INDUSTRIAL ENGINEERING AND MANAGEMENT, MSC.

FORECASTING DEMAND AND CREATING AN INVENTORY POLICY FOR ECORUS HOME

MASTER THESIS

B. SMEEKES

S1695169

26-04-2024

UNIVERSITY OF TWENTE.

Management Summary

Ecorus, a pioneering entity in the area of sustainable energy solutions, is committed to fostering a greener world through the widespread adoption of solar power. As a leader in this field since its establishment in 2014, Ecorus strives to contribute positively to society by delivering solar park projects that align with its mission for sustainability. This research will be conducted for one of their markets, Ecorus Home. Despite its proactive stance in the renewable energy sector, Ecorus Home faces challenges in accurately forecasting demand and optimizing inventory management to meet the increasing demand. The main research question is defined as: "How can an accurate forecast and an inventory management policy be created for Ecorus to match their capacity to the increasing demand?"

Current Situation

Ecorus Home predominantly imports its stock from overseas suppliers, with lead times averaging around six weeks. However, this necessitates timely ordering to ensure adequate inventory levels. Out of the five main categories in which the SKUs are divided, only three categories have usable historical data, which are meters, invertors, and panels. After analyzing the historical data there are 8 SKUs that will be included in this research. Due to the short SKU lifetime and the often replacement of SKUs by newer product types, it could be beneficial to forecast based on product type.

Forecasting Demand

The SKUs are classified by using their strength of trend and strength of seasonality. Each SKU and the product groups are forecasted using 7 different methods, which are Naïve, Seasonal Naïve, Moving Averages, Simple Exponential Smoothing, Holt's, Holt-Winters and ARIMA. These first two methods result in one forecast per method, which is used as a benchmark for the other forecasting methods. For the latter five methods, several combinations of parameters are analyzed, which resulted in 135 different forecasts for each SKU and the product groups. For each SKU and the product groups the RMSE, bias, and MAE% are calculated to determine which method and parameters provide the most accurate forecasts. The study identifies SES and ARIMA as consistently superior methods, offering adaptable solutions capable of capturing diverse demand patterns. Despite its popularity, the Holt-Winters method generates in general the least accurate forecasts. When selecting the most accurate method per SKU with corresponding parameters, an average increase in the accuracy of 21% can be obtained compared to selecting the overall best forecasting method with its parameters.

Inventory Policy

Using the most accurate forecasts, an (s, Q) policy was created for each SKU and the product groups. In this policy the safety factor, the reorder-point, and the safety stock per SKU were determined. Implementing these parameters resulted in an average fillrate of 91%. Furthermore, a sensitivity analysis regarding the effect of the safety factor and the order quantity on the total costs as well as to the fillrate is performed. By systematically adjusting parameters such as safety stock levels, reorder points, and order quantities, decision-makers can explore the impact of variations on key performance metrics. Through iterative refinement of inventory policies based on sensitivity analysis results, Ecorus can strike an optimal balance between cost efficiency and service level attainment, thereby ensuring operational agility in response to evolving market dynamics and demand patterns.

Recommendations

It is recommended that Ecorus Home implements SES and ARIMA forecasting methods per SKU, with a keen focus on parameter optimization to fine-tune inventory policies. Expanding forecasting methods to other departments and aligning orders across departments could mitigate variability and streamline supply chain operations. The Excel tool developed for this research proves valuable and should be continuously utilized, with parameters adjusted based on insights gleaned to optimize forecasting accuracy and inventory turnover.

Limitations and Future Research

Despite notable findings, challenges such as limited historical data availability and computational constraints are acknowledged. Future research should concentrate on enhancing data integration, exploring alternative inventory policies, and developing methodologies for disaggregating forecasts at a more granular level.

Table of Contents

List of Symbols

a	level
A	fixed order costs in €/year
b	trend per period
F_t	seasonal factor in period t
C_t	cyclical variation in period t
ϵ_t	irregular fluctuations in period t
D	demand in units/year
$G_u(k)$	a special function of the unit normal (mean 0, standard deviation 1) variable
k	safety factor
L	replenishment lead time in years
$p_{u\geq(k)}$	probability that a unit normal (mean 0, standard deviation 1) variable takes on a value of k or larger
p	unit price €/unit
Q	prespecified order quantity in units
r	inventory carrying charge, in €/€/year.
s	order point in units
S	order-up-to-level in units
SS	safety stock in units
v	unit variable cost in €/unit
$\hat{x}_L$	forecast or expected demand over a replenishment lead time in units
σ_L	standard deviation of errors of forecasts over a replenishment lead time in units
B_1	cost per stockout occasion in €
B_2	fractional charge per unit short in €
z	undershoot in units
$E(t)$	average transaction size in units
D_t	demand in period t in units
$e_{t-1,1}$	error of the forecast made at period $t-1$ for period t in units
x_t	realized demand in period t
$\bar{X}$	average demand in units
n	number of periods

List of Abbreviations

ARIMA	Autoregressive Integrated Moving Average
ARMA	Autoregressive-Moving Average
CI	Confidence Interval
CoV	Coefficient of Variation
CW	Central Warehouse
EOQ	Economic Order Quantity
EPA	Euro Pallet
GBM	Gradient Boosting Methods
IP	Inventory Position
KPI	Key Performance Indicator
LW	Local Warehouse
MA	Moving Average
MAD	Mean Average Deviation
MAE	Mean Absolute Error
MAPE	Mean Absolute Percentage Error
ML	Machine Learning
MSE	Mean Squared Error
NPV	Net Present Value
OH	On-hand
PI	Prediction Interval
PV	Photovoltaics
RMSE	Root Mean Squared Error
SES	Simple Exponential Smoothing
SKU	Stock Keeping Unit
SS	Safety Stock

Chapter 1 Introduction

In this chapter, the research problem and the context will be described. This will be done by first describing the company in Section 1.1 and in Section 1.2 the existing problems of the company are analyzed. Next, in Section 1.3, the scope of the research will be determined and in Section 1.4 some essential concepts will be explained. Next, in Section 1.5, the corresponding research questions are stated and their relation to the methodology is explained, which is combined with the methodology in Section 1.6 to create a reading guide. Furthermore, in Section 1.7, the deliverables of this research are listed and finally, in Section 1.8, it is explained how the validity and reliability of this research can be assured.

1.1 Company Description

Supply Value is a consultancy firm located in Zeist (Supply Value, 2022). They specialize in the field of network optimization, considering Procurement, Supply Chain & Operations, and Digital. Their clients operate in different industries, such as Mobility & Logistics, Public Sector, Health, Fast Moving Consumer Goods & Retail, and High Tech Industry. This research is motivated by Ecorus, one of the clients of Supply Value. Ecorus develops solar park projects that provide added value for society (Ecorus, 2022). Ecorus was founded in 2014, making them a precursor in this market. Ecorus wants to contribute to creating a sustainable world, and they are convinced that this is possible with the large-scale roll-out of solar power.

They operate in multiple markets, which they specify as Home, Energy, Business, and Development. The Home department mainly specializes in placing solar panels on social renting houses, which will be the focus of this research. Especially the panels installed at social rental houses contribute greatly to the public interest, in financial and environmental ways. Since Ecorus wants to keep placing their solar projects on social rental housing, they are in need of assistance with some problems they face.

1.2 Problem Identification

In this section, the problems at Ecorus will be investigated. In Section 1.2.1 a description of the situation and the problems at Ecorus will be given. In Section 1.2.2, the problem cluster of the mentioned problems will be displayed. Finally, in Section 1.2.3, the core problem of this research will be selected and explained.

1.2.1. Problem Context

Over the last year, a lot of negative events occurred in the world, one of the most recent being the war in Ukraine. This has led to an enormous increase in gas and energy prices. While some sources state that this is only a short-term spike (Tollefson, 2022), others state that the effects will be continuing for a long period (Isidore, 2022).

Preventing an energy crisis is one of the most crucial issues of the 21st century (Devabhaktuni et al., 2013). Many of the resources that future generations will require, such as oil or coal, are currently being depleted. Renewable energy sources such as wind, solar, biomass, and hydropower can provide a sustainable energy supply. Even though the

Netherlands is already a leading country when it comes to solar power (Sévano, 2022), there is still a lot to gain. Studies have shown that solar energy can supply about 1,000 times more energy than the demand worldwide, but only 0.02% of its full potential was used in 2013 (Devabhaktuni et al., 2013). In 2020, 5.2% of the total electricity production in the European Union (EU) came from solar energy. Based on current market trends, it has the potential to meet up to 20% of the electricity demand of the EU by 2040 (European Commission, 2022).

The Dutch Government recently introduced a netting arrangement for solar energy (RTL Nieuws, 2023; van Gastel, 2023). The netting arrangement is for small consumers who feed electricity into the grid, which is mainly used by small consumers who generate their own electricity with solar panels. The Dutch Government has also introduced a package of measures specifically for social renting houses. By 2028 all houses should have energy label D or higher and by 2050 all houses should have zero CO_2 emissions (Rijksoverheid, 2022b). Furthermore, the landlord taxes have been removed to create a 2 billion euros budget to improve the sustainability of the houses (Rijksoverheid, 2022a). These aforementioned reasons lead to an increase in the demand for solar panels (RTL Nieuws, 2023; van Gastel, 2023). The expected growth in photovoltaic (PV) capacity in The Netherlands can be seen in Figure 1.

In general, just before a new calendar year starts, Ecorus makes agreements with each of the housing associations regarding the number of projects that are aimed to be installed per month for the next calendar year. Since these predictions have a timespan between one and twelve months, not every prediction is accurate, which is especially the case for the predictions at the end of the year. Therefore, each month the predictions are made more exactly for the coming months. Since the projects on each of the houses are unique, different amounts and types of parts are used for individual projects. An example of this is the number and type of panels that are installed on the roof. Based on data from Ecorus' previous projects, it can be calculated how many items of each part are on average required per house. This however is hugely influenced by a project being for a high-rise or a low-rise building. Currently, it is assumed that for every housing association, this distribution will be the same as in previous years. However, this could lead to wrong forecasts. If this ratio deviates from the expectation, the number of items required per part could differ enormously, since often different parts are required for a high-rise building compared to a low-rise building.

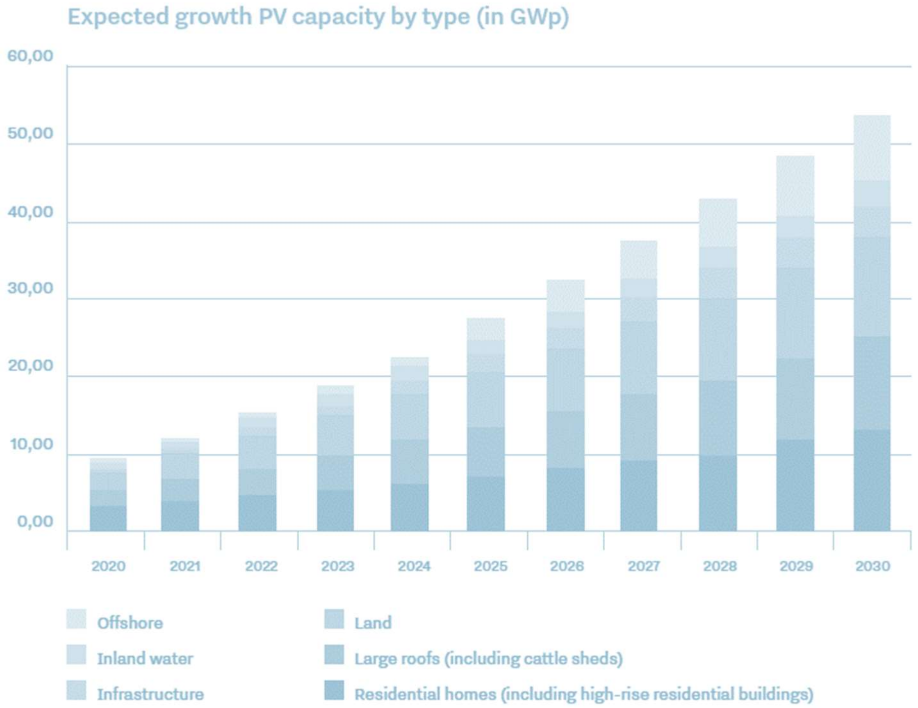

Figure 1: Expected Growth PV Capacity in The Netherlands (Netherlands Enterprise Agency (RVO), 2020)

There are several reasons the number of installed projects deviates from the number of projects that were supposed to be installed. Some of these reasons are listed below:

- Tenants of the social renting houses only decide if they want solar panels after the housing association estimated the number of projects.

- The infeasibility to install projects in periods in which it freezes, since it is impossible to install solar panels on the roof during this time.
- The Bouwvak, which is the period in which construction industries have their summer break. However, this period is known in advance.
- The roofs of houses turn out to be infeasible to install panels on, this could be for example due to the construction of the roof or because of shade from nearby trees or houses.
- The work preparation is not done properly.
- The installation depends on whether the right materials are present.
- Capacity limitations of the instalment partners. If these partners do not have the right personnel at their disposal, it is not possible to have the projects installed. However, it is important to notice that the lack of availability of installers is not caused by poor planning, but rather by the overstrained market (UWV, 2022). Since demand is increasing, Ecorus is also actively searching for new instalment companies to start cooperating with.

At this moment, it is not clear to Ecorus how often, for what reason, and by how many projects the agreements are not met. Insights into these deviations could help Ecorus to create better predictions on how many projects will be realized and therefore how many parts will be required.

Currently, Ecorus has one Central Warehouse (CW) and several Local Warehouses (LW) they can use to store solar panels and parts. This central warehouse is outsourced by Ecorus. The other local warehouses are owned by the instalment partners. Using the warehouses of the instalment partners to store the parts may lower the time it requires to install a project, but this requires timely planning and might increase total holding costs due to the loss of inventory pooling benefits (Axsäter, 2006). Furthermore, these local warehouses have less storage capacity compared to the central warehouse. Most of the parts Ecorus uses for the instalment of solar panels are imported from overseas. This significantly reduces the costs, but this induces high lead times. Furthermore, the suppliers promise a certain delivery time, but it is unclear to which extent these promised lead times are complied with. When the parts are not ordered timely or are frequently delivered later than agreed, there is a risk Ecorus becomes out of stock, which would lead to lost sales.

Furthermore, the central warehouse has moved from Noord-Holland to Friesland recently. One disadvantage of this new location for the warehouse is that the costs for a rush delivery to the local warehouses have increased enormously, due to the average distance to the instalment partners. This makes it important to deliver the parts to the local warehouses timely. Previously, multiple trips per week could be arranged from the central warehouse to the local warehouses, but since the central warehouse has moved, this would ideally be only once a week. At this moment, there is no inventory policy for the central warehouse and the local warehouses. It is desired to have a model which includes and balances different objectives, such as minimizing the stock levels or minimizing the risk of lost sales by maximizing the fraction of demand that can be satisfied directly from stock.

With an increasing number of installed projects, the total number of breakdowns in projects also increases. Since it is the responsibility of Ecorus to repair them, it is desirable to fix the breakdowns as soon as possible, to prevent high compensation costs to the tenants. Currently, the breakdowns are logged, but this information is not yet used. Ecorus could benefit from an analysis of these past breakdowns, in order to forecast the number of expected breakdowns. These forecasts can be used such that parts can be ordered in advance instead of when the breakdown has occurred.

1.2.2. Problem Cluster

To get an overview of the problems in a process, it is useful to create a problem cluster (Heerkens & van Winden, 2021). The relations between the problems stated in Section 1.2.1 are shown in the problem cluster in Figure 2.

