the least data to observe. The further we go, the more similar they become, as the amount of data provides a clear trend, and each method manages to capture part of it. This again emphasizes the importance of having enough data.

6 COMPARISON OF DIFFERENT DEMAND FORECASTING AND APPROXIMATIONS METHODS

In this chapter, I delved into forecasting methods and their application to the data set. The essence of forecasting is not just in predicting future values, but also in ensuring that the predictions are as accurate as possible.

I started by implementing two forecasting methods: Triple exponential smoothing (Holt-Winters method) and auto ARIMA. In this chapter I will explain how I used these specific forecasting methods. Both forecasting methods are applied six times, that is, once for each approximation method.

In addition, I have broken down the complexity of the methods chosen. For instance, while using the Holt-Winters method, I tailored it to our data set, manually adjusting the alpha, beta, and gamma parameters to optimize for the trend, seasonal, and residual components. On the other hand, the auto.arima function in R software had an inherent ability to choose the best-fitting ARIMA model.

Using a loop that cycles through results of each approximation methods, I was able to evaluate the performance of both approximation and forecasting methods. As a part of this loop, I calculated forecasting accuracy metrics such as MAE, MSE, RMSE, and MAPE to compare the performance of our models against the actual test data for 2023. These results are used to compare different approximation methods as well as these two forecasting methods.

I will also apply the Holt-Winters and auto ARIMA forecasting methods on the original data set with stockouts still present. This will emphasize the difference between doing a forecast on “clean” data where demand in stockout periods is approximated against the original unadjusted data set.

6.1 Application of different demand forecasting methods

Previously, in chapter 3, I addressed the challenge of forecasting time series data characterized by both trend and seasonality. To accomplish this, I selected two appropriate forecasting methods that are suitable for this specific data set: triple exponential smoothing (Holt-Winters) and auto ARIMA.

The reason why I am using auto ARIMA and not simply SARIMA, which is suitable for data with trend and seasonality, is because as good demand planners and analysts, we want to be as efficient as possible. Auto ARIMA is a convenient function in the forecast package that automatically determines the best ARIMA model and its parameters for a given time series, using a combination of unit root tests, minimization of the AIC and maximum likelihood estimation. It makes the method selection process quicker and more automated, especially when handling multiple time series.

The challenge with using only SARIMA is that we would need to determine the right parameters either through domain knowledge, diagnostics plots, or through a grid search, which can be time-consuming.

6.1.1 Data preparation

The initial step involves merging all approximation methods, that is, the introduced data set optimized using each approximation method, into a single data frame along with the date column. This merge is performed to facilitate the implementation of the forecasting process, as explained in the next sub-chapter. Additional column in this data frame is column with original sales data. This allows us to see the forecast results in case when stockout periods are not approximated.

The second step is dividing data into training and test set. Training set uses data from the beginning of data set (01-01-2019) until the beginning on 2023. This is the period where we had stockouts and in which different approximation methods were applied. On the other hand, the test set uses the data from the beginning of 2023 till the last point in the data set which is 31-12-2023.

The test set is important because it simulates real-world scenarios in which the method encounters new unseen data. Essentially, it helps us to evaluate the performance of the method, that is, to compare performance of different methods. Prior to forecast, it is smart to convert regular sales data to time series, as this ensures preservation of chronological order, which is essential for analyzing trends, seasonal patterns, and cyclic behavior. This step requires parameter for frequency of data set to be provided. Since the data set have yearly seasonality and consists of daily sales, this parameter should be set to 365. This is essentially the number of data points in one seasonal cycle.

6.1.2 Forecasting process

In this section, demand forecasting is performed using both Holt-Winters and Auto-ARIMA for the data frame described in the previous subsection. As I will describe in a moment, both forecasting methods have a similar application process in the R software. Essentially, both demand forecasting methods are applied to all columns of the data frame except the date column. The results of each forecast are stored along with the values of the accuracy measurements and the parameters used so that they can be compared in the next subchapter.

Let's describe code written in R software for the application of the Holt-Winters method by observing code in Figure 17. There is a comment before every part of code, which enables easier reading and evaluation.

The code in Figure 17 consists of five nested "for" loops. The initial "for" loop iterates through every column in the mentioned data frame except the date column. Before, I mentioned that every column represents data optimized using a specific approximation method. The creation of this forecasting method takes several parameters, as we know from chapter 3.2. Therefore, for each approximation method, the second, third, and fourth "for" loops are iterating through different values of parameters alpha, betta, and gamma. Fifth loop test every combination of parameters further using additive or multiplicative seasonality type. For each combination of these parameters Holt-Winters model is created (hw_model). Frequency parameter is set to 365, and as already mentioned, this corresponds to the number of data points in a seasonal cycle.

The next step after model creation is the forecast itself, which is based on the created method (hw_method). Additional parameter that needs to be specified is the length of the forecast. The length of the forecast corresponds to the length of the test set, that is, one year. Forecast is followed by evaluation of its accuracy, using different forecasting accuracy measurements. These measurements are MAE, MSE, RMSE, and MAPE. Here, these measurements are based on comparison between the forecasting results and actual sales from the test data set.

As final step in the loop, accuracy of the approximation method, with given set of parameters is evaluated according to MAPE. If the error is lower than the previous one, those parameters are saved as current best. This iterative approach results in obtaining the best set of parameters for a specific approximation method.