1.2.3. Core Problem

Once the problem cluster is created, a core problem can be chosen. Since a core problem can never be a problem that you cannot influence (Heerkens & van Winden, 2021), the problems in the white boxes cannot be chosen as the core problem. This results in a number of possible core problems, as shown in the light blue boxes in Figure 2. These four influenceable problems that can be derived are "No insights in the accuracy of the agreements made with the housing associations", "Absence of an inventory policy considering the central warehouse", "No insights in the delivery performance of suppliers", and "No forecasts for the number of parts required to repair the failures". Since a forecast

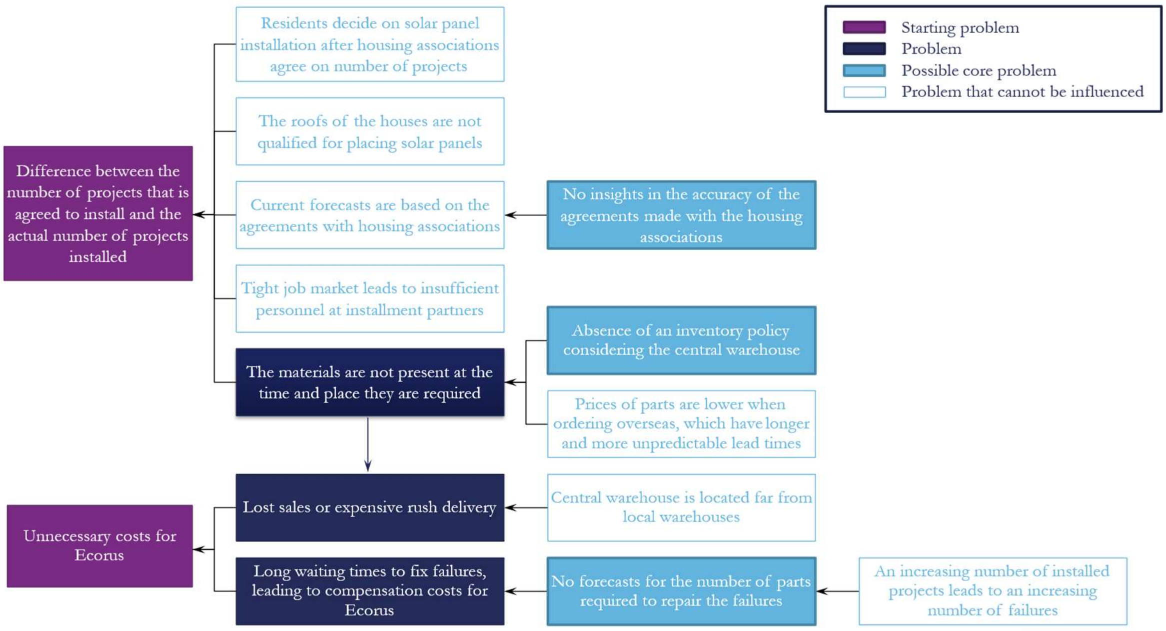

Figure 2: Problem Cluster displaying the relations between the problems at Ecorus

is often used in an inventory management policy, these problems are highly related. Therefore, the core problem for this research is a combination of two subsequent problems, and has been chosen as:

"The absence of an accurate forecast and an inventory management policy limits Ecorus in matching their capacity to the increasing demand."

The other possible core problems, regarding the absence of insights in the accuracy of the agreements with housing associations and the absence of insights in the delivery performance of the suppliers, will be shortly addressed later, but these are not directly part of the core problem.

1.3 Scope of Research

In this section, the scope of the research will be discussed. As mentioned before, Ecorus has several markets in which they operate. The focus of this research will be on the Home market, which places solar projects on social rental houses. Furthermore, the research will be based on the data from January 2021 until August 2023, since this is the timespan of the available data.

As mentioned before, Ecorus makes use of a number of local warehouses, which are owned by the instalment partners. This research will mainly focus on the central warehouse they currently have and will have a scalability which can include local warehouses.

The focus of this research will be on the operational and tactical level. Therefore, no strategic choices will be discussed such as the locations of the local warehouses and the products that are offered.

1.4 Theoretical Framework

In the problem context, previously discussed in Section 1.2, multiple times the concept 'forecasts' has been discussed, which are the basis of all supply chain planning. For all processes the level of activity must be anticipated, for example in production, capacity, inventory, or transportation. To forecast demand, first the factors which influence demand must be identified, and then a relationship between these factors must be ascertained (Chopra & Meindl, 2015). Knowledgeable factors related to the demand forecast include for example past demand, planned advertising or planned price discounts (Chopra & Meindl, 2015).

An inventory management policy is a set of guidelines and procedures that is used to manage inventory effectively. The purpose of an inventory management policy is to ensure that the right amount of inventory is on hand to meet demand, minimize inventory carrying costs, and avoid stockouts or overstocking. This involves addressing a complex array of factors both internal and external to the organization (Axsäter, 2006). Examples of this are the lead time, which is the time required to deliver a replenishment order, or the ordering costs. When determining the appropriate stock levels for a specific item at a particular location, three critical considerations must be taken into account: the frequency of inventory assessments, the point at which a replenishment order should be placed, and the size of those orders (Axsäter, 2006).

Section 1.2 also mentions the concept 'parts'. These parts, which are specific units of stock, are called Stock Keeping Units (SKUs). A SKU is defined as an item of stock that is completely specified considering its characteristics, such as function, style, size, color, and location (Silver et al., 2016). In the case of Ecorus, an example is a certain solar panel of a given size, color, and watt peak. If a solar panel has a different color, this would result in the items being two different SKUs. It is important to note that this classification system can result in high correlations between the demand for two SKUs. This is caused by the possibility to switch between SKUs with the same functionality if another SKU is out of stock, for example if a solar panel is only in stock in white and not in black. Given the distinct features of each SKU, it is desired to develop a customized forecast and inventory management policy for each SKU. However, considering the substantial number of SKUs that are frequently maintained in inventory, it may be feasible to categorize similar SKUs based on their shared characteristics, which can streamline the management process (Silver et al., 2016).

1.5 Research Questions

In this section, the research question and the subquestions that have been formulated to support and answer the main research question will be discussed. The core problem as stated in Section 1.2.3 can be translated into the following research question:

"How can an accurate forecast and an inventory management policy be created for Ecorus to match their capacity to the increasing demand?"

Subquestions have been formulated to support and answer the main research question in a structured manner. The chapters in this research and the questions they cover will be described in this section.

In Chapter 2, the focus will lie on the exploration of the current situation at Ecorus. This is done by examining the characteristics of the SKUs, the demand, and the central warehouse. The research questions that will be answered in this chapter are:

1. What is the current status of Ecorus with respect to forecasting their demand and an inventory management policy?

i. What are the characteristics of the SKUs Ecorus provides?
ii. What are the characteristics of the demand of Ecorus?
iii. What are the characteristics of the central warehouse?

Next, in Chapter 3, insights in forecasting and inventory management policies are provided based on the available methods in literature. First the focus will lie on the literature on forecasting. This will be done by first exploring the methods to classify the different SKUs, since it might be beneficial to categorize similar SKUs to optimize the management process, as described in Section 1.4. Next, the literature on forecasting demand will be examined since. To conclude the research question regarding forecasting, the methods available to analyze the performance of forecasts are explored.

Additionally in Chapter 3, the third research question regarding the inventory management policies will be discussed. This is done by first focusing on the relevant inventory control systems and their corresponding parameters. This includes the methods to determine the frequency of inventory assessments, the point at which a replenishment order should be placed, and the size of those orders, as mentioned in Section 1.4. Since there are numerous internal and external factors influencing the inventory management policies, as briefly mentioned in Section 1.4, the methods known to calculate the safety stocks are analyzed and the modelling of the uncertainty in lead times is examined. To conclude this part, the methods known to evaluate an inventory management policy are addressed.

2. What is known in literature about forecasting?

i. Which methods can be used to classify the SKUs?
ii. Which methods can be used to forecast demand?
iii. Which methods can be used to measure the performance of a forecast?

3. What is known in literature about inventory management policies?

i. Which inventory control systems are known in literature that would be applicable for Ecorus?
ii. What are the required parameters for these policies?
iii. Which methods are known to calculate the required safety stocks?
iv. What is known about modelling uncertainty in the lead times?
v. Which methods are known to evaluate an inventory management policy?

Once the relevant literature has been analyzed, the fourth research question focusses on the methodology of creating a forecast and inventory policy, which will be done in Chapter 4.

4. How can forecasts and inventory models be created for Ecorus?

i. What is the best applicable method to forecast demand at Ecorus?
ii. Which inventory control systems can be used by Ecorus?
iii. What are the corresponding parameters of Ecorus for these inventory control systems?

Once it is clear how the forecasts and the inventory model need to be created, the results can be analyzed. This will be done in Chapter 5. This is first done by analyzing the forecasting results, and next creating the inventory models based on the forecasting results. The corresponding research question and its subquestions are listed below.

5. What are the results of the forecasts and the inventory models?

i. What is the performance of the forecasts?
ii. Which insights can be gained regarding the forecasts?
iii. What is the performance of the inventory management policy for Ecorus?
iv. Which insights can be gained regarding this inventory management policy?

Once the results are analyzed, the sixth research question will be addressed in Chapter 5. This question concentrates on providing recommendations for the implementation of the forecasting tool and the inventory management policy.

6. What recommendations can be given to Ecorus regarding the solution and its implementation?

Furthermore in Chapter 6, this research is concluded. The final research question in this study focuses on identifying the limitations of this research and providing recommendations for future research that could benefit Ecorus.

7. What limitations are there in this research and what recommendations can be given to Ecorus for future research?

1.6 Reading Guide

Table 1 summarizes the research questions as described in Section 1.5 and the chapter in which each of these phases and questions are discussed. Due to confidentiality, required because of the concurring market, the suppliers and the warehouse are anonymized. In this research, keys will be used to identify the suppliers, indicated by "S.". The decoding key can be found in Appendix A.

Chapter	Research question
1. Introduction	NA
2. Current Situation	1
3. Literature Review	2, 3
4. Design and Development	4
5. Results	5
6. Evaluation	6, 7

Table 1: Reading Guide

1.7 Deliverables

This research will deliver the following:

- An analysis of the current situation.
- A report discussing the design and the development of the forecasting and inventory policy tool.
- A manual on how to use the tool.
- A forecasting and inventory policy tool which incorporates, and balances multiple objectives created in Excel.
- A generalized version of this forecasting and inventory tool with the aim of extending Supply Value their building blocks to help their clients better and quicker.

1.8 Validity and Reliability

To ensure reliable and valid research, several actions need to be taken. If research is repeated, the same results should be achieved. If this is the case, the research is reliable. If a research measures what was intended to be measured, it is valid research. Validity can be categorized into three types, which internal validity, construct validity and external validity (Heerkens, 2017).

Internal validity is the soundness of the research design, which means if it is measured what is aimed to be measured. Threats to internal validity are for example an unrepresentative sample, incorrect statistical methods, unwanted artificiality, extreme events, and unreliable measuring instruments. To ensure internal validity, several methods can be used, such as cross-validation, which involves dividing the data in subsets and using each of them to validate the results obtained from the other set, or a sensitivity analysis, which involves evaluating the robustness of the results by changing the inputs, assumptions, or methods used in the research (Salkind, 2010).

Construct validity concerns the operationalization of the variables, if the correct indicators are used for the correct variables, and if the variables and indicators are based on the scientific body of knowledge. Threats to construct validity are for example the incorrect use of literature, inadequate definition of variables, one dimensional operationalization, and one-dimensional measuring.

External validity concerns the generalizability of the research, especially if the research is valid outside the specific research environment. Threats to external validity are for example stimuli being unique for the research context, the population being unique, the environment being unique, and the time being unique. As this research focusses on Ecorus specifically, there are some threats to the generalizability. This is caused by the research environment which is unique for Ecorus. Due to confidentiality caused by the concurring market, the suppliers, the installation partners, and the clients of Ecorus are anonymized. The decoding key is available to Ecorus, however, this will not be published. To ensure external validity, several methods can be used, examples are a collaboration with stakeholders, such as users or experts, to ensure that the prototypes or designs are grounded in real-world needs and requirements, or user testing in naturalistic settings to assess the effectiveness and usability in actual use (Salkind, 2010). Furthermore, to perform valid research, the steps taken in this research are carefully documented.

Chapter 2 Current Situation

In this chapter the first research question, "What is the current status of Ecorus with respect to demand forecasting?" will be answered. This will be done by first describing the characteristics of the SKUs of Ecorus in Section 2.1. Next, the relevant characteristics of the demand are described in Section 2.2. Furthermore, in Section 2.3, the characteristics of the central warehouse will be described. This chapter will be concluded in Section 2.4.

2.1 Characteristics of the SKUs

The SKUs that Ecorus has can be roughly divided into five categories, which are panels, inverters, small materials, mounting materials, and meters. The SKUs of each of these categories will be described separately. An overview of the suppliers of each of these SKU categories can be found in Table 2.

The first SKU category is the solar panels. Almost every panel Ecorus Home buys is manufactured in China and transported from there by barge to the harbor of Rotterdam. Ecorus Home has several purchase contracts which are negotiated and established by the procurement department. The logistics coordinator arranges the operational purchasing. The main suppliers of solar panels are S.6 and S.8, when solar panels are bought these suppliers take care of the sea transport. The lead times of these panels are usually around 6 weeks. When these panels arrive in the port of Rotterdam, they are first tested by S.10. Next, Ecorus Home arranges the transport in collaboration with their central warehouse and the warehouse in the port of Rotterdam. In case the preferred suppliers do not have any panels available, panels can be bought at suppliers located in The Netherlands, but this is at a much higher price. However, if panels are urgently required, these local suppliers can deliver with a short lead time of a few days. Two main types of solar panels can be distinguished, panels with a white back sheet or panels with a black back sheet.

Another SKU category is inverters. Two types of inverters are used for the installations at Ecorus Home: string and micro inverters. Within the house, string inverters are used, which are available in various shapes and sizes. On the roofs micro inverters are used, which can be roughly divided into two types, inverters suitable for 2 panels and inverters suitable for 4 panels. Since these inverters can be used for multiple panels, fewer inverters are required compared to panels. If inverters are urgently required, they can be delivered with a short lead time of a few days when ordered at suppliers located in the Netherlands.

SKU group	Group code	Main supplier	Back-up supplier	Price increase at back-up supplier	Number of SKUs in group	Number of SKUs used in 2023	Number of SKUs with > 150 demand in '21- '23
Solar Panels	003	S.6, S.8	S.7, S.9	25-30%	31	13	8
Inverters	002			30%	27	19	13
Small materials	006	S.2, S.4, S.5	N/A	N/A	22	N/A	N/A
Constructions	007		N/A	N/A	11	N/A	N/A
Meters	001	S.2, S.5	N/A	N/A	11	7	3

Table 2: Suppliers and Details of the Different SKU Groups

In the category small materials items such as extension cables, extension pieces, and end pieces are included. In the category constructions, all parts required to attach the solar panels to the roof are included, such as roof hooks, bases, and ballast plates. Even though these materials are essential for the installation of projects, there is no historical demand data of this category and therefore these materials will not be further discussed. The last category are meters, which are used inside the house to measure the current and are ordered through S.3. This supplier also sends an order to S.1 for the SIM cards that will be mounted in the meters. Once these SIM cards have been placed, the meters will be sent to the central warehouse.

It is important to note that most SKUs used by Ecorus have a limited lifespan due to the development of better parts, which is in particular seen within the solar panels, the inverters and the meters. The solar industry is an industry that continuously develops better solar panels and converters, making the older products obsolete.

2.2 Characteristics of the Demand

In this section, the second subquestion "What are the characteristics of the demand of Ecorus?" will be discussed. Currently, Ecorus has contracts with 13 housing associations in the Netherlands. Per housing association, the type and the size of the contract differs. The type of contracts can roughly be divided into three categories, which are 1) contracts with a structural demand, 2) contracts with a differing demand, and 3) private customers.

The first category, contracts with a structural demand, are the contracts with the most predictable demand. These types of contracts usually have a fixed and constant number of projects that will be installed per year. The second type of contracts, contracts with a differing demand, are more flexible. There are several reasons the demand could differ per year. One of those reasons is that only on the first of January, the financial year is closed, so housing associations only know at that moment how much budget is available in the next financial year. Since the beginning of 2021, there have been projects installed for 19 different housing associations, ranging from only a few projects to several hundreds of projects per year per housing association. For each of the projects, several SKUs are required. However, which SKUs are required is determined once the laying plans are drawn by Ecorus. Data on the SKUs used for each individual project is available from January 2021 to August 2023.Before further analyzing this data, outliers will be removed. As mentioned in Section 2.1, the SKUs can have a limited lifespan. This also results in SKUs that were not used at the beginning of 2021. If this is applicable to a SKU, these datapoints have been removed. Furthermore, observations of zero demand have been ignored in weeks Ecorus is closed, such as the week between the Christmas Holidays and New Year.

It is important to further analyze the demand of the individual SKUs. As previously mentioned in Section 2.2, there are numerous SKUs that have not been used in 2023 anymore. Furthermore, there are also SKUs which were introduced in 2023. For these latter SKUs the forecasts could be made, however, they cannot be validated. Since data is available from January 2021 to August 2023 and 25% of data is required to measure the performance of the forecast, historical demand data is required prior to April 2023. Furthermore, the total demand per SKU can be determined. SKUs with less than 150 units of demand since 2021 have been excluded. Appendix B gives an overview of all SKUs and if their data is sufficient to be included in the forecasts. After selecting the appropriate SKUs, the other SKUs will be further excluded from this analysis. This leads to 8 remaining SKUs, of which 2 are meters, 4 are invertors, and 2 are panels. The demand of these SKUs per product group is shown in Figure 3 to Figure 5. In the remainder of this research these SKUs will be referred to as SKU 1 to SKU 8, Appendix B shows the actual SKU number.

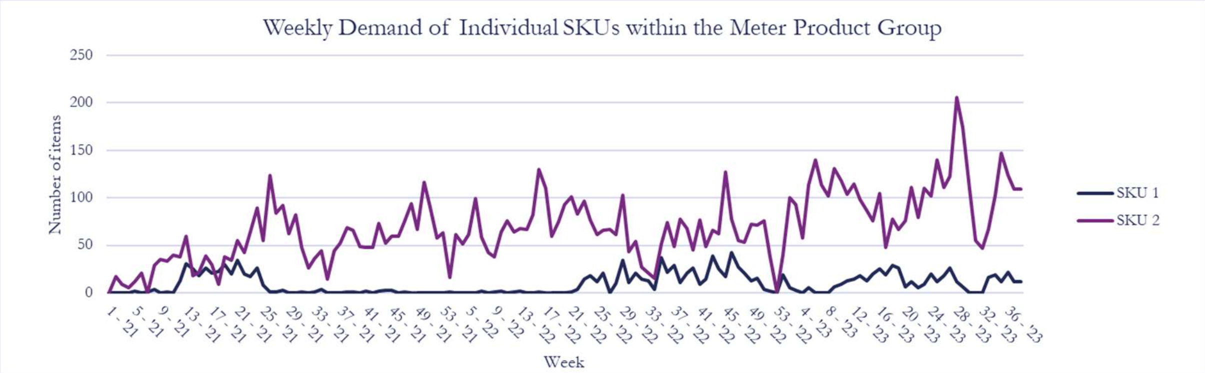

Figure 3: Weekly Demand of the remaining SKUs in the Meter Product Group

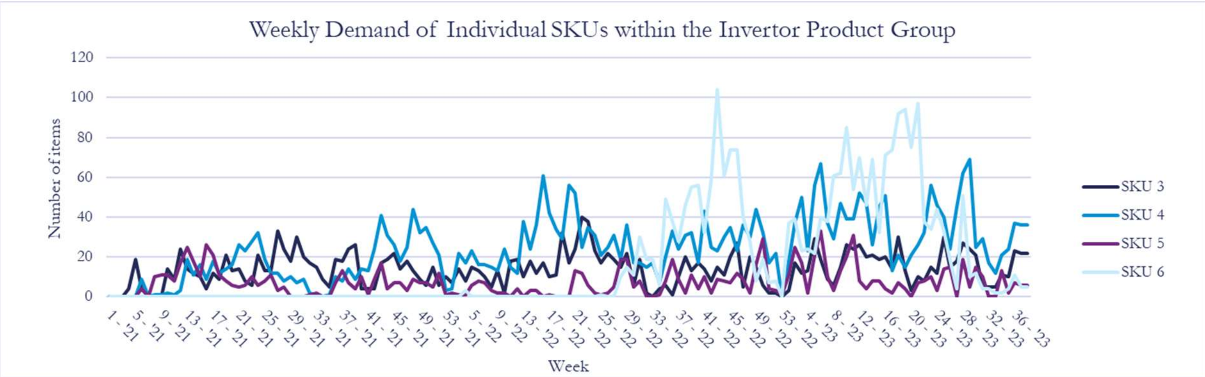

Figure 4: Weekly Demand of the remaining SKUs in the Invertor Product Group

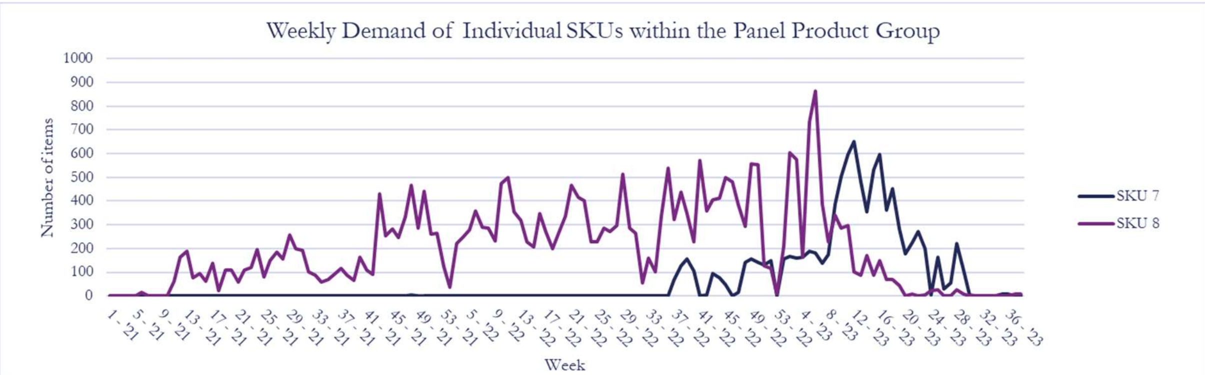

Figure 5: Weekly Demand of the remaining SKUs in the Panel Product Group

As described in Section 2.1, SKUs have a limited lifespan and therefore SKUs are often replaced by a more recent designed type. This could make it beneficial to look at the SKU groups instead of separate SKUs. To provide insight in the demand, Figure 6 shows an overview of the number of items used per week grouped per type of product. SKUs that have been excluded on individual level are included in this combined demand.

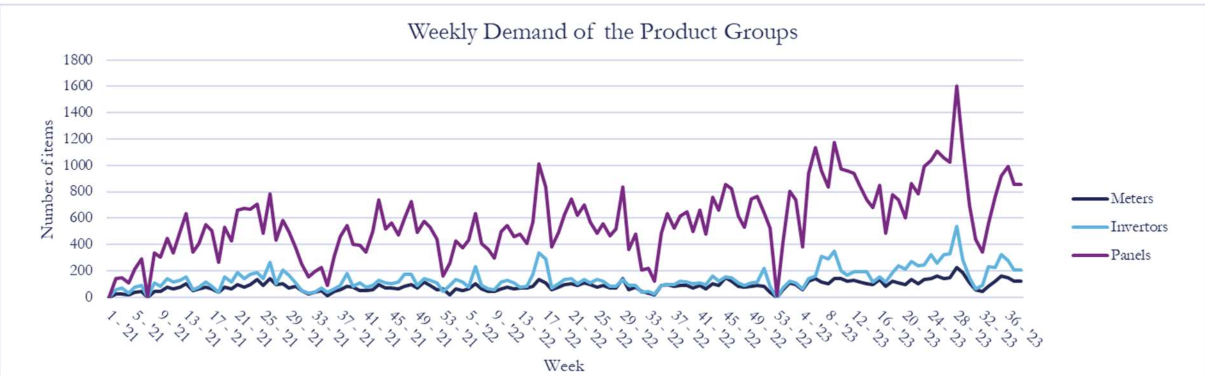

Figure 6: Combined Demand for the Product Groups

As mentioned in Section 2.1, the panels can be divided in two main groups. The demand of each of these groups is shown in Figure 7. Regarding the last 20 weeks of demand for these panels, there has been a huge shift in the type of panel that has been used. When comparing these two types of panels and the demand of the panels in total, as shown in Figure 6, it is assumed that more accurate forecasting results are obtained when the solar panels are considered as one product group.

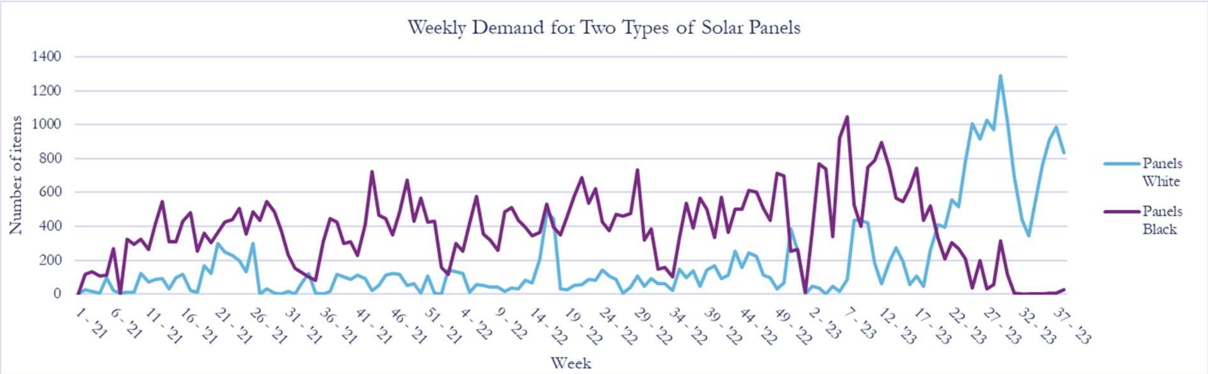

Figure 7: Combined Demand for the Two Main Categories of the Solar Panels

2.3 Characteristics Central Warehouse

In this section, the fourth subquestion "What are the characteristics of the central warehouse?" will be discussed. The warehouse is outsourced by Ecorus, therefore most costs incurred are variable to the number of parts that are stored. When storing parts three main types of costs are considered, which are fixed and variable costs for entry, variable costs for storage, and fixed and variable costs for dispatch. The fixed costs for entry and dispatch are independent of the size of the order and are incurred for checking and administering the order. The variable costs for entry and dispatch are based on the number of pallets. For storing, a variable fee for the number of pallets per week is incurred. There is a distinction made between regular euro pallets (EPA) and solar panel pallets with dimensions. Currently, a maximum of 200 EPAs and 200 solar panel pallets can be stored at the warehouse. However, additional storage might be possible in consultation with and in agreement by the warehouse owner.

2.4 Conclusion

Usually, most SKUs are imported from overseas with lead times around 6 weeks, however, if Ecorus is urgently in need of parts they can also be bought at suppliers located in The Netherlands. In this case, the lead time will only be several days but the price per part increases by around 30%. Out of the five main categories in which the SKUs are divided, only three categories have usable historical data, which are meters, invertors, and panels. After removing the outliers, the historical data is analyzed, which results in 8 SKUs that have sufficient data to be included further in this research. Due to the short SKU lifetime and the often replacement of SKUs by newer product types, it could be beneficial to forecast based on product type.

Chapter 3 Literature Review

In this chapter, a literature review will be performed. First, the focus will lie on the literature on forecasting. The second research question "What is known in literature about forecasting?" will be answered. This will be done by first examining the methods to classify SKUs in Section 3.1. Furthermore, in Section 3.2, the methods which can be used to forecast demand will be analyzed and in Section 3.3 the methods to determine the performance of forecasts will be discussed.

Next, the third research question "What is known in literature about inventory management policies?" will be answered. This will be done by first looking into inventory control systems that would be applicable for Ecorus in Section 3.4 and their corresponding parameters in Section 3.5. Next, in Section 3.7, the known literature about modelling uncertainty in the lead times will be examined and in Section 3.6 the focus will lie on the known methods to calculate the safety stocks. This chapter will be concluded in Section 3.8, where an answer will be provided to the second and third research question.

3.1 Classification of SKUs

Considering the substantial number of SKUs that are frequently maintained in inventory, it may be feasible to categorize similar SKUs based on their shared characteristics, which can streamline the management process (Silver et al., 2016). Therefore, the first subquestion "Which methods can be used to classify the SKUs?" will be answered in this section.

The ABC classification, or Pareto analysis, is a commonly used method for classifying SKUs based on the attention required (Waters, 2003). The ABC classification considers the combination of SKU volume and sales value to group SKUs into three categories, each corresponding to the degree of attention and care that should be given (Axsäter, 2006; Silver et al., 2016; Waters, 2003). In this research only SKUs that require much attention are considered, which makes the ABC classification not sufficient. Furthermore, since traditional ABC methods only classify parts based on their sales value (Zhang et al., 2001), it is not applicable in situations where parts are not sold individually.

The XYZ analysis is a dynamic extension of the ABC analysis that divides SKUs into classes based on the predictability and regularity of demand. The designated SKU classes enable informed decision-making when determining the quantity to order, which is crucial for products with short life cycles (Nowotynska, 2013). The XYZ analysis has several limitations, one of which is the difficulty in categorizing new products (Dhoka & Choudary, 2013). Since the SKUs discussed in this research all have a limited lifespan and are often replaced by newer generation SKUs, the XYZ analysis would not be a suitable method to classify SKUs.

In the domain of intermittent demand forecasting, Williams (1984) introduced an analytical technique to categorize demand into smooth, slow-moving, or intermittent patterns. This method involved breaking down the variance in lead-time demand into its fundamental elements, which include transaction variability, fluctuations in demand size, and variations in lead time. Syntetos et al. (2005) introduced a method which allows for the classification of time-series based on two key factors: the average demand interval length and the coefficient of variation in demand sizes when demand does occur.

Another possibility to classify SKUs is based on the strength of their trend and the strength of their seasonality. To determine this, the time-series can be decomposed and consequently the strength of the trend and the strength of seasonality can be determined, described by F_T and F_S (Hyndman & Athanasopoulos, 2018). This can be done by using

Equations (1) – (2), in which the trend component is described by T_t, the seasonal component is described by S_t, and R_t is the residual after decomposition.

$$F_T = \max\left\{0; 1 - \frac{var(R_t)}{var(T_t + R_t)}\right\} \tag{1}$$

$$F_S = \max\left\{0; 1 - \frac{var(R_t)}{var(S_t + R_t)}\right\} \tag{2}$$

3.2 Forecasting Demand

In this section, the second subquestion "Which methods can be used to forecast demand?" will be answered. Methods to forecast demand will be discussed in Section 3.2.1. Next, in Section 3.2.2 the advantages and disadvantages of these methods to forecast demand will be discussed. As described in Section 2.2, it might be beneficial to forecast the SKUs as a SKU group, therefore Section 3.2.3. focusses on the required adaptations to forecast demand on SKU group level.

3.2.1. Methods to Forecast Demand

Forecasting can be achieved through a combination of both statistical forecasting, which involves extrapolating past observations, and expert judgment, which considers future events and circumstances (Silver et al., 2016). For items or groups of related items for which accurate historical demand data is available, techniques known as time-series forecasting models can be utilized. These models use the demand history for an item or group of items to construct a forecast for the future demand of that item or group. This approach is particularly useful when the demand for an item or group of items exhibits temporal patterns that can be modeled and projected into the future. When time-series models are appropriately selected and calibrated, they have the potential to generate high-quality forecasts (García et al., 2010). According to Axsäter (2006), demand can be modelled in three ways; constant demand, trend model, and a trend-seasonal model, where the latter is the more general model since it can describe more demand patterns than the constant demand model. However, the use of a more general model should be avoided unless there is evidence that it will provide certain advantages (Axsäter, 2006). In Section 3.2.1.1 to Section 3.2.1.6 several categories of forecasting methods are discussed.

3.2.1.1 Naïve and Seasonal Naïve

The Naïve method is a relatively straightforward forecasting technique that relies solely on the most recent observation to predict the future value of a time-series for the current period. This method is commonly employed as a statistical benchmark in forecasting practice (Hyndman & Athanasopoulos, 2018; Makridakis et al., 2020; Paldino et al., 2021). A disadvantage of this method is that the forecasts for multiple periods ahead will not differ in value compared to the forecast of one period ahead. A variation of this approach is known as the seasonal Naïve method (sNaïve), specifically useful for data with pronounced seasonality patterns. In this method, the forecast is set to be identical to the last observed value from the same season (Hyndman & Athanasopoulos, 2018; Makridakis et al., 2020). While one of the advantages of the sNaïve Method is that seasonality can be included, a disadvantage is that no trend can be taken into consideration. While both methods cannot capture the trend in the time-series, both methods will be used to forecast with the purpose of benchmarking the other forecasting methods.