Figure 17: Holt-Winters forecasting process

```
# Forecasting and Accuracy Measurement
for (col in names(Data)[!names(Data) %in% "Date"]) {

  alpha_values <- seq(0.1, 1, by = 0.1)
  beta_values <- seq(0.1, 1, by = 0.1)
  gamma_values <- seq(0.1, 1, by = 0.1)

  seasonality_types <- c("additive", "multiplicative")  #

  best_mape <- Inf
  best_params <- c()
  # Iterate through parameters
  for (alpha in alpha_values) {
    for (beta in beta_values) {
      for (gamma in gamma_values) {
        for (seasonality_type in seasonality_types) {

          # Fit Holt-Winters model
          hw_model <- HoltWinters(ts(train_data[[col]], frequency =
365),
                                  alpha = alpha, beta = beta, gamma =
gamma,
                                  seasonality = seasonality_type)

          # Make forecast
          hw_forecast <- forecast(hw_model, h = nrow(test_data))

          # Calculate all forecasting error metrics
          hw_metrics <- c(
            MAE = mae(test_data[[col]], hw_forecast$mean),
            MSE = mse(test_data[[col]], hw_forecast$mean),
            RMSE = rmse(test_data[[col]], hw_forecast$mean),
            MAPE = mape(test_data[[col]], hw_forecast$mean)
          )

          # Check if current combination yields a better MAPE
          current_mape <- hw_metrics["MAPE"]
          if (current_mape < best_mape) {
            best_mape <- current_mape
            best_params <- c(alpha = alpha, beta = beta, gamma = gamma,
seasonality = seasonality_type)
          }
        }
      }
    }
  }

  # Save parameter choices to the dataframe
  parameter_results <- rbind(parameter_results, c(col, best_params))

  # Save forecasting errors
  results_hw[[col]] <- list(hw_metrics = hw_metrics)
}
```

Source: Own work.

Auto-ARIMA forecast method follows the same logic as the Holt-Winters forecast, however it is somewhat simpler. The main difference is that there is no second, third, fourth and fifth “for” loop. This is because Auto-ARIMA selects the most optimal method (ARMA, ARIMA or SARIMA) and its parameters for a given data set by iterating through different combination of parameters and selecting the most optimal ones based on AIC score. As before, Auto ARIMA also creates a forecast method with suitable parameters for each column (optimized data set). The forecast is then created based on the method created, with the length of the forecast corresponding to one year. The created forecast is followed by an evaluation of its accuracy, using different forecasting accuracy measurements, same as before. As mentioned, these measurements are based on the comparison between the forecast results and the actual sales from the test data set. Auto ARIMA did not choose the same parameters for each approximated data set. However, it chose the SARIMA method for each of them, which is to be expected given the characteristics of the data set.

As a result of described forecasting process, Table 4 has been generated, displaying the parameters for the Holt-Winters forecasting method, which are selected by the algorithm from Figure 17, as well as the parameters optimized by Auto ARIMA method. As can be observed, these parameters can be quite different for datasets optimized with different approximation methods.

Table 4: Evaluation of parameters for the data set adjusted by different approximation methods

Approximation method used in the data set	**Holt-Winters** $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma})$	**Auto ARIMA** $(\boldsymbol{p}, \boldsymbol{d}, \boldsymbol{q})\ (\boldsymbol{P}, \boldsymbol{D}, \boldsymbol{Q})\boldsymbol{m}$
Un-approximated data set	(0.8, 0.7, 0.3)	(2,0,3)(0,1,0)365
Average of last known periods	(0.9, 0.5, 0.1)	(5,0,0)(0,1,0)365
Polynomial interpolation	(0.8, 0.4, 0.1)	(2,0,1)(0,1,0)365
Spline interpolation	(0.9, 0.5, 0.1)	(5,0,0)(0,1,0)365
Kalman Struct TS	(0.8, 0.2, 0.1)	(1,0,2)(0,1,0)365
AutoARIMA TS	(0.9, 0.2, 0.1)	(4,0,0)(0,1,0)365
Weighted moving average	(0.1, 0.2, 0.1)	(5,0,0)(0,1,0)365

Source: Own work.

For the Holt-Winters method, every approximation method performed more accurately with multiplicative seasonality. This is because, as I mentioned earlier in chapter 3.2, multiplicative seasonality is appropriate when the seasonality effect increases or decreases proportionally with the overall trend, which is the case for our data set. One note is that multiplicative seasonality could not be performed for unapproximated original data set, since in R software data must be non-zero for multiplicative Holt-Winters method. Here, the forecast is done using additive seasonality. The parameters α and β, which control the level and the trend component respectively, are similar across each approximation method. In practice, these parameters should generally be smaller than 0.7, otherwise the history is not given much importance. In Table 4, we can see that α is usually not smaller than 0.7, and this might be justified with the constant increasing trend. On the other hand, β is usually smaller than 0.7. The parameter that is fairly constant is γ, which fluctuates around 0.1 to 0.3. This is actually good, since in practice, this parameter should be low (<0.3). This is because from a business perspective, we should not assume that seasonality changes drastically from one year to the next. With high γ there is more room for overfitting (Vandeput, 2021).