3.2.1.2 Moving Average

When it is assumed the underlying demand structure for a SKU is constant and has no observable trend or seasonality, the moving average (MA) method can be used to forecast (Chopra & Meindl, 2015). The MA is a simple and widely used technique in which the average of the N most recent observations is used to estimate the level (Axsäter, 2006). The simple moving average (SMA) uses equally assigned weights, and the weighted moving average (WMA) assigns weights to the different periods. Furthermore, N can be adjusted to find the best model for each SKU. One of the main advantages of the MA is its simplicity, which makes it easy to understand and implement. Additionally, it does not require a large amount of historical data. However, it also has some limitations such as not considering any underlying patterns or trends in the data, which can result in poor forecasting accuracy. Furthermore, it is not able to adapt to changes in the underlying data when N is large, meaning it may not perform well when there are significant changes in the data trend. Additionally, it can be sensitive to outliers and may not be able to handle large fluctuations in the data. Furthermore, similarly to the Naïve method, the forecasts for multiple periods are equal to the forecast one period ahead, making it most suitable for forecasting only a few periods into the future (Kapgate, 2014). Since most SKUs exhibit a trend as described in Section 2.2, the moving average procedure might therefore not be suitable to forecast.

3.2.1.3 Exponential Smoothing Methods

Exponential smoothing is a variation of the moving average method, and similarly to the moving average method, the Simple Exponential Smoothing (SES) method is based on a constant demand model (Chopra & Meindl, 2015). This method creates a forecast for the next period by using a weighted average of the historical data, where the weights decrease exponentially as the data becomes older. However, since most SKUs exhibit a trend as described in Section 2.2, the Simple Exponential Smoothing method might not be suitable to forecast demand at Ecorus. Exponential smoothing can be extended which allows for the inclusion of a trend component in the forecast, often referred to as Holt's method (Holt, 2004). The basic underlying model is the model with a trend. The initial values of the level and trend in Holt's method can be determined through a least-squares regression on historical data, with the aim of minimizing the difference between the historical data and the predicted values using the chosen level and trend values. Like Exponential Smoothing, Holt's method requires to update the estimates for the trend and level.

In many organizations, there are numerous demand patterns of SKUs that display significant seasonality (Silver et al., 2016). For a trend-seasonal model the procedure referred to as Winters can be used, which is useful in aggregate, medium-range forecasts (Winters, 1960). This method is a natural extension of the Holt procedure for a trend model and is intuitively appealing (Silver et al., 2016). To handle the presence of trend and seasonal factors in demand patterns, the initialization process becomes more complex since one must properly separate the trend and seasonal effects in historical data. A commonly used initialization method is the ratio to moving average procedure, which effectively addresses changes in underlying trends during the historical period and tends to eliminate cyclical effects (Silver et al., 2016). Ideally, from a statistical perspective, several seasons worth of data should be used to accurately separate out the trend and seasonal effects in the historical data. However, using too much history can lead to the risk of the seasonal pattern having changed, making early data no longer representative of current or future conditions. A compromise is to use a minimum of four complete seasons, with five or six seasons used with care (Silver et al., 2016). As previously described in Section 2.2, demand data of Ecorus is available over the period January 2021 – December 2022. Since this would not cover four complete seasons, including seasonality might not give accurate forecasting results.

Exponential smoothing is widely used in practice due to its simplicity and effectiveness. Exponential smoothing techniques have demonstrated remarkable performance in forecasting competitions, often outperforming more complex and sophisticated approaches (Billah et al., 2006). Some of the advantages of exponential smoothing include its ability to adapt to changes in the underlying data trend, its ability to incorporate the uncertainty associated with the forecast, and its computational simplicity. However, there are also some disadvantages of exponential smoothing. One of the disadvantages is that methods may not be appropriate for certain types of time-series data, such as data with multiple seasonal patterns or high levels of volatility.

3.2.1.4 Regression

When creating a medium-term time-series forecast, a possible method is regression analysis. There are several types of regression models, such as linear, logistic, polynomial, and quantile regression models. Regression models can include seasonal factors and multiple predictor variables. A least-squares criterion is used to estimate values-of the parameters (Silver et al., 2016). The advantages of regression procedure are that it allows for the modeling of complex relationships between variables, it can be used to make predictions about future values of the dependent variable, and it can be used to identify the most important predictor variables. Some of the main disadvantages of the linear regression procedure are that it can be sensitive to outliers and other types of noise in the data.

3.2.1.5 Autoregressive Integrated Moving Average

More complex models compared to simple regression models are the Autoregressive-Moving Average (ARMA) and Autoregressive Integrated Moving Average (ARIMA) class of models, introduced by Jenkins et al. (1994). The ARMA and ARIMA class of models combine regression with moving averages and represent the demand in the current period as a weighted sum of past demands and unpredictable random components (Anderson, 1977). The models are time-series forecasting models that make very few assumptions and are very flexible (Jenkins et al., 1994). These attributes have contributed to their remarkable effectiveness in both academic research and practical industrial applications (Ho & Xie, 1998; Ramos et al., 2015). The ARIMA models are an extension of the ARMA models, which uses differencing to handle time-series which are non-stationary, making the ARIMA models more common since time-series are often non-stationary. The Box-Jenkins approach proposes a broad class of underlying statistical models for understanding demand patterns, as well as a procedure for selecting an appropriate model based on historical data. These models offer the potential for improved forecasting accuracy, particularly at the medium-term, aggregate level. One of the main advantages of the ARIMA approach is its ability to handle complex data patterns that may not be captured by simpler forecasting methods such as moving averages. Furthermore, similarly to previously mentioned regression methods, the approach allows for the estimation of multiple parameters, which can lead to more accurate forecasts. However, one of the main disadvantages of the ARIMA approach is that it requires a large amount of historical data to estimate the necessary parameters.

3.2.1.6 Machine learning

Recently, there has been a shift towards using Machine Learning (ML) methods for demand forecasting, as an alternative to traditional statistical methods. The primary advantage of ML methods lies in their utilization of nonlinear algorithms which can learn from historical data through trial and error, and improve their performance over time, without making many assumptions about the data generation process (Spiliotis et al., 2022).

Decision tree learning is a supervised learning approach employed in statistics, data mining, and machine learning. It involves the use of a classification or regression decision tree as a predictive model to make conclusions about a set of observations. Decision trees are highly favored in machine learning given their simplicity and intelligibility and are among the most widely used machine learning algorithms (Rokach & Maimon, 2005).

The Random Forest (RF) is a machine learning technique commonly employed to solve regression and classification problems by using an ensemble learning approach that combines multiple classifiers to address complex problems (Breiman, 2001). The RF algorithm is composed of numerous decision trees, and the forest generated by this algorithm is trained via bagging or bootstrap aggregating. The RF algorithm determines the output based on the forecasts of the decision trees and calculates the prediction by averaging or taking the mean of the output from multiple trees. The accuracy of the outcome can be increased by increasing the number of trees used in the model. RF algorithms overcome the limitations of decision tree algorithms, reduce overfitting of datasets, and enhance precision without requiring excessive configuration in packages.

Gradient Boosting Methods (GBMs) are machine learning techniques used in regression and classification problems. When a decision tree is the weak learner, the resulting algorithm is called gradient-boosted trees, which usually outperforms Random Forest. In the recent forecasting competition M5 'Accuracy', the majority of well performing methods used LightGBM (Makridakis et al., 2022). LightGBM is a machine learning algorithm designed for nonlinear regression using gradient boosted trees (Ke et al., 2017). GBMs offer several advantages over other ML alternatives in the context of forecasting tasks, such as its ability to effectively handle multiple features, encompassing past sales and explanatory variables of various types, such as numeric, binary, and categorical. Furthermore, it requires the optimization of a relatively small set of parameters, including the learning rate, number of iterations, maximum number of bins, number of estimators, and loss functions.

By using ML methods in demand forecasting it has been managed to outperform the Naïve benchmark, the sNaïve benchmark, and the top performing benchmark ES_bu (Makridakis et al., 2022), making ML methods very promising for forecasting demand. However, due to the requirement of large samples to effectively capture the dynamics and inter-connections of demand volume and inter-demand intervals, ML methods are very data-intensive (Spiliotis et al., 2022). As a result, ML methods are not suitable for demand series with few observations or demand series with many observations of zero. Furthermore, since one of the requirements of this research is the ability to create forecasts in Excel, which is not possible with aforementioned ML methods, Machine Learning methods are further excluded.

3.2.2. Advantages and Disadvantages

The previously mentioned methods to forecast demand have their advantages and disadvantages, which are listed in in Table 3.

	Naïve method	Moving Average	Exp. Smoothing	Regression	ARIMA
Limited amount of historical data required	Yes	Yes	Yes	No	No
Possibility to include trend	No	No	Yes	Yes	Yes
Possibility to include seasonality	Yes	Yes	Yes	Yes	Yes
Optimizes parameters to reduce error terms	No	Yes	Yes	Yes	Yes
Different forecasts for more than 1 period ahead forecasts	No	No	Yes	Yes	Yes
Model complex relation between variables	No	No	No	Yes	Yes
Suitable for non-stationary time-series	No	No	Yes	No	Yes
Computationally intensive	No	No	No	No	No

Table 3: Advantages and Disadvantages of the Methods to Forecast Demand

3.2.3. Forecasting Demand on SKU Group Level

Time-series data can often be broken down into subcategories based on specific attributes of interest, creating a hierarchical structure (Hyndman & Athanasopoulos, 2018). In numerous scenarios, there exist multiple hierarchically structured time-series that can be grouped together at various levels based on factors like products, geography, or other attributes (Hyndman et al., 2011). More complex structures can emerge when attributes are both nested and crossed, as in the case of a manufacturer interested in sales by both product type and geographic division (Hyndman & Athanasopoulos, 2018). In grouped time-series, the structure does not naturally disaggregate in a unique hierarchical manner. Typically, hierarchical time-series are forecasted using either a bottom-up or a top-down approach. In the case of forecasting as a SKU group and later disaggregating this forecast into a forecast on SKU level, a top-down approach is applicable. Most literature on aggregating and disaggregating time-series forecast is respective to the geographical location (Athanasopoulos et al., 2009; Silveira Netto et al., 2023). A simple example of a hierarchical SKU structure is shown in Figure 8. At the top of the hierarchy lies the "Total", representing the most aggregated level of the data. The t-th observation of the Total series is denoted as y_t for $t = 1,\dots,T$. The Total series is further broken down into two series, which are subsequently divided into three and two series, respectively, at the lowest level of the hierarchy. Below the top level, observations for the series corresponding to node j at time t are represented as $y_{j,t}$.

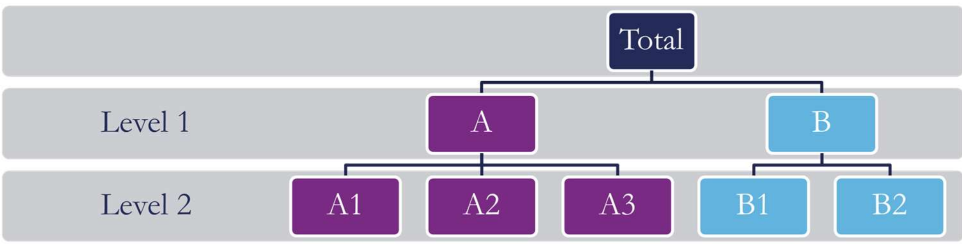

Figure 8: Example of a Simple Hierarchical SKU Structure

For any time t, the observations at the bottom level of the hierarchy will sum to the observations of the series above as in the example: $y_t = y_{A1,t} + y_{A2,t} + y_{A3,t} + y_{B1,t} + y_{B2,t}$. In top-down approaches, the initial step is to create forecasts for the total y_t. Subsequently, these forecasts are distributed down the hierarchy. The set $p_1, \dots, p_m$ represents the disaggregation proportions that dictate how the forecasts for the total series should be allocated to generate forecasts for each series at the lowest level of the structure. For instance, when applying proportions $p_1, \dots p_5$ to the hierarchy illustrated in Figure 8, the resulting forecasts are as shown in Equation (3).

$$\tilde{y}_{A1,t} = p_1\hat{y}_t, \quad \tilde{y}_{A2,t} = p_2\hat{y}_t, \quad \tilde{y}_{A3,t} = p_3\hat{y}_t, \quad \tilde{y}_{B1,t} = p_4\hat{y}_t, \quad \tilde{y}_{B2,t} = p_5\hat{y}_t \tag{3}$$

Various methods exist for determining the disaggregation proportions. Gross & Sohl (1990) provide insights into several approaches for selecting these proportions. Within the top-down approach, the most prevalent methods for determining the disaggregation proportions are rooted in historical data proportions: average historical proportions and (Hyndman & Athanasopoulos, 2018). For average historical proportions, each proportion represents the mean of historical proportions pertaining to the bottom-level series throughout the entire specified period, in relation to the overall aggregate. This relationship is mathematically defined in Equation (4). On the other hand, when considering proportions of the historical averages, each proportion captures the historical average value of the lower-level series concerning the average total value of the overall aggregate, which is shown in Equation (5).

$$p_x = \frac{1}{T}\sum_{t=1}^{T}\frac{y_{X,t}}{y_t} \tag{4}$$

$$p_x = \sum_{t=1}^{T}\frac{y_{X,t}}{T} \Big/ \sum_{t=1}^{T}\frac{y_t}{T} \tag{5}$$

3.3 Performance of Forecasts

In this section, the subquestion "Which methods can be used to measure the performance of a forecast?" will be answered. Regardless of the method used to generate forecasts, evaluating their quality is of considerable importance for two reasons. Firstly, analyzing the past performance of a forecasting process provides valuable information for making informed decisions about resource allocation based on a description of future demand, and secondly, monitoring the performance of a forecasting process over time may reveal opportunities for improvement (Silver et al., 2016). There are several methods that can be used to measure the performance of a forecast, also known as Key Performance Indicators (KPIs), which will be discussed in the remainder of this section.

Bias, shown in Equation (6): It is possible for a forecasting model to exhibit low bias despite a lack of precision, as positive errors of an item may counterbalance negative errors of this item (Vandeput, 2021). Merely evaluating the bias of the forecast may not provide sufficient insight into its accuracy. However, a highly biased forecast can serve as a warning signal of potential flaws in the model.

Mean Absolute Deviation (MAD), shown in Equation (7): The MAD can be beneficial when evaluating the accuracy of predictions in terms of the same unit as the original time-series (Khair et al., 2017).

Mean Squared Error (MSE), shown in Equation (8): Several algorithms utilize MSE as a performance metric due to its computational efficiency and manipulability (Vandeput, 2021). However, MSE does not accurately reflect the scale of the original error as it involves squaring the error values, resulting in a KPI that cannot be directly related to the original demand scale.

Mean Absolute Percentage Error (MAPE), shown in Equation (9): While MAPE is one of the most commonly used KPIs to measure forecast accuracy, it is recognized as having limited accuracy as an indicator (Vandeput, 2021). The calculation of MAPE involves dividing each error by the respective demand value, leading to potential biases. Specifically, high errors occurring during periods of low demand will have a significant impact on MAPE (Chicco et al., 2021; Gross & Sohl, 1990). Thus, an optimization of MAPE may result in a suboptimal forecast that is prone to undershooting demand.

Mean Absolute Error (MAE), shown in Equation (10): One of the primary challenges with utilizing this KPI is its lack of scalability relative to the average demand (Vandeput, 2021). To solve this, the MAE can be divided by the average demand, which results in the KPI MAE% (Vandeput, 2021).

Root Mean Squared Error (RMSE), shown in Equation (11): When comparing the RMSE to the MAE, both are not scaled to demand, but the RMSE metric places a greater emphasis on the magnitude of the most significant errors, while the MAE gives equal weight to each error in the calculation of the metric (Vandeput, 2021). However, the RMSE is considered a more appropriate metric for evaluating the performance of a model when the error distribution is predicted to be Gaussian, compared to the MAE (Chai & Draxler, 2014). Similarly to MAE%, the RMSE can also be normalized to the NRMSE (Chai & Draxler, 2014).

Coefficient of Determination (R^2), shown in Equation (13): This method calculates the proportion of the variance in the dependent variable that is predictable from the independent variables used in the model (Chicco et al., 2021). To determine R^2, the average demand is required, of which the determination can be found in Equation (12). Despite their usefulness, the aforementioned metrics share a common limitation, since their values can range from zero to positive infinity, making it difficult to interpret the performance of a regression model based on a single score. The R^2 provides a more informative and accurate assessment of the model's performance, as it yields a high score only if most of the ground truth elements have been correctly predicted (Chicco et al., 2021). Furthermore, R^2 is less prone to interpretational limitations compared to the other methods. For this reason, it is recommended to use R^2 as the standard metric for evaluating regression analyses (Chicco et al., 2021).

$$bias = \frac{1}{n}\sum\nolimits_{t=1}^{n}\left(x_t - \hat{x}_{t-1,t}\right) = \frac{1}{n}\sum\nolimits_{t=1}^{n}\left(e_{t-1,1}\right) \tag{6}$$

$$MAD\frac{1}{n}\sum\nolimits_{t=1}^{n}\left|x_t - \hat{x}_{t-1,t}\right| = \frac{1}{n}\sum\nolimits_{t=1}^{n}\left|e_{t-1,t}\right| \tag{7}$$

$$MSE = \frac{1}{n}\sum\nolimits_{t=1}^{n}\left(x_t - \hat{x}_{t-1,t}\right)^2 \frac{1}{n}\sum\nolimits_{t=1}^{n}\left(e_{t-1,1}\right)^2 \tag{8}$$

$$MAPE = \frac{1}{n}\sum\nolimits_{t=1}^{n}\left|\frac{x_t - \hat{x}_{t-1,t}}{x_t}\right| = \frac{1}{n}\sum\nolimits_{t=1}^{n}\left|\frac{e_{t-1,t}}{x_t}\right| \tag{9}$$

$$MAE = \frac{1}{n}\sum\nolimits_{t=1}^{n}\left|x_t - \hat{x}_{t-1,t}\right| = \frac{1}{n}\sum\nolimits_{t=1}^{n}\left|e_{t-1,t}\right| \tag{10}$$

$$RMSE = \sqrt{\frac{1}{n}\sum\nolimits_{t=1}^{n}\left(x_t - \hat{x}_{t-1,t}\right)^2} = \sqrt{\frac{1}{n}\sum\nolimits_{t=1}^{n}\left(e_{t-1,1}\right)^2} \tag{11}$$

$$\bar{X} = \frac{1}{n}\sum\nolimits_{t=1}^{n}x_t \tag{12}$$

$$R^2 = 1 - \frac{\sum_{t=1}^{n}\left(\hat{x}_{t-1,t} - x_t\right)^2}{\sum_{t=1}^{n}(\bar{X} - x_t)^2} \tag{13}$$

These aforementioned methods can be used to evaluate the performance of a forecast and help to determine which method is most accurate and reliable for a given dataset and scenario. However, the reduction of a set of error values into a single numerical value using aforementioned KPIs results in a loss of significant information. The ideal statistical metrics should not only provide a measure of performance, but also offer a representation of the error distribution (Chai & Draxler, 2014). The main characteristics are listed in Table 4.

	Bias	MAD	MSE	MAPE	MAE	MAE%	RMSE	R^2
Same scale as original demand	Yes	Yes	No	No	No	No	No	No
Positive and negative errors cancel out	Yes	No	No	No	No	No	No	No
Indicates if there is an over- or underestimation	Yes	No	No	No	No	No	No	No
Emphasis on the most significant error	No	No	Yes	No	No	No	Yes	Yes
Resulting value on a fixed scale	No	No	No	No	No	No	No	Yes

Table 4: Main Characteristics of the Methods to Evaluate Forecasts

3.4 Inventory Control Systems

In this section, the topic of discussion will be the question "Which inventory control systems are known in literature that would be applicable for Ecorus?". First, in Section 3.4.1 background information on the topic is given. Section 3.4.2 will describe the physical operation of the five most common types of inventory control systems and their advantages and disadvantages.

3.4.1. Background Information Inventory Control Systems

The primary objective of an inventory control system is to address three key challenges or concerns, as proposed by Silver et al. (2016):

1. The frequency at which the inventory status should be assessed.
2. The optimal timing for placing a replenishment order.
3. The appropriate order quantity for replenishment.

In some control systems a review interval is used, which determines how often the inventory status should be checked. Continuous review, where stock levels are always known, is not always necessary or feasible in practice. Instead, transactions reporting, where stock levels are updated with each transaction, is commonly used (Silver et al., 2016). Periodic review, where inventory status is checked at set intervals, is also widely used.

Coordination of replenishments may be more attractive in periodic review, where items in a coordinated group can be given the same review interval. This allows for a reasonable prediction of workload on staff. Continuous review can be more expensive in terms of reviewing costs and errors, especially for fast-moving items. However, data collection systems have significantly reduced these costs and errors. Periodic review may be more effective in detecting spoilage or pilferage of slow-moving items because it forces an occasional review, whereas transactions reporting does not (Silver et al., 2016).

The major advantage of continuous review is that it requires less safety stock, hence lower carrying costs, to provide the same level of customer service (Silver et al., 2016). This is caused by the period over which safety protection is required being longer under periodic review, and the stock level could drop significantly between review instants without any reordering action being possible in the interim. Overall, physical constraints or limits on time often determine the review interval used in practice.

3.4.2. Common Types of Inventory Control Systems

The (s, Q) system is a type of inventory control system in which a fixed quantity, Q, is ordered whenever the inventory position drops to or below the reorder point, s. Unlike inventory control systems that use net stock to trigger an order, the (s, Q) system uses inventory position. This ensures the on-order stock is taken into account, to ensure that an order is placed only when necessary. The (s, Q) system ensures that the replenishment of inventory is on order when the inventory position reaches the reorder point, eliminating the need for frequent checks and orders. The advantages of the (s, Q) system include its simplicity and the predictability of production requirements for suppliers. However, the primary disadvantage of the (s, Q) system in its unmodified form is its inability to handle situations where individual transactions are large. If the transaction that triggers the replenishment is large enough, then a replenishment of size Q may not raise the inventory position above the reorder point. In such cases, an integer multiple of Q could be ordered instead to ensure that the inventory position is above the reorder point. Overall, the (s, Q) system offers a practical approach to inventory control, particularly for products with consistent demand patterns. Its simplicity and predictability make it a popular choice for businesses with limited resources for inventory management.

The (s, S) system is a continuous review inventory control method in which a replenishment order is made when the inventory position drops to or below the order point s. Unlike the (s, Q) system, the (s, S) system uses a variable replenishment quantity that is ordered to raise the inventory position to the order-up-to-level S. Since the SKUs considered in this research are ordered in fixed batch sizes, the (s, S) system is not suitable.

The periodic-review system (R, S), also referred to as the replenishment cycle system, is commonly used in companies without advanced computer control, particularly for items ordered from the same supplier or requiring resource sharing (Silver et al., 2016). In this system, at every review instant, which occurs every R units of time, inventory is ordered to raise the inventory position to the level S. Compared to order point systems, the periodic-review system is preferred in coordinating the replenishment of related items, especially when ordering from overseas and needing to fill a shipping container to control shipping costs. Like the (s, S) system, the (R, S) system is not suitable due to the required ordering in batches of the SKUs considered in this research.

The (R, s, S) system combines the features of the (s, S) and (R, S) systems, however, this (R, s, S) system does not use fixed batch sizes. An adaption of this model is the (R, s, Q) inventory model, which has undergone extensive research in the past few decades. This inventory policy involves monitoring the inventory position at regular intervals of R time units to decide on replenishment (Janssen et al., 1998). If the inventory position falls below s an integral multiple of Q is ordered to bring the inventory position to a value between s and $s + Q$. Several heuristic and optimal techniques have been developed to determine the optimal values of the control parameters R, s, and Q (Janssen et al., 1998). This system is closely related to the (R, s, S) system, but instead of ordering up to a level, a fixed quantity is ordered, which is required for the SKUs considered.

The two possible inventory models therefore are the (s, Q) and the (R, s, Q) inventory model. The difference between these models is using a continuous review period versus using a periodic-review system. A periodic-review system allows for placing orders for multiple SKUs simultaneously, which results in discounts at the suppliers which provide multiple SKUs, by taking advantage of bulk discounts and reducing transaction costs. Furthermore, continuous review systems require real-time data integration. Therefore, the periodic-review model with a fixed-order quantity is considered the most suitable inventory policy.

3.5 Parameters Inventory Control Systems

In this section, the research question "What are the required parameters for these policies?" will be answered. Independent of the chosen policy, there are several costs that should be taken into consideration, which are holding costs, ordering costs, and the unit price.

By holding stock, several direct and indirect costs arise and should be taken into consideration. When keeping SKUs in stock an opportunity cost for capital is tied up in inventory (Axsäter, 2006). These costs are often closely related to the return on an alternative investment and are referred to as cost of capital. An approach to determine this is to use the Net Present Value (NPV) (Axsäter, 2006), or the Weighted-Average Cost of Capital (WACC), which takes into account the return on equity of a firm (Brealey et al., 2022; Chopra & Meindl, 2015). The holding costs of an inventory consist also of obsolescence costs, handling costs, occupancy costs, and miscellaneous costs (Chopra & Meindl, 2015). While it is ideal to analyze these costs for each SKU separately, it is often assumed that a single value for holding costs applies to all SKUs (Silver et al., 2016). Holding costs can be expressed as a total value per unit, or as a percentage of the unit price (Chopra & Meindl, 2015). Ordering costs refer to all the costs associated with replenishing items and are incurred regardless of the order size, and are comprised of buyer time, transportation costs, receiving costs, and other expenses (Chopra & Meindl, 2015). The unit price or the unit value is the average price paid per unit purchased, which is the key component in the lot size decision (Chopra & Meindl, 2015). The lot size could be increased if this results in a reduction of the unit price per unit (Chopra & Meindl, 2015). Taking these costs into consideration, the total annual costs can be determined, which furthermore depends on the ordering quantity (Chopra & Meindl, 2015). When optimizing this equation, the Economic Order Quantity (EOQ) is found, as shown in Equation (14).

$$EOQ = \sqrt{\frac{2AD}{vr}} \tag{14}$$

When dealing with probabilistic demand, it is helpful to categorize inventories into the following concepts (Silver et al., 2016):

- On-hand (OH) stock: the inventory physically available on the shelf, which can never be negative. It determines whether a specific customer demand can be directly fulfilled.
- Net stock: this is calculated as on-hand stock minus backorders and can become negative if backorders exist. It is used in mathematical derivations and is an essential component of the Inventory Position.
- Inventory Position (IP): this is also known as available stock and is calculated as on-hand stock plus on-order minus backorders minus committed. On-order stock refers to inventory that has been requisitioned but not

yet received by the stocking point. Committed quantity is required when this stock cannot be used for other purposes in the short run. Inventory position is a critical quantity in determining when to replenish stock.

- Safety Stock (SS): this is the average level of net stock just before a replenishment arrives. A positive SS provides a buffer against larger-than-average demand during the effective replenishment lead time. The appropriate method of calculation and numerical value of SS depends on what happens to demand during a stockout. When backorders are not permitted, net stock and thus SS are always nonnegative.

3.6 Determine Required Safety Stock

In this section, the subquestions "Which methods are known to calculate the required safety stocks?" will be discussed. In the following subsections this will be discussed for the different suitable systems described in Section 3.4.2.

3.6.1. Parameters of a (s, Q) System

The safety factor, denoted as k, is recognized for its significance in determining the appropriate value of the reorder point s. The calculation of k results in a corresponding value of s, as determined by the two interrelated equations shown in Equation (15) and (16). It is important to include that the system has a manual override option to enable the user to adjust the reorder point to account for factors not considered in the model (Silver et al., 2016). The total yearly expected costs are comprised of ordering, holding, and shortage costs, as shown in Equation (17). It is possible to calculate the total costs for a given k, however, this may be infeasible if many SKUs are kept in stock. To overcome this, Silver et al. (2016) have developed two decision rules using either the shortage costs B_1, or the specified fractional charge per unit short, indicated by B_2 (Silver et al., 2016). As described in Section 2.1, when stock-outs occur, SKUs can be bought at a back-up supplier in the Netherlands instead of overseas. These price increases and thus additional costs can be seen as the costs per unit short, therefore the decision rule that utilizes B_2 should be used to determine k, which is shown in Equation (18). If this statement in Equation (18) is true, k should be set to the lowest allowable value, which should be determined by management. If Equation (18) is not satisfied, k can be chosen such that Equation (19) is satisfied. Several other methods exist, such as using the specified fractional charge per unit short per unit of time, the probability of no stockout per replenishment cycle, or the specified fraction of demand satisfied directly from shelf (Silver et al., 2016). Regarding Equation (19), a Normal distribution can be assumed if the demand during lead time is over 10 items (Silver et al., 2016), which is the case for all SKUs. Using Microsoft Excel, Equations (20)-(21) can be used to determine the relation between $p_{u\geq(k)}$ and k.

$$s = \hat{x}_L + SS \tag{15}$$

$$SS = k * \sigma_L \tag{16}$$

$$TC = \frac{AD}{Q} + \left(\frac{Q}{2} + k\sigma_L\right) vr + \frac{B_2 v\, \sigma_L p_{u\geq(k)} D}{Q} \tag{17}$$

$$\frac{Qr}{DB_2} > 1 \tag{18}$$

$$p_{u\geq(k)} = \frac{Qr}{DB_2} \tag{19}$$

$$p_{u\geq(k)} = 1 - \text{NORMDIST}(k) \tag{20}$$

$$k = 1 - \text{NORMSINV}(p_{u\geq(k)}) \tag{21}$$

3.6.2. Parameters of a (R, s, Q) system

To calculate the safety stock as the required parameters for an (R, s, Q) policy, the method proposed by Janssen et al. (1998) can be used. This method uses the fraction of demand fulfilled directly from the shelf, known as the fill rate or P_2 service level, as the service criterion. The assumption is made that any customer orders that cannot be satisfied directly from shelf will be backordered. The expected shortage at the beginning of a replenishment cycle is not neglected, so the approach includes the undershoot. The importance of including the undershoot lies in the considerable impact on performance levels, particularly when the probability of zero demand during the lead time is high, which often occurs considering the demand of the SKUs considered in this research. Equation (22) can be used to determine the values for s and Q.

$$\beta = 1 - \pi_{\hat{L}} \frac{E(Z_2 - s)^+ - E(Z_1 - s - Q)^+}{Q} + (1 - \pi_{\hat{L}}) \frac{E(U_2 - s)^+ - E(U_1 - s - Q)^+}{Q} \tag{22}$$

3.7 Uncertainty in Lead Times

In this section the subquestion "What is known about modelling uncertainty in the lead times?" will be answered. In Sections 3.4 and 3.5 various decision systems have been discussed, in all these systems the total demand in an interval of length $R + L$ or L is the critical variable in setting a reorder point or order-up-to-level. The decisions are primarily based on a known replenishment lead time L, with the only uncertainty being the demand rate during L or $R + L$. However, if L is not known with certainty, it is apparent that increased safety stock is required to protect against this additional uncertainty (Silver et al., 2016).

It is important to note that if the pattern of variability is known, such as seasonally varying lead times, there is no additional concern (Ammar et al., 2013). The lead time at any given calendar time is known, and the safety stock and reorder point can be adjusted accordingly. Similarly, if lead times are increasing in a known fashion due to reduced availability of raw materials, the safety stocks and reorder points should be appropriately adjusted.