In Table 4, we can observe something similar when we look at the Auto ARIMA parameters. The parameters vary slightly between the different approximation methods. However, the Auto ARIMA model has chosen SARIMA for each of them, which is expected due to nature of the data set. To summarize briefly, the non-seasonal ARIMA parameters (p, d, q) reflect the emphasis on past observations and the adjustment for past errors, while the seasonal parameters $(P, D, Q)m$ account for the seasonality in the data. m reflects number of observations per year, which is 365, as explained earlier. Seasonal parameters are constant for all approximation methods. This indicates first-order differencing at the seasonal level, suggesting that the model takes into account the difference between an observation from this season and the previous one. For the non-seasonal parameters, we see the constant 0 for the d component, indicating that the model has not recognized the need for high differentiation to achieve stationarity. For the p component, we see fluctuations indicating that different values were assigned to past observations. The q component is generally low, indicating a low recall of forecast errors.

6.2 Comparison of results

Now, when demand forecast is performed on the data set in which the demand for stockout periods is approximated by various methods, as well as with the original data set without approximated demand, the final results can be observed and compared.

Interestingly, the results of the individual approximation methods do not vary significantly, as they all perform well. This can be attributed to the careful selection of the approximation methods, where I immediately excluded the methods that are considered ineffective for our specific problem and data set.

Demand forecast based on each approximation method managed to closely follow trend and seasonal patterns. The least accurate forecast was achieved by relying on the original data set without approximating demand during stockout periods, as expected. Although the forecast based on the original data set managed to also capture trend and seasonal patterns, the omission of demand approximation during stockout periods resulted in the model predicting similar stockout occurrences in the future. Although there may be a pattern to the occurrence of stock-outs, stock-outs are generally random. Therefore, not adjusting for stockouts in the time series data results in a high forecasting error . This emphasizes the importance of approximating demand during stockout periods, as even a reliable forecasting method will give inaccurate results if this aspect is overlooked.

In terms of forecasting accuracy, the demand forecast based on original data set leads to an accuracy of 92.12% when forecast is done using Auto ARIMA method, and accuracy of 90.89% when forecast is done using Holt-Winters method. The accuracy measure used here is based on the MAPE formula, and is expressed as (100 – (MAPE *100)). On the other hand, demand forecasts based on each approximation method have much higher accuracy. Here, the most accurate forecast have an accuracy of 98.79%, while the least accurate forecast have an accuracy of 97.43%. The approximation method that resulted in highest accuracy is spline interpolation when performed with Holt-Winters forecasting method, while the least accurate approximation method is polynomial interpolation when performed with the Auto ARIMA forecasting method.

Except for the demand forecast based on the original data set, Holt-Winters outperformed Auto ARIMA across all approximation methods. The demand forecasts generated by Holt-Winters have an accuracy of approximately 98%, as measured by MAPE, whereas the demand forecasts produced by Holt-Winters showed an accuracy of around 97%. Since the difference between most and least accurate approximation method is less than 1.5%, for the remainder of this chapter, I will concentrate only on the most and least accurate approximation method, comparing them against demand forecast based on the original unapproximated data set.

It is important to mention that, even though Holt-Winters performed slightly better for this specific situation, this does not mean that we should always use it instead of Auto ARIMA. It is always a good idea to test several options and to see which forecasting method is the best for the current data set.

Let’s now visually compare the results of the demand forecast based on the least accurate and the most accurate approximation method and the forecast based on the nonapproximated original data set. I will plot only results of these three methods so that the differences can be easily compared without overcrowding the graph. Figure 18 below shows only the forecasted part, that is, forecasted daily demand for one year period, from the beginning of 2023 to the end of 2023. Here I have also added the original sales from the test set so that we can observe the discrepancy between the forecasts and the actual sales. In this figure, we look at the daily sales without any aggregation.

Figure 18: Forecast by most and least accurate approximation method vs. original data set