In this section the possibility to incorporate stochastic lead times into inventory models is discussed. Previously, only constant lead times were considered, but in reality, the supply process is often stochastic. According to Axsäter (2006) two types of stochastic lead times can be distinguished, which are sequential deliveries independent of lead-time demand and independent lead times. The former is most common, where orders cannot cross in time, and the stochastic lead time for a certain order may depend on previous demand but not on future demand. The latter occurs when orders are served by many independent servers, and later demands may trigger orders delivered earlier than the current order. Modeling the first type of lead times can be complex and an integrated model that considers both inventory and queues may be needed. Evaluating lead time variations can also be challenging compared to demand variations, so it is often reasonable to replace a stochastic lead time with its mean, especially when dealing with independent stochastic lead times.

A straightforward yet effective technique that is helpful when it is feasible to measure the overall demand during the lead time has been developed by Lordahl & Bookbinder (1994). This method relies on the distribution-free characteristics of order statistics and directly utilizes these measurements. The primary goal is to determine the reorder point while the parameters and structure of the lead time demand distribution are unknown, and P_1 is the desired performance metric. The procedure mandates that the user keeps track of the actual demand during the lead time over a period. The following steps outline the procedure:

- Rank the observed lead time demand from least to greatest. $x(1) \leq x(2) \leq \cdots \leq x(n)$ given n observations of lead time demand.
- Let $(n + 1)P_1\ y + w$, where $0 \leq w < 1$, and y is an integer.
- If $(n + 1)P_1 > n$, set $s\ x(n)$. Otherwise, set $s\ (1 - w)x(y) + wx(y + 1)$, which results in a weighted average of two observations.

Another method involves modeling the lead time and the demand as independent random variables and requires measurements or estimates of both (Silver et al., 2016). Despite the occasional occurrence of positive correlation between high demand and long lead times due to heavy workload, and negative correlation between low demand and long lead times due the need to accumulate sufficient orders for the desired production run size, the assumption of independence is deemed a reasonable approximation (Silver et al., 2016). If it is assumed L and D are independent random variables, Equation (23) can be used to determine σ_x which can substitute σ_L in Equation (16) to determine s.

$$\sigma_x = \sqrt{E(L)\, var(D) + [E(D)]^2 var(L)} \tag{23}$$

3.8 Conclusion

In this section, the second and third research question will be answered. This will be done by first concluding the second research question "What is known in literature about forecasting?" in Section Literature on Forecasting3.8.1. Next, in Section 3.8.2, the third research question "What is known in literature about inventory management policies?" will be concluded.

3.8.1. Literature on Forecasting

Common methods to classify SKUs are the ABC and the XYZ method. However, in this research only SKUs that require much attention are considered, which makes the ABC classification not applicable. Furthermore, the ABC method is not applicable in situations where parts are not sold individually. Since the SKUs discussed in this research all have a limited lifespan and are often replaced by newer generation SKUs, the XYZ analysis would also not be a suitable method to classify SKUs. A more useful method to classify SKUs is based on the strength of their trend and the strength of their seasonality. Additionally, this classification can also guide in selecting the most appropriate forecasting model.

Based on the literature analyzed, there are several forecasting methods applicable for Ecorus. Exponential smoothing is a widely used method for forecasting time-series data and has been proven to be adaptable to different situations. Furthermore, Holt's method extends exponential smoothing to handle trends, which is important for Ecorus as trends can greatly impact demand for renewable energy solutions. Additionally, the Winters procedure can be used to include seasonality in demand patterns, which is also important for Ecorus given that demand for renewable energy solutions may vary depending on seasonal factors. Although exponential smoothing assumes independent and identically distributed forecast errors, this method is simple and adaptable, making it a suitable choice for Ecorus as they may encounter various forms of demand patterns. Ecorus may want to initialize models with both trend and seasonal factors, which can be complex but can lead to more accurate forecasts. ARIMA models have the ability to capture other patterns than seasonality, which could provide good results considering the demand patters of the SKUs Ecorus uses.

There are several methods to measure the performance of forecasts. It is recommended for Ecorus to use a combination of MAE%, RMSE, and bias to measure the performance of their forecast model. MAE is a useful performance measure as it calculates the average absolute difference between the predicted and actual failure rates, which can help Ecorus determine the level of accuracy of their forecast. However, since MAE is not scaled to demand, it is recommended to use MAE% for better comparison across SKUs. RMSE is a more appropriate metric for placing a greater emphasis on the magnitude of the most significant errors and can help Ecorus identify potential outliers in their forecast. Finally, bias is recommended as it provides insights into the over- or underestimation of the forecasts. By using a combination of these performance measures, Ecorus can get a more comprehensive evaluation of their model its performance and better understand its strengths and weaknesses. While the other performance measures discussed can be useful for evaluating different aspects of forecast accuracy, they may not provide a comprehensive evaluation of the performance of a model on their own.

3.8.2. Literature on Inventory Models

Regarding Ecorus, the most suitable policies are the (s, Q) and the (R, s, Q) policies. Since the SKUs considered in this research are ordered in fixed batches, such as full pallets, it is not feasible to order-up-to a certain level. The (R, s, Q) policy is advantageous for this type of company because it offers the flexibility to adjust the timing and quantity of orders based on changing demand patterns. The fixed batch ordering approach means that the inventory position may not need to be monitored as frequently, so R can be set accordingly. By setting a reorder point s and an order quantity Q, Ecorus can ensure that it always has enough inventory to meet demand without overstocking. Moreover, the (R, s, Q) offers a good balance between inventory holding costs and order costs. With the fixed batch ordering approach, Ecorus can take advantage of economies of scale in ordering and transportation costs and the possibility of receiving a wrong number of SKUs due to a variable order size are reduced.

In inventory management, there is a possibility of incorporating stochastic lead times into the models. If the pattern of variability is known, such as seasonally varying lead times, there is no additional concern. Two types of stochastic lead times can be distinguished, sequential deliveries independent of lead-time demand and independent lead times.

Modeling the first type of lead times can be complex, and an integrated model that considers both inventory and queues may be needed. Evaluating lead time variations can also be challenging compared to demand variations, so it is often reasonable to replace a stochastic lead time with its mean. A technique which relies on the distribution-free characteristics of order statistics and directly utilizes measurements of the overall demand during the lead time to determine the reorder point can be used. Another method involves modeling the lead time and the demand as independent random variables and requires measurements or estimates of both.

Chapter 4 Design and Development

In this chapter, a forecasting and inventory policy will be designed for Ecorus. In this chapter, the focus will lie on the research question "How can forecasts and inventory models be created for Ecorus?" will be answered. This will be done by first focusing on forecasting the demand in Section 4.1, which will answer the sub question "What is the best applicable method to forecast demand at Ecorus?". Next, the focus will lie on the development of an inventory policy in Section 4.2, which will answer the subquestions "Which inventory control systems can be used by Ecorus?" and "What are the corresponding parameters of Ecorus for these inventory control systems?". An inventory management policy is a set of guidelines and procedures that is used to manage inventory effectively. The purpose of an inventory management policy is to ensure that the right amount of inventory is on hand to meet demand, minimize inventory carrying costs, and avoid stockouts or overstocking. This involves addressing a complex array of factors both internal and external to the organization (Axsäter, 2006). When determining the appropriate stock levels for a specific item at a particular location, three critical considerations must be taken into account: the frequency of inventory assessments, the point at which a replenishment order should be placed, and the size of those orders (Axsäter, 2006). Finally, this chapter will be concluded in Section 4.3, where an answer will be provided to the fourth research question.

4.1 Forecast Demand at Ecorus

In this section, the subquestion "What is the best applicable method to forecast demand at Ecorus?" will be answered. First, the SKUs are further analyzed in Section 4.1.1, and next, in Section 4.1.2, the methodology of generating the different forecasts is discussed. Furthermore, in Section 4.1.3, finding the most accurate forecasting method is explained.

4.1.1. Further Analysis of the SKUs

As mentioned in Chapter 2, there is historical data available from January 2021 until August 2023. 80 percent of this data will be used to initialize the different forecasting methods, if required. This results in 114 weeks for initializing the models. The forecasts for the SKUs will be created over the review period and lead time. Therefore, to determine the safety stock and reorder point, forecasts over the period consisting of the review period and lead time are required. As explained in Section 3.6, there is a direct connection between the forecast over a period and the required safety stock and corresponding reorder point. The lead times of the different SKUs depend on the group if the item, the corresponding lead times and the product groups of each of the SKUs are shown in Table 5.

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
Group	1	1	2	2	2	2	3	3	1	2	3
R	1	1	1	1	1	1	1	1	1	1	1
L	3	3	14	14	14	14	3	3	3	14	3
$R + L$	4	4	15	15	15	15	4	4	4	15	4

Table 5: Review Period and Lead Time of the Different SKUs in Weeks

Before continuing to the initialization of the individual forecast methods, it is important to further analyze the different SKUs. The previously mentioned methods to forecast demand have their advantages and disadvantages, which are discussed more elaborately in Section 3.2.

One distinction between the forecast methods is the ability to include trend and seasonality. To determine if the SKUs should be forecasted using trend and seasonality, the time-series can be decomposed and consequently the strength of the trend and the strength of seasonality can be determined. When decomposing the time-series, different elements of the time-series can be distinguished, such as trend and seasonality (Hyndman & Athanasopoulos, 2018). Figure 9 shows two examples of the decomposed time-series of SKUs 3 and 7. As can be seen in Figure 9, the strength of the trend and the fluctuations in the seasonality can differ significantly between different SKUs. Looking at SKU 1, a meter, the trend is quite stable after a short initial period. When looking at the trend of SKU 7, a panel, initially a strong positive trend is present. However, in the last periods there has been a strong negative trend. The change of the direction of the trend can be explained by the short SKU lifetime and the replicability of SKUs with more recent developed types, as explained in Section 2.2. As this replicability of SKUs is in particular seen in solar panels and less in invertors and meters, the trends as shown in Figure 9 are as expected.

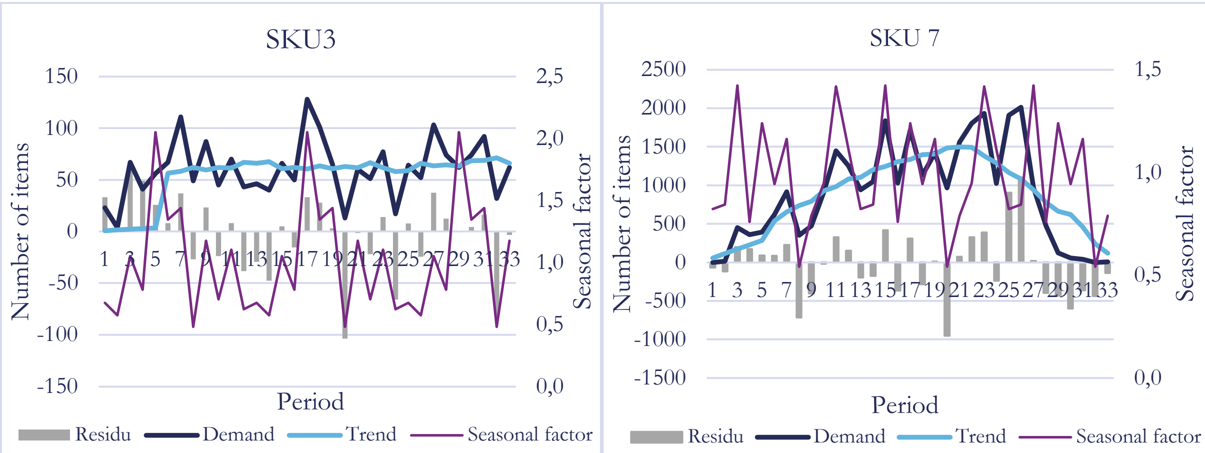

Figure 9: Decomposed time-series of SKU 3 and SKU 6

Once these time-series are decomposed, the strength of the trend and seasonality can be determined, as described in Section 3.1. The resulting strength of the trend and the seasonality for each of the SKUs as well as for the SKU groups can be found in Table 6. It can be concluded from this analysis that including seasonality could be beneficial regarding the performance of the forecasts for all SKUs as well as for the SKU groups. For SKUs 3 and 5 adding seasonality might not lead to improvement, due to the relatively lower strength of seasonality. Including the trend could be beneficial for all SKUs, however, implementing a trend for SKUs 1, 4 and 6 might not result in a significant improvement.

	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
F_T	0.699	0.841	0.899	0.729	0.822	0.724	1.000	0.901	0.814	0.920	0.844
F_S	0.977	0.999	0.882	0.996	0.714	0.998	1.000	1.000	0.998	0.998	1.000

Table 6: Strength of Trend and Strength of Seasonality per SKU and per SKU group

4.1.2. Generating Forecasts

To generate the forecasts, several forecast techniques will be considered, as described in Section 3.2. Both the Naïve and seasonal Naïve forecasting methods will be used as a benchmark for evaluating other forecasting models. Since there are SKUs that do not have a very high strength of trend, Moving Average (MA) models will also be included. Regarding Moving Averages methods, exclusively the Simple Moving Average (SMA) will be used, and various time windows will be examined. As mentioned in Section 3.2.1.3, there are different variations of exponential smoothing, which are Simple Exponential Smoothing (SES), Holt, and Holt-Winters (HW). These three types of exponential smoothing will all be included since the SKUs all have a different combination between the strength of the trend and the strength of seasonality. Machine Learning methods will not be further examined in this research, since these models are computationally intensive making it infeasible to execute in Excel, which is one of the requirements of Ecorus. Furthermore, ARIMA models will be used, since they are able to capture patterns in the data that might not be captured with only seasonality. For the ARIMA model, several combinations of the MA and AR orders will be used. To summarize, there are 7 methods that will be considered to generate forecasts for each of the time-series, which are:

1. Naïve
2. Seasonal Naïve (sNaïve)
3. Simple moving average (MA)
4. Simple exponential smoothing (SES)
5. Holt's linear trend method (Holt)
6. Holt-Winters method with multiplicative trend and multiplicative seasonality (HW)
7. Autoregressive Integrated Moving Average (ARIMA)

For each of these forecast methods, the corresponding models can be found in Appendix C. In Section 4.1.2.1 to Section 4.1.2.7, the approach for each of the forecasting methods specifically is discussed.

4.1.2.1 Naïve

Regarding the Naïve forecasting method, there are no initializations required since the forecast simply states the last observation. With this forecast method, there are no parameters required. Using the Naïve forecasting method, the forecast is the same as the last observation, in this case the observation of the previous week.

4.1.2.2 Seasonal Naïve

Similar to the Naïve forecasting method, the seasonal Naïve forecasting method does not need any initialization of parameters. When using the Seasonal Naïve method when forecasting on a weekly basis, the forecasts are the last observation of that same week, which is the observation of 1 year ago.

4.1.2.3 Simple Moving Average (MA)

Regarding the simple moving average, the initialization only requires taking the average of the last n number of periods. Once a forecast is made for the test period, these moving averages are updated before generating the next forecast. To find the best value for the time-window, a minimum of 2 weeks is tested until 13 weeks, and all values in between for each SKU and the SKU groups. This results in a total of 12 experiments for the Moving Average.

4.1.2.4 Simple Exponential Smoothing (SES)

For initializing the Simple Exponential Smoothing, finding the value of alpha is required. This is done by setting the first observation as the level at $t = 0$, and updating this value with a fraction α by the next observation, until the end of the observations is reached, which is at $t = 114$. The level at $t = 114$ is equal to the first one period ahead forecast. After each forecast is generated, the level is updated again by using the level of the latest observation with a fraction α before generating the next forecast. For the SES method, values for α between 0.05 and 0.35 with steps of 0.05 are tested.

4.1.2.5 Holt's Linear Trend Method (Holt)

For the initialization of Holt's linear trend model, the level and the trend need to be determined. This is done by decomposing the timeseries, as previous described in Section 4.1.1. After each forecast for the test period is made, the estimate for the level and trend are updated with the latest observation before a new forecast is generated. For Holt's method, combinations for α and β between 0.05 and 0.35 with steps of 0.05 are tested. This results in a total of 49 experiments for Holt's method.

4.1.2.6 Holt-Winters method with multiplicative trend and multiplicative seasonality (HW)

To determine the seasonal factors of the SKUs, first the timeseries is decomposed, as described in Section 4.1.1. Next the seasonal factors can be determined and normalized, such that the average is 1. With these factors the forecasts can be generated similar to Holt's method, but with multiplying the forecasted demand by the seasonal factor. For Holt Winters method, combinations of α and β between 0.05 and 0.35 with steps of 0.05 are tested. Since the period used for testing these forecasts is 6 months, γ does not influence the performance of these forecasts, as this would only influence the forecast which uses the same seasonal factor, which is one year later. This results in a total of 49 experiments for Holt-Winters method.

4.1.2.7 Autoregressive Integrated Moving Average (ARIMA)

When forecasting using ARIMA, it is highly dependent on the AR order, MA order and the differences for the initialization of the parameters. Once the coefficients are determined, the forecasts can be generated. After each forecast, the coefficients are redetermined prior to generating the forecast for the next period. For the determination of the best parameters, combinations for the AR and MA order are tested from 0 to 3, which results in a total of 16 experiments.

4.1.3. Finding the Most Accurate Forecasts

When generating the forecasts as explained in Section 4.1.2, a total of 135 forecasts will be generated. After the forecasts are created, the performance of each of the forecasts will be determined. This will be done by calculating the RMSE, the bias and MAE% for each of these forecasts, as described in Section 3.3. Once these KPIs are calculated, the most accurate forecast method can be determined.

4.2 Inventory Management Policy

In this section, the subquestions "Which inventory control systems can be used by Ecorus?" and "What are the corresponding parameters of Ecorus for these inventory control systems?" will be answered.

4.2.1. Selection of Inventory Policy

As described in Section 3.4, the two most suitable inventory polices are the (s, Q) system and the (R, s, Q) system. Since the time window used to forecast is one week, the smallest possible time window used in the inventory policy is also one week. As described in Section 1.2.1, the items that are required at Ecorus their instalment partners are sent out every week. Therefore, the stock levels should also be analyzed every week. This results in a review period of $\mathrm{R} = 1$. Since an (R, s, Q) system with $R = 1$ can also be seen as an (s, Q) policy, the model that will be used for Ecorus is the (s, Q) policy.

4.2.2. Determination of Parameters

To create an (s, Q) policy, several parameters are required, as discussed in Section 3.6.1. These parameters are demand per year, unit variable costs, inventory carrying charge, fixed order costs, costs per unit short, order quantity, safety factor, standard deviation of forecasts over lead time demand, safety stock, and reorder point. The determination of each of these parameters is done in the remainder of this subsection. An overview of all parameters discussed in this section and their values are shown in at the end of this sub section in Table 13.

Demand per Year (D)
To determine the demand per year, the demand data over the test period is used. The total demand over this period is then divided by the number of periods, which is 114 weeks, and multiplied by the number of periods in a year, which is 52 weeks. The resulting values for the demand per year can be found in Table 13.

Unit Variable Costs (v)
The unit variable costs of the SKUs consist of the unit purchase price, indicated by p, and any additional costs incurred to use the SKUs. As described in Section 2.3, the warehouse is outsourced by Ecorus, making most costs variable to the number of parts. To determine the unit variable costs for Ecorus, the transportation costs to the warehouse and the entry and exit costs incurred by storing the pallets in the warehouse are added as well. An overview of the number of items per pallet, the costs per pallet, and the corresponding costs per unit are shown in Table 7.

	Meters	Invertors	Panels
Pallet size (units)	20	24	31
Transport costs (€/pallet)	73.44	18.36	26
Transport costs (€/unit)	3.67	0.77	0.84
Entry costs (€/pallet)	2.62	2.62	2.62
Exit costs (€/pallet)	3.00	2.25	2.25
Entry and exit costs (€/unit)	0.15	0.09	0.07
Total unit variable costs (€/unit)	p + 0.49	p + 0.97	p + 1.00

Table 7: Unit Variable Costs per Product Group

Inventory Carrying Charge (r)

To determine the inventory carrying charge, several costs aspects need to be taken into consideration, which are the WACC, the insurance costs per unit, and the warehouse costs per unit, as described in Section 3.5. For the WACC, a percentage of 3.3% of the unit purchase price is taken into consideration, which is used within the different departments of Ecorus for the WACC. The charge for keeping a pallet in stock differs per product group. For meters and invertors € 1.19 is charged by the warehouse per pallet per week, for solar panels €2.97 is charged per pallet per week. An overview of these costs is given in Table 8. Usually, insurance costs are also included in the inventory carrying charge, however, the warehouse does not provide this. Therefore, additional insurance is taken out from a third party. For this insurance, a fixed amount per year is charged, which is independent of the number of items stored. Since these insurance costs are not variable but fixed, the insurance costs are not included in the inventory carrying charge.

	Meters	Invertors	Panels
Pallet size (units)	160	24	31
Inventory costs (€/pallet/week)	1.19	1.19	2.97
WACC (€/unit/week)	0.033 p	0.033 p	0.033 p
Inventory carrying charge(€/unit/week)	0.033p + 0.01	0.033p + 0.05	0.033p + 0.10

Table 8: Inventory Carrying Charge per Product Group

Fixed Order Costs (A)

The fixed order costs are, as described in Section 3.5, the costs that are made by placing an order, regardless of the size of this order. There are no fixed order costs at the supplier, however, there are fixed costs per order at the warehouse. Regardless of the size of the order, €5 is charged for entry of the order. Furthermore, Ecorus uses a fixed order costs to express the labor required by placing an order, which is set to €120 per order. This amount is set to cover the time required for ordering the items and handling the paperwork of it. An overview of the fixed order costs is shown in Table 9.

	Meters	Invertors	Panels
Entry costs (€/order)	5	5	5
Labor costs (€/order)	120	120	120
Fixed order costs (€/order)	125	125	125

Table 9: Fixed Order Costs per Product Group

Costs per Unit Short (B_2)

One distinguishment between the two methods to determine k as described in Section 3.6.1, is the use of the shortage costs, B_1 or the costs per unit short, B_2. As described in Section 2.1, when stock-outs occur, SKUs can be bought at a back-up supplier in the Netherlands instead of overseas. These price increases and thus additional costs can be seen as the costs per unit short. The price increase as percentage of the unit price can be found in Table 2, the resulting values for B_2 can be found in Table 10.

	Meters	Invertors	Panels
Price increase (%/unit)	30%	30%	25%
Costs per unit short (€/unit)	0.30p	0.30p	0.25p

Table 10: Unit Variable Costs per Product Group

Order Quantity (Q)

The order quantity can be determined by using the EOQ, as shown in Equation (14). The required input parameters are the fixed order costs, the yearly demand, the unit variable price, and the inventory carrying charge, which are determined as described above. When implementing these values in Equation (14), the resulting EOQ as shown in Table 11 can be obtained. However, these quantities cannot directly be used as the order quantity, as the SKUs are only procured per pallet. Therefore, the EOQ is rounded to the nearest multiple of the pallet size. The resulting values for the order quantity are shown in Table 11.

	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
Pallet size (units)	160	160	24	24	24	24	31	31	160	24	31
EOQ	16	46	11	11	55	24	86	161	45	36	252
Q	160	160	24	24	48	24	72	155	160	24	248

Table 11: Determination of Q by using EOQ

Safety Factor (k)

To determine the safety factor, Equation (18) can be used. The required input for using this equation are the order quantities, the holding costs, the yearly demand, and the shortage costs per unit short, which are determined as described above. The results of the left-hand side of this equation are shown in Table 12. As can be seen in this table, for none of the SKUs or SKU groups Equation (18) is satisfied. Therefore, Equation (19) should be used to determine k. Since the probability that a unit normal variable takes on a value of k or larger, denoted by $p_{u\geq(k)}$, is equal to the left-hand side of Equation (18), the final step is to determine at which value of k Equation (19) is true. This can be done by inserting $p_{u\geq(k)}$ in Equation (21). The resulting values of k can be found in Table 12.

	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
$\frac{Qr}{DB_2}$	0.12	0.02	0.01	0.01	0.04	0.01	0.01	0.00	0.01	0.00	0.00
$p_{u\geq(k)}$	0.12	0.02	0.01	0.01	0.04	0.01	0.01	0.00	0.01	0.00	0.00
k	2.19	3.13	3.28	3.47	2.73	3.45	3.32	3.65	3.21	4.01	3.73

Table 12: Determination of k by implementing Equation (18)

Standard Deviation of Forecasts Over Lead Time (σ_L)

As discussed in Section 3.6.1, a Normal distribution can be assumed if the demand during lead time is over 10 items (Silver et al., 2016), which is the case for all SKUs and the product groups. To use the Normal distribution, two parameters are required, which are the average μ, and the standard deviation σ. Once the forecasts are generated and thus the forecast errors are known, the average and standard deviation over these errors can be calculated.

Safety Stock (SS)

The safety stock can be determined by using Equation (16). Once the values of σ_L and k are known, the safety stock can be determined by multiplying these two values with each other.

Reorder Point (s)

The reorder point can be determined by using Equation (15). The reorder point is equal to the safety stock and the forecasted demand during the lead time. The safety stock can be determined as described above, the forecasted demand during the lead time will be the result of the forecasts, as will be discussed in Chapter 5.

Total Costs of Inventory Policy (TC)

The total costs of an inventory policy can be determined by using Equation (17). All parameters used in this equation can be determined by using the methods described in this subsection. The standard deviation of forecasts over lead time, the safety stock, and the reorder point should be determined after the forecasts are generated, this will be done in Chapter 5.

Overview of Required Parameters

Table 13 provides an overview of all parameters required for implementing the (s, Q) that are known before the forecasts can be generated.

	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Inverto rs	Panels
D	447	3167	705	1169	387	1101	2448	13247	3853	6058	26416
v	p+0.49	p+0.49	p+0.97	p+0.97	p+0.97	p+0.97	p+1.00	p+1.00	p+0.49	p+0.97	p+1.00
r	0.033p +0.01	0.033p +0.01	0.033p +0.05	0.033p +0.05	0.033p +0.05	0.033p +0.05	0.033p +0.10	0.033p +0.10	0.033p +0.01	0.033p +0.05	0.033p +0.10
A	125	125	125	125	125	125	125	125	125	125	125
B_2	0.30p	0.30p	0.30p	0.30p	0.30p	0.30p	0.25p	0.25p	0.30p	0.30p	0.25p
Q	160	160	24	24	48	24	72	155	160	24	248
k	2.19	3.13	3.28	3.47	2.73	3.45	3.32	3.65	3.21	4.01	3.73

Table 13: Parameters Required for the (s, Q) Inventory Policy

4.2.3. Evaluation of Inventory Policy

To evaluate the different inventory policies, the total costs can be compared. Furthermore, the fillrate can be determined. When adjusting the values of k and Q both the total costs and the fillrate might change. Therefore, several adjustments of k and Q will be examined. For of k an increase of +0.5, +1.0, -0.5, and -1.0 will be examined. Q will be increased and decreased with 1 and two pallet sizes, if possible. For SKUs that already have Q at the same quantity as a pallet, no decrease of Q will be analyzed. The resulting values of the total costs and the fillrate can then plotted in a scatterplot, which shows the relationship between the fillrate and the total costs. From this analysis it can be determined what the increase in costs would be if it is desired to increase the fillrate, or on the contrary, what the financial gains would be if it would be accepted to have a lower fillrate.

4.3 Conclusion

After further analyzing the SKUs, different characteristics for each of the time-series are found. In total 7 different forecast methods will be used to find the most accurate forecasting method. Naïve and sNaïve will be used as a benchmark, the other forecasting methods will be included as well, since they can capture trend, seasonality, or other patterns in the data. In total 135 forecasts will be generated, and for each of them the RMSE, bias, and MAE% will be determined. Based on these KPIs, the most accurate method to forecast demand for each SKU will be determined.

The best applicable inventory policy for Ecorus is the (s, Q) inventory policy. To implement this policy, several parameters can be determined prior to generating the forecasts. These parameters are demand per year, unit variable costs, inventory carrying charge, fixed order costs, costs per unit short, order quantity, and safety factor. Once the forecasts are generated, the standard deviation of forecasts over lead time demand, safety stock, and reorder point can be determined, which will be done in Chapter 5. Once these parameters are determined, the total costs of the inventory policy can be determined, as well as the corresponding fillrate. When performing a sensitivity analysis of the total costs compared to the fillrate, the costs of increasing the fillrate, or the financial gains when decreasing the fillrate can be determined.

Chapter 5 Analysis of Results

In this chapter, the results from the methods described in Chapter 4 will be discussed. In Section 5.1, the results of the forecasts will be discussed, and in Section 5.2 insight is these results are discussed. Next, in Section 5.3 the analysis of the results of the inventory policy will be discussed. Finally, this chapter will be concluded in Section 5.4.

5.1 Analysis of Forecasting Results

In this section, the subquestion "What is the performance of the forecasts?" will be answered. As described in Chapter 4, seven different methods to forecast demand at Ecorus are analyzed. The results of each of these methods are described in Section 5.1.1. to Section 5.1.7.

5.1.1. Naïve and sNaïve

When using the Naïve and the sNaïve method, the RMSE values as shown in Figure 10 can be obtained. Using the Naïve method, the forecasting results as shown in Table 14 are obtained. Using the sNaïve method, the results shown in Table 15 are obtained. As can be seen in these tables, for all SKUs except SKUs 3, 4 and 5, the Naïve method outperforms the sNaïve method. For SKU 4 this can be explained by the relatively low strength of trend and the high strength of seasonality. For all other SKUs it is as expected that Naïve performs better than sNaïve, due to the high strength of trend. With Naïve forecasting the observation of 1 week previously is used, with sNaïve the observation of 1 year ago is used. With a strong trend it can be expected that the observation for last year is significantly different from the current value.

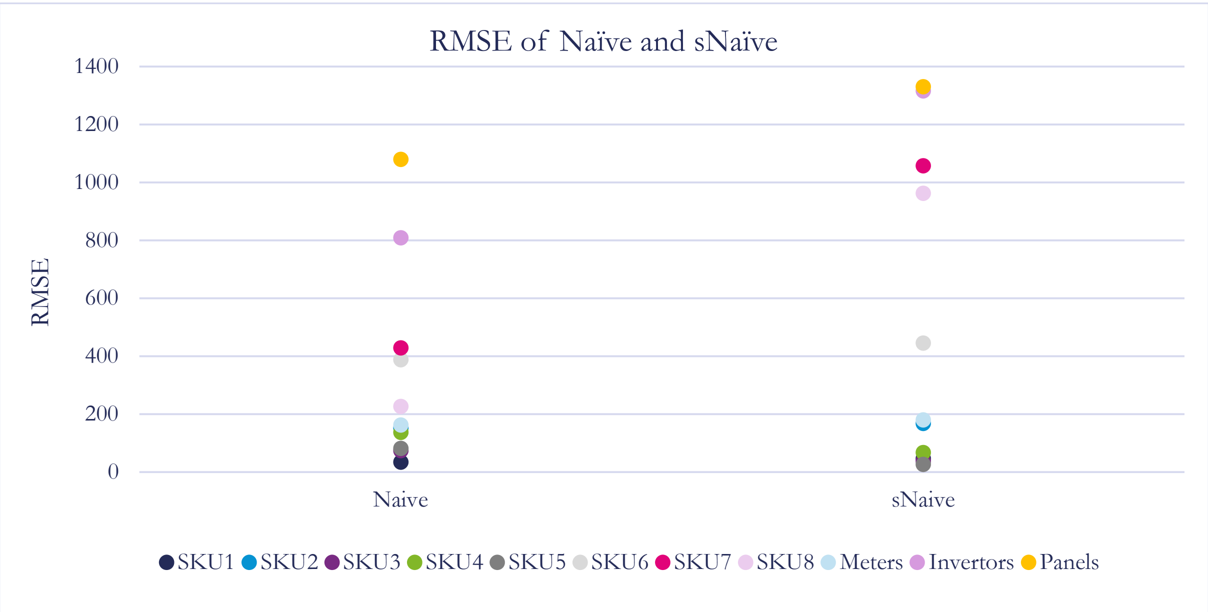

Figure 10: RMSE of the Forecasts Using Naïve and sNaïve Forecasting Method

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
RMSE	34.8	153.4	74.3	137.3	82.2	389.0	429.1	227.7	162.8	810.1	1080.4
MAE%	181	103	240	246	434	540	141	248	87	195	84
bias	-2.6	-0.3	11.4	-10.4	8.9	177.4	162.0	97.2	-1.5	-288.9	59.4

Table 14: Forecasting Results Naïve Method

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
RMSE	46.8	168.5	38.0	67.6	27.5	445.5	1058.5	963.3	180.4	1316.9	1331.9
MAE%	249	122	142	100	200	680	335	1646	115	381	122
bias	-16.8	-105.2	21.8	-32.0	-16.3	-267.7	-722.5	841.1	-142.5	-861.9	-1028.7

Table 15: Forecasting Results sNaïve Method

5.1.2. Moving Average (MA)

When using the Moving Average method, the RMSE values as shown in Figure 11 can be obtained. As can be seen in this figure, for the product groups and SKUs 4 and 5, the RMSE values decrease when the MA window is increased. For SKU 4, SKU 5, the meters and the panels this can be explained by the relatively low strength of trend. With a higher moving average window less, trend is captured in the forecasts. On the other hand, SKUs 7 and 8 have a strength of trend of 1.000. It can therefore be expected that the moving average window with the lowest value performs the best, as it is most reactive to the most recent observations. An overview of the best settings and the KPIs for all items is shown in Table 16.