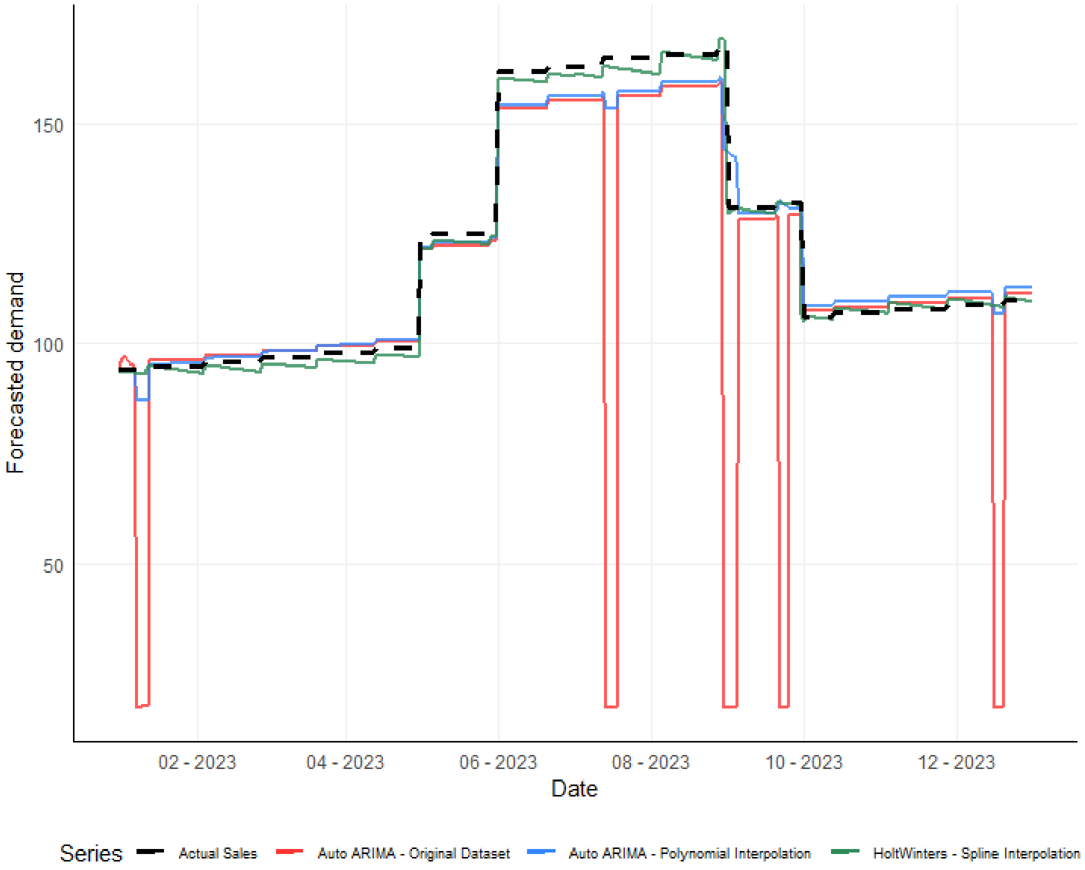

Source: Own work.

In Figure 18, we see that the Auto ARIMA method based on the original data captures some trends and seasonal fluctuations, however, it predicts five periods of significant drops in demand. This is of course due to the fact that the demand during the stockout periods was not approximated, causing the model to predict similar periods of "low demand" also in the future. This shows the magnitude of negative influence that the unapproximated stockouts have on the forecast. Interestingly, Holt-Winters forecast, utilizing the original data, introduced only two low periods of demand. On the other hand, it did not manage to capture correct trend and seasonality, resulting in higher forecasting error.

The most accurate approximation method was able to accurately forecast demand and capture correctly both trend and seasonality when done with Holt-Winters. The least accurate approximation method, using the Auto ARIMA forecasting method, also provided satisfactory results by capturing the trend and part of the seasonality. However, the highest forecasting error results from the inability to effectively predict the sharp increase in demand on days influenced by seasonality. Overall, the demand based on the approximated datasets exhibits a clear absence of noise, sudden drops in demand, or outliers. Instead, it follows the trend smoothly and displays a more stable pattern.

Figure 19: Forecasting accuracy measurements

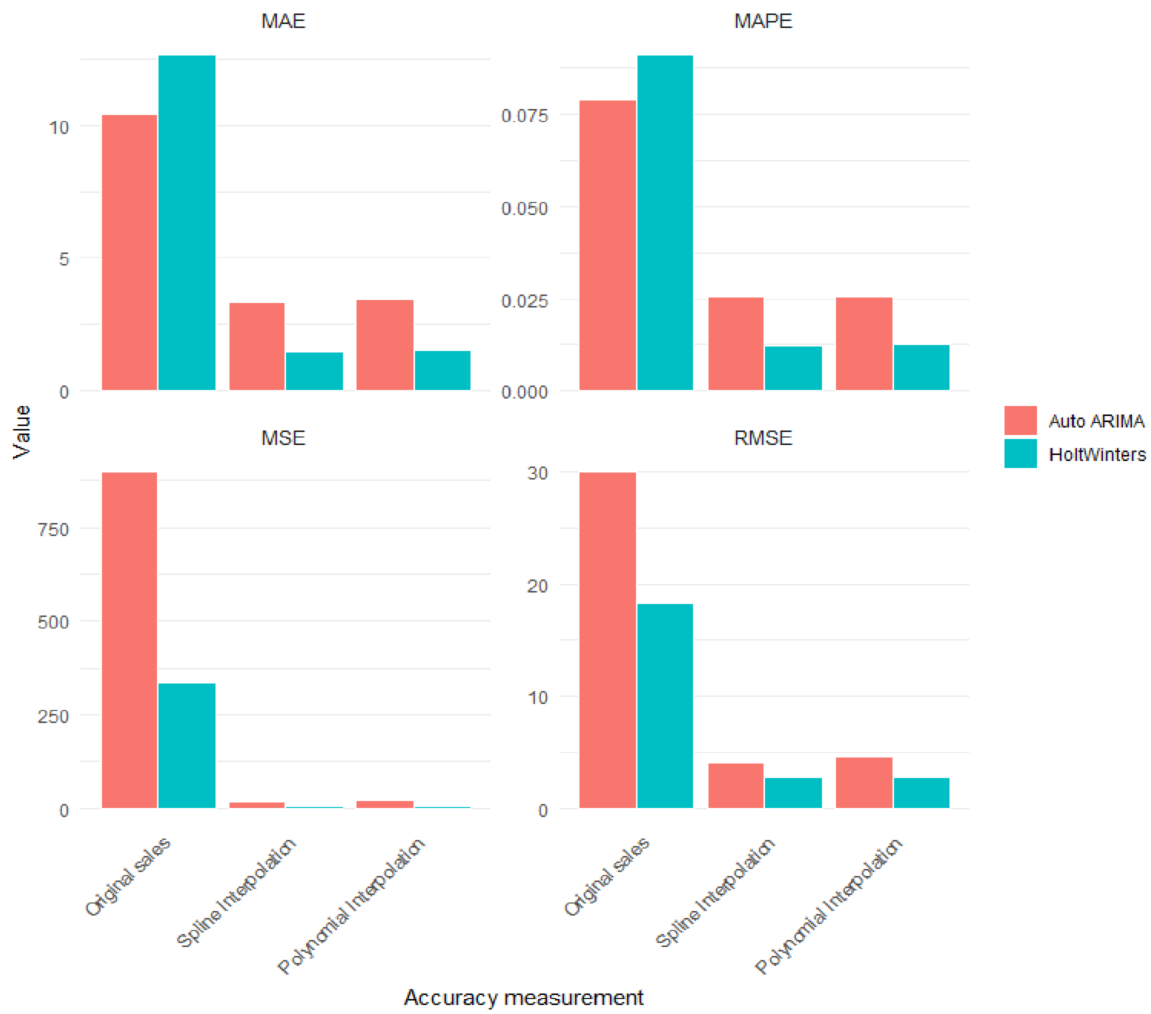

Source: Own work.

As we visually observed differences of demand forecasts against the original sales, now it is time to quantify them. I already mentioned the accuracy measured using MAPE, but let's consider other measures of forecasting accuracy.

Figure 19 shows measures of forecasting accuracy discussed in chapter 3.4. MSE, RMSE, MAE, and MAPE are used to evaluate demand forecasts based on accuracy of most and least accurate approximation method, as well as the original data set. Results of both, Auto ARIMA and Holt-Winters forecasting methods are now displayed. In this figure, every block represents one measurement of forecasting accuracy, each having its own scale. In this representation, the higher the bar, the greater the forecasting error, indicating lower accuracy.

What immediately gets attention is high error measurements of forecast done by both Auto ARIMA and Holt-Winters using original unapproximated original data set. We can also observe how small the difference is between the most and least accurate approximation methods. Actually, the difference between the most and least accurate approximation method done using the Holt-Winters method is less than 0.10%. Same goes for forecasts done by Auto ARIMA.

In actual business, the accuracy achieved using both Auto ARIMA and Holt-Winters is highly desirable. As mentioned, none of the used approximation methods resulted in poor forecasting results. This again emphasizes the importance of careful selection of both approximation method as well as forecasting method.

On the other hand, this example really emphasizes the negative effects of not optimizing the data set by approximating demand during stockout periods. Even though forecasts on original, unapproximated data set gave more or less acceptable accuracy, this might not always be the case. Especially if we are working with datasets that have more complex seasonal and trend patterns, alongside a significant amount of noise and outliers.

7 CONCLUSION

The objective of the master's thesis was to demonstrate the process of effectively forecasting demand based on sales data, particularly in instances of stockouts. I successfully implemented several demand approximation methods and then I applied two forecasting methods to the adjusted data set.

Drawing conclusions from my research, the approximation methods showcased remarkable performance, with all reporting an accuracy over 97%. Some even achieved an accuracy of approximately 98.5%. This can be attributed to careful selection of approximation methods, as well as their implementation. It is also attributed to careful selection of forecasting methods as well as parameters and their implementation. For instance, the fairly simple algorithm used for the Holt-Winters method was running for over several hours in order to test each combination of parameters and get the best results.

However, in order for companies to create an optimal demand forecast, a comprehensive understanding of their data is crucial. Recognizing elements such as seasonality and when it occurs, the magnitude of trends, anomalies such as stockouts, and similar is essential. For instance, while Auto ARIMA displayed significant success, it was not as easy as just applying the code to the data. When deployed with approximated data without specific frequency, it failed to capture seasonality, choosing a regular ARIMA method. Once I specified its frequency parameters, the performance dramatically improved. At this point, I would like to emphasize that one does not necessarily need high programming skills to achieve high forecasting accuracy. The more important part is to understand the data, forecasting methods, and to know how to deal with anomalies such as stockouts or outliers.

In the context of my research, I recommend using any of these approximation methods with both Auto ARIMA and Holt-Winters forecasting methods. While spline interpolation approximation method leveraging the Holt-Winters forecasting method gave the best results for this specific example, it is important to understand that every data set is a story for itself. That is why it is hard to give a general recommendation. However, this thesis provides companies with a structured template on how to deal with forecasting in the case of stockouts. A deep understanding of their data is the initial step. Some of the companies might have data where their trend is not straightforward, maybe they have ups and downs, maybe the trend was not constant for an entire year, etc. In such cases they should, for example, adjust the relevant parameters in the forecasting method and in doing so adjust the method better to the characteristics of the demand.

Additionally, I want to emphasize the importance of constantly measuring forecasting accuracy, as well as constantly reevaluating chosen forecasting method. Data and its attributes can change, potentially leading to a suboptimal performance of the existing forecasting method. Given such scenarios, it is wise to maintain stored scripts, similar to those used to determine Holt-Winters parameters or polynomial degrees in polynomial interpolation. In this way, companies can quickly re-evaluate the parameters used in a particular forecasting method as well as the test results of different forecasting and/or approximation methods.

An additional important recommendation for the companies is to keep track of their inventory levels, so they can flag shortages as they occur. This practice holds significant importance because relying on manual detection of shortages is not a good practice. Of course, the method can be trained to recognize some small sales and mark them as outliers, that is, potential stockouts, but there are better practices than this.

Leveraging the tracking of inventory levels, companies should ideally start collecting the demand as soon as stockout period occurs, that is, when inventory levels drop to zero. For example, companies can do this by order collection, i.e., by using order management system in case of shortages to handle incoming orders and client requests. In this way, collection of demand will continue even if the outbound shipments decline. By doing this, entire demand might not be capture, but significant portion will. Hence, companies should not just rely on approximating demand during the stockout periods, but try to collect it. Of course, they can still approve that collection by applying specific approximation method.

This master's thesis offers a comprehensive theoretical framework for companies to understand and capture demand during stockouts, enabling them to employ demand forecasting methods on their real-life data sets. Additionally, I have practically addressed the challenge of approximating demand during stockout periods for businesses that might have overlooked this step. In this way an accurate forecast can be made, and consequently better business decisions can be taken.

REFERENCE LIST

1. Adhikari, N. C. D., Domakonda, N., Chandan, C., Gupta, G., Garg, R., Teja, S., Das, R., & Misra, A. (2019). An intelligent approach to demand forecasting. In *International Conference on Computer Networks and Communication Technologies: ICCNCT 2018* (pp. 167-183). Springer Singapore.
2. Asur, S., & Huberman, B. A. (2010). Predicting the future with social media. In *2010 IEEE/WIC/ACM international conference on web intelligence and intelligent agent technology* (Vol. 1, pp. 492-499). IEEE.
3. Barratt, S. T., & Boyd, S. P. (2020, July). Fitting a Kalman smoother to data. In *2020 American Control Conference (ACC)* (pp. 1526-1531). IEEE.
4. Bell, P. C. (2000). Forecasting Demand Variation When There Are Stockouts. *The Journal of the Operational Research Society*, *51*(3), 358–363.
5. Bingham, N. H., & Fry, J. M. (2010). *Regression: Linear models in statistics*. Springer Science & Business Media.
6. Brandimarte, P. (2012). *Quantitative Methods*: *: An introduction for business management*. John Wiley& Sons.
7. Brubacher, S. R., & Wilson, G. T. (1976). Interpolating time series with application to the estimation of holiday effects on electricity demand. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, *25*(2), 107-116.

8. Burden, R. L., J Douglas Faires, & Burden, A. M. (2016). *Numerical analysis*. Cengage Learning.

9. Cowpertwait, P. S., & Metcalfe, A. V. (2009). *Introductory time series with R*. Springer Science & Business Media.

10. DeLamater, J. D., Myers, D. J., & Collett, J. L. (2018). *Social psychology*. Routledge.

11. Eroglu, C., & Croxton, K. L. (2010). Biases in judgmental adjustments of statistical forecasts: The role of individual differences. *International Journal of Forecasting*, *26*(1), 116-133.

12. Hawkins, D. M. (2004). The problem of overfitting. *Journal of chemical information and computer sciences*, *44*(1), 1-12.

13. Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: principles and practice*. OTexts.

14. Hyndman, R., Koehler, A. B., Ord, J. K., & Snyder, R. D. (2008). *Forecasting with exponential smoothing: the state space approach*. Springer Science & Business Media.

15. Ivanov, D., Tsipoulanidis, A., & Schönberger, J. (2019). *Global supply chain and operations management: A decision-oriented introduction to the creation of value* (Vol. 2). Springer International Publishing.

16. Kim, S., Kim, H., & Lu, J. C. (2019). A practical approach to measuring the impacts of stockouts on demand. *Journal of Business & Industrial Marketing*, *34*(4), 891-901.

17. Lepot, M., Aubin, J. B., & Clemens, F. H. (2017). Interpolation in time series: An introductive overview of existing methods, their performance criteria and uncertainty assessment. *Water*, *9*(10), 796.

18. Ali, M. H., Al-Adili, A. S., & Sivakugan, N. (2018). Comparison between deterministic and stochastic interpolation methods for predicting ground water level in Baghdad. *Engineering and Technology Journal*, *36*(12A), 1222-1225.

19. Dorne, R., Malpass, J., & Shah, M. (2008). Forecasting and Demand Planning. *Service Chain Management: Technology Innovation for the Service Business,* 51-64.

20. McKinley, S., & Levine, M. (1998). Cubic spline interpolation. *College of the Redwoods*, *45*(1), 1049-1060.

21. Miklos, P. (2004, October). Image interpolation techniques. In *2nd Siberian-Hungarian joint symposium on intelligent systems*.

22. Moritz, S., & Bartz-Beielstein, T. (2017). imputeTS: Time Series Missing Value Imputation in R. *The R Journal*, *9*(1), 207.

23. Moroff, N. U., Kurt, E., & Kamphues, J. (2021). Machine Learning and statistics: A Study for assessing innovative demand forecasting models. *Procedia Computer Science*, *180*, 40-49.

24. Murphey, S. (2022). *The Art of Polynomial Interpolation*. Pennsylvania State University.

25. Parasuraman, A., Zeithaml, V. A., & Berry, L. L. (1988). SERVQUAL instrument. *Journal of retailing*, 64(1), 12-40.

26. Pennings, C. L., van Dalen, J., & Rook, L. (2019). Coordinating judgmental forecasting: Coping with intentional biases. *Omega*, *87*, 46-56.

27. Ritz, C., & Streibig, J. C. (Eds.). (2008). *Nonlinear regression with R*. Springer New York.

28. Hyndman, R. J., & Kostenko, A. V. (2007). Minimum sample size requirements for seasonal forecasting models. *foresight*, *6*(Spring), 12-15.

29. Rossi, R. (2021). *Inventory Analytics* (p. 184). Open Book Publishers.

30. R-project.org. (2019). The Comprehensive R Archive Network. https://cran.r-project.org/.

31. Salomon, D. (2011). *The computer graphics manual*. Springer Science & Business Media.

32. Sanders, N. R. (2017). *Forecasting fundamentals*. Business Expert Press.

33. Shumway, R. H., Stoffer, D. S., & Stoffer, D. S. (2000). *Time series analysis and its applications* (Vol. 3). Springer.

34. Vandeput, N. (2023a). *Demand Forecasting Best Practices.* De Gruyter.

35. Vandeput, N. (2023b). *Forecasting Demand Despite Shortages*. Medium. https://nicolas-vandeput.medium.com/forecasting-demand-despite-shortages-fee899120c08.

36. Vandeput, N. (2020). *Inventory optimization: Models and simulations*. Walter de Gruyter GmbH & Co KG.

37. Vandeput, N. (2021). *Data Science for Supply Chain Forecasting*. De Gruyter.

38. Webster, A. (2013). *Introductory regression analysis: With computer application for business and economics*. Routledge.

39. Wecker, W. E. (1978). Predicting demand from sales data in the presence of stockouts. *Management Science*, *24*(10), 1043-1054.

40. Woodward, W. A., Gray, H. L., & Elliott, A. C. (2017). *Applied time series analysis with R*. CRC press.

41. Young, P. (1988). Recursive extrapolation, interpolation and smoothing of nonstationary time-series. *IFAC Proceedings Volumes*, *21*(9), 35-46.
42. Wang, Z., & Bovik, A. C. (2009). Mean squared error: Love it or leave it? A new look at signal fidelity measures. *IEEE signal processing magazine*, *26*(1), 98-117.

APPENDICES

Appendix 1: Summary of the thesis on Slovene Language

Napovedovanje povpraševanja je ključni poslovni proces, ki se osredotoča na napovedovanje količine blaga ali storitev, ki jih bodo kupci verjetno kupili v določenem časovnem okviru. Napovedovanje povpraševanja je nekaj, kar je danes prisotno v vsakem podjetju, in tudi ljudje, ki se ne ukvarjajo z dobavno verigo ali analitiko, so verjetno že naleteli na ta izraz. Napovedovanje lahko opredelimo kot natančno napovedovanje prihodnosti ob upoštevanju vseh razpoložljivih informacij, vključno z zgodovinskimi podatki, ter znanja in zavedanja o vseh dogodkih, ki bi lahko vplivali na napoved (Hyndman & Athanasopoulos, 2018).

Živimo v svetu, ki je poln nepredvidljivosti, zato je povpraševanje, tako kot vse ostalo, podvrženo negotovosti in ni nikoli konstantno. Ne glede na njihovo velikost ali vrsto, morajo biti sposobni natančno napovedati povpraševanje. Zakaj poudarjam prav danes? Preprosto zato, ker lahko v današnjem poslovnem svetu preprosta napaka, kot je pomanjkanje zalog, stane veliko, saj konkurence ne predstavlja le eno ali dve podjetji, temveč veliko več, kupec pa lahko med podjetji prehaja lažje kot kdaj koli prej. Zato sposobnost natančnega napovedovanja povpraševanja ni le dejavnost z dodano vrednostjo, temveč je nujno potrebna za preživetje. Zaradi natančnega napovedovanja bo manj pomanjkanja, manj neuporabnih zalog in več prodaje. Vendar je napovedovanje povpraševanja lahko izziv, zlasti v primerih, ko pride do izčrpavanja zalog. Izpadi zalog skrajšajo porazdelitev povpraševanja in če se ne popravijo, dobimo napačne rezultate (Bell, 2000). Na srečo živimo v svetu, polnem podatkov, in obstajajo prava orodja, ki lahko pomagajo pri pridobivanju pomembnih vpogledov iz teh podatkov.

Kot sem že omenil, lahko natančna napoved zmanjša stroške. Stroške zmanjšuje tako, da podjetju omogoča uspešno upravljanje zalog in preprečevanje prevelikih ali premajhnih zalog. Za uspešno upravljanje zalog se podjetja zanašajo na pretekle podatke, predvsem na preteklo prodajo, ki se uporabljajo za izračun napovedi povpraševanja. Če bi bil določen izdelek, za katerega želimo napovedati prihodnje povpraševanje, vedno na zalogi in na voljo za prodajo, potem bi podjetje lahko izvajalo napovedovanje na podlagi podatkov o pretekli prodaji brez kakršnih koli sprememb in prilagoditev. Žal pa ni vedno tako. Občasno se zgodi, da se določen izdelek izprazni, zaradi česar je prodana količina manjša od količine povpraševanja (Wecker, 1978).

Na natančnost napovedi lahko poleg zalog vplivajo tudi različni drugi dejavniki. Nekatere od teh dejavnikov je težko predvideti, nekatere pa je mogoče opaziti in upoštevati pred napovedjo. Nekateri dejavniki, ki jih je mogoče upoštevati, so na primer vreme, prazniki, veliki športni dogodki, promocije, tržne kampanje itd (Ivanov et al., 2018). Na srečo se večina praznikov ali velikih športnih dogodkov vsako leto zgodi na isti datum, zato je v nekaterih primerih dovolj, da se zanesemo na pretekle podatke.

Težava, s katero se soočajo številna podjetja, je, da ne razlikujejo med napovedovanjem povpraševanja in napovedovanjem prodaje. Povpraševanje se nanaša na to, kaj želi stranka, in ne na to, kaj je bilo prodano ali odpremljeno. Spremljanje samo prodaje je v redu, dokler ima podjetje dovolj zalog (se ne sooča z izpadom zalog), in v tem primeru sta povpraševanje in prodaja enaka. Običajno napovedovanje brez upoštevanja zalog daje sliko, da je povpraševanje po izdelku manjše, kot je v resnici. Vendar to ni vedno tako. Kadar se soočamo z zalogami enega izdelka, pride do učinka substitucije. Zaradi tega učinka se bo prodalo več nadomestnega izdelka, kot je prvotno povpraševanje (Vandeput, 2023).

Kot magistrski študent oskrbne verige in logistike ter podatkovni inženir in podatkovni analitik sem ugotovil, da je ta tema zelo zanimiva in primerna zame. Moja strast do te teme izvira iz želje, da bi združila teoretično znanje, ki sem ga pridobila v zadnjih dveh letih kot študentka magistrskega študija oskrbne verige na Ekonomsko-poslovni fakulteti Univerze v Ljubljani, ter praktično znanje in izkušnje pri delu s podatki. Za namen te magistrske naloge sem uporabil R (programsko opremo), da bi prikazal, kako lahko z različnimi tehnikami podatkovne analitike izboljšamo natančnost napovedovanja povpraševanja, kako ravnati z zalogami in kako meriti natančnost napovedi. Programska oprema R vsebuje številne uporabne pakete za napovedovanje časovnih vrst, statistično modeliranje ter izrisovanje grafov, zato je priljubljena izbira številnih podatkovnih analitikov in podatkovnih znanstvenikov.

V tej magistrski nalogi sem se poglobil v temo napovedovanja povpraševanja. Razložil sem, katere so najboljše prakse, kako napovedati povpraševanje, katere statistične metode uporabiti, kaj vključiti v napoved in čemu se izogniti, kako obravnavati izpade zalog v prodajnih podatkih in kako primerjati rezultate različnih metod napovedovanja. To magistrsko delo je sestavljeno iz teoretične raziskave in praktičnega preverjanja. Teoretični del zajema opisno razlago napovedovanja povpraševanja in različnih metod napovedovanja z uporabo obstoječe literature. Za praktični prikaz tega vprašanja sem uporabil nabor podatkov, ustvarjen s programom Python. Nabor podatkov je podrobneje opisan v nadaljevanju diplomske naloge.

Pri analizi sem uporabil dve metodi napovedovanja. Na podlagi nabora podatkov, ki predstavljajo podatke o prodaji z naključnimi izpadi zalog, ki trajajo od štiri do sedem dni v 25 obdobjih, želim napovedati povpraševanje. Metodi napovedovanja, uporabljeni za ta poseben problem in niz podatkov, sta trojno eksponentno glajenje (Holt-Winters) in metoda Auto ARIMA. Prva metoda je bila uporabljena zaradi značilnosti trenda in sezonskosti podatkovnih nizov. Čeprav so značilnosti trenda in sezonskosti podatkovnega niza nakazovale uporabo metode SARIMA, sem se za drugo metodo odločil za metodo Auto ARIMA zaradi njenih prednosti v primerjavi z izključno metodo SARIMA. Auto

ARIMA je namreč priročna funkcija v paketu za napovedovanje, ki samodejno določi najboljše parametre ARIMA za dano časovno vrsto s kombinacijo testov enotnega korena, minimizacije informacijskega kriterija Akaike (AIC) in največje verjetnostne ocene. Funkcija Auto ARIMA je seveda izbrala metodo SARIMA, ne da bi mi bilo treba ročno izbrati najbolj optimalne parametre.

Seveda sem moral pred izvajanjem teh metod napovedovanja zapolniti vrzeli v podatkih o prodaji z različnimi metodami aproksimacije. Te se delijo predvsem na metodo povprečja zadnjih znanih obdobij, dve metodi interpolacije in algoritme imputiranja časovnih vrst z uporabo paketa ImputTS v programski opremi R. Na koncu so primerjani rezultati in podan zaključek.

Na podlagi zgoraj opisanega dela sem lahko odgovoril na glavno raziskovalno vprašanje:

- Kako napovedati povpraševanje na podlagi podatkov o prodaji v primeru izpadov zalog?

Poleg tega sem lahko odgovoril na naslednja raziskovalna podvprašanja:

- Kako približno oceniti povpraševanje med obdobji izčrpavanja zalog?
- Kako izmeriti natančnost napovedi z ustreznimi statističnimi merami in kako se izboljša natančnost napovedi povpraševanja, ko se prilagodi za zaloge?

Vsa koda, napisana v Pythonu in R, je na voljo tudi na mojem osebnem GitHubu. Različne kode s posebnimi težavami bodo ločene v posameznih datotekah, tako da je lažje najti, kar se išče. Kode so predstavljene v dodatkih in so dostopne na GitHubu prek te povezave.

Celoten nabor podatkov je na voljo na mojem osebnem profilu Kaggle na tej povezavi. Povezava do mojega profila in s tem tudi nabora podatkov je na voljo tudi na GitHubu v datoteki readme.md.

Iskreno verjamem, da bi moja magistrska naloga lahko prispevala številnim podjetjem kot smernica za najboljše prakse pri napovedovanju povpraševanja, pa tudi kot odgovor na napačne rezultate njihove trenutne napovedi. V prihodnosti nameravam izkoristiti svoje strokovno znanje s področja upravljanja oskrbovalne verige in podatkovne analitike ter še naprej pomembno prispevati k temu področju.