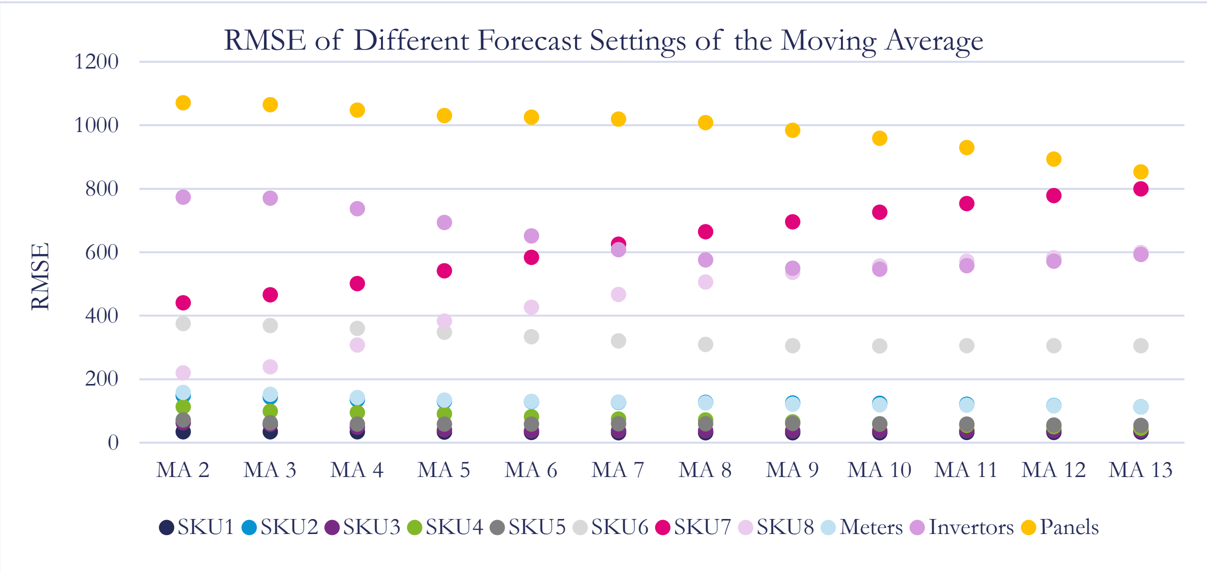

Figure 11: RMSE of the Forecasts Using Different Configurations for the Moving Average Forecasting Method

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
Window	8	13	8	13	13	10	2	2	13	10	13
RMSE	32.0	113.7	38.1	46.1	55.8	305.1	441.3	220.5	114.1	547.6	854.2
MAE%	179.7	82.0	122.6	76.1	402.2	433.3	142.1	254.1	65.7	154.8	77.8
bias	-5.0	-10.5	15.6	17.0	32.7	95.3	173.6	112.3	-19.2	-350.0	17.4

Table 16: Forecasting Results MA Method

5.1.3. Simple Exponential Smoothing (SES)

Using Simple Exponential Smoothing, the RMSE values as shown in Figure 11 are obtained. Similar to the Moving Average method, SKU7, SKU 8, and the invertors perform better once α is increased. With a higher α, the forecasts become more reactive to the most recent observations. Since these SKUs have the highest possible strength of trend, it is expected that more reactive models perform better. For all other SKUs the strength of trend is significantly lower, which explains that a less reactive model with a lower α performs better. The best settings for each of the SKUs and the SKU groups and the corresponding KPIs can be found in Table 17.

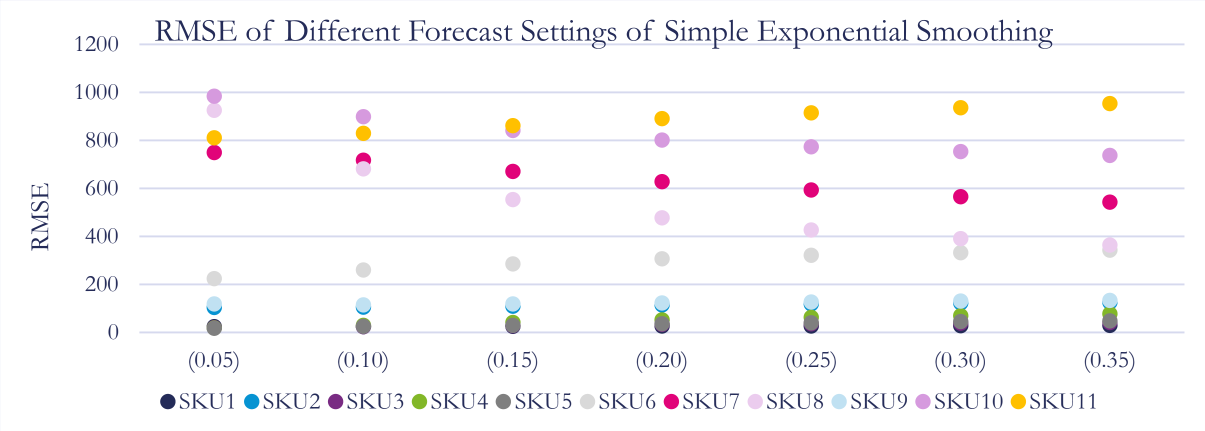

Figure 12: RMSE of the Forecasts Using Different Configurations for the Simple Exponential Smoothing Forecasting Method

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
Window	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.35)	(0.35)	(0.10)	(0.35)	(0.05)
RMSE	24.2	104.2	21.0	19.2	18.2	224.6	543.8	365.2	115.9	738.7	811.4
MAE%	134.4	72.4	70.7	32.1	114.7	297.6	178.2	428.2	69.2	212.1	69.7
bias	-4.8	-35.8	-3.6	5.1	8.9	71.4	159.0	218.8	-35.9	-479.4	-265.0

Table 17: Forecasting Results Simple Exponential Smoothing Method

5.1.4. Holt's Linear Trend Method (Holt)

When using Holt's linear trend model, the RMSE values for each combination of α and β can be found in Figure 13. SKU 8 is the only SKU that has a better performance if α and β are increased to the maximum value of 0.35. For SKU 7 this pattern is also expected since both SKUs had a similar performance and settings in the previous discussed methods, however, SKU 7 performs best when α and β are decreased to the minimum value of 0.05. Even though all SKUs have a significant trend as shown in Table 6, most SKUs have a better performance when SES is used compared to Holt's method. The best performing values of α and β and the resulting KPIs are shown in Table 18. When comparing these results with the results of SES, as shown in Table 17, it can be concluded that adding a trend does not improve the accuracy of the forecasts for any of the SKUs of the groups.

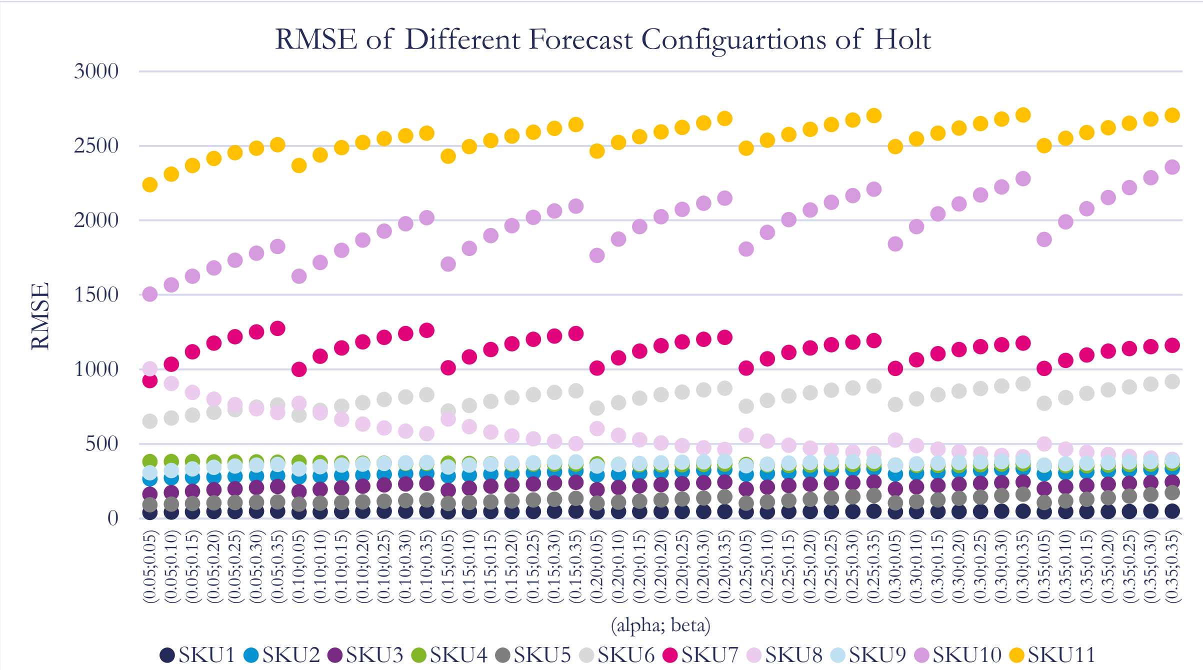

Figure 13: RMSE of the Forecasts Using Different Configurations for the Holt Forecasting Method

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
Alpha	0.05	0.05	0.05	0.35	0.05	0.05	0.05	0.35	0.05	0.05	0.05
beta	0.05	0.05	0.05	0.10	0.05	0.05	0.05	0.35	0.05	0.05	0.05
RMSE	39.6	266.1	161.9	353.3	91.3	650.9	923.6	397.1	306.3	1503.7	2240.2
MAE%	248.9	231.6	639.1	632.9	758.6	1121.5	359.5	304.3	222.3	453.1	237.4
bias	34.8	239.4	110.0	219.6	61.7	441.3	778.3	128.2	274.5	1024.1	2015.6

Table 18: Forecasting Results Holt Method

5.1.5. Holt-Winters Method with Multiplicative Trend (HW)

When using Holt-Winters to forecast, the RMSE values as shown in Figure 14 are obtained. As with previous methods, SKU 8 performs best when α and β are increased to the maximum value of 0.35, which is expected since the high strength of trend and seasonality. For SKU 1 and the meters there is no significant difference between the used values of α and β. The invertors and panels show a similar pattern as seen with Holt's method, since the forecasts perform better once α and β are decreased to the minimum value of 0.05. Given the high strength of trend for most SKUs, it is unexpected to see that for most SKUs the forecasts are most accurate if β is decreased to the minimum of 0.05. An overview of the best values for α and β and the resulting KPIs are shown in Table 19. When comparing these KPIs to Table 18, it can be concluded that adding seasonality is not improving the accuracy for any of the SKUs of groups.

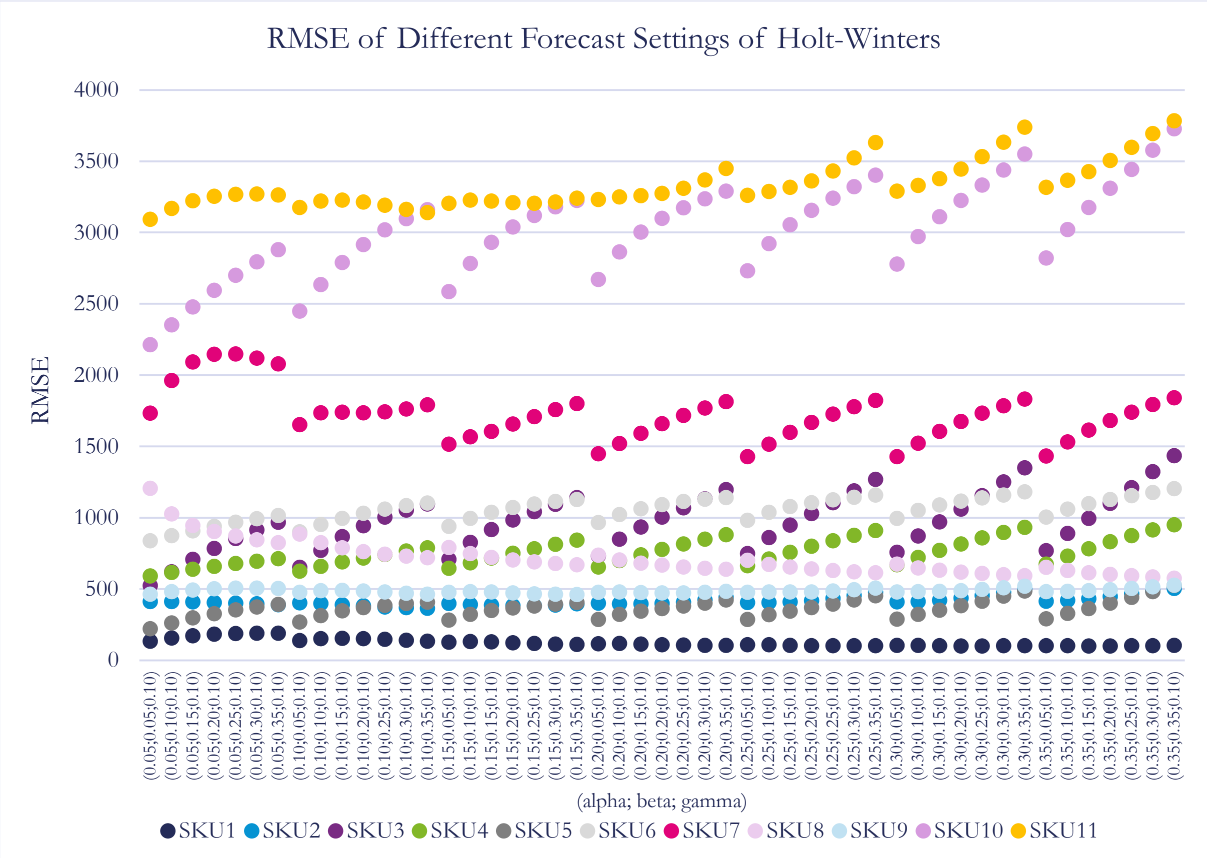

Figure 14: RMSE of the Forecasts Using Different Configurations for the Holt-Winters Forecasting Method

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
Alpha	0.35	0.10	0.05	0.05	0.05	0.05	0.30	0.35	0.15	0.05	0.05
Beta	0.15	0.35	0.05	0.05	0.05	0.05	0.05	0.35	0.35	0.05	0.05
Gamma	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10
RMSE	101.0	365.9	524.9	590.7	222.3	838.2	1427.4	577.0	462.3	2214.2	3093.3
MAE%	491.4	294.6	1991.8	1147.6	1770.7	1388.4	476.7	400.1	307.2	662.3	311.5
bias	66.3	304.5	342.7	398.1	144.1	546.3	1023.8	173.9	379.3	1497.0	2645.4

Table 19: Forecasting Results Holt-Winters Method

5.1.6. Auto Regressive Integrated Moving Average (ARIMA)

When implementing ARIMA, the RMSE values as shown in Figure 15 are obtained. For SKUs 1, 3, 4, and 5 low RMSE values are achieved. For all SKUs except SKU 8, a more accurate forecast is achieved when implementing ARIMA compared to Holt-Winters and Holt. The best settings for each of the SKUs and the SKU groups is shown in Table 20.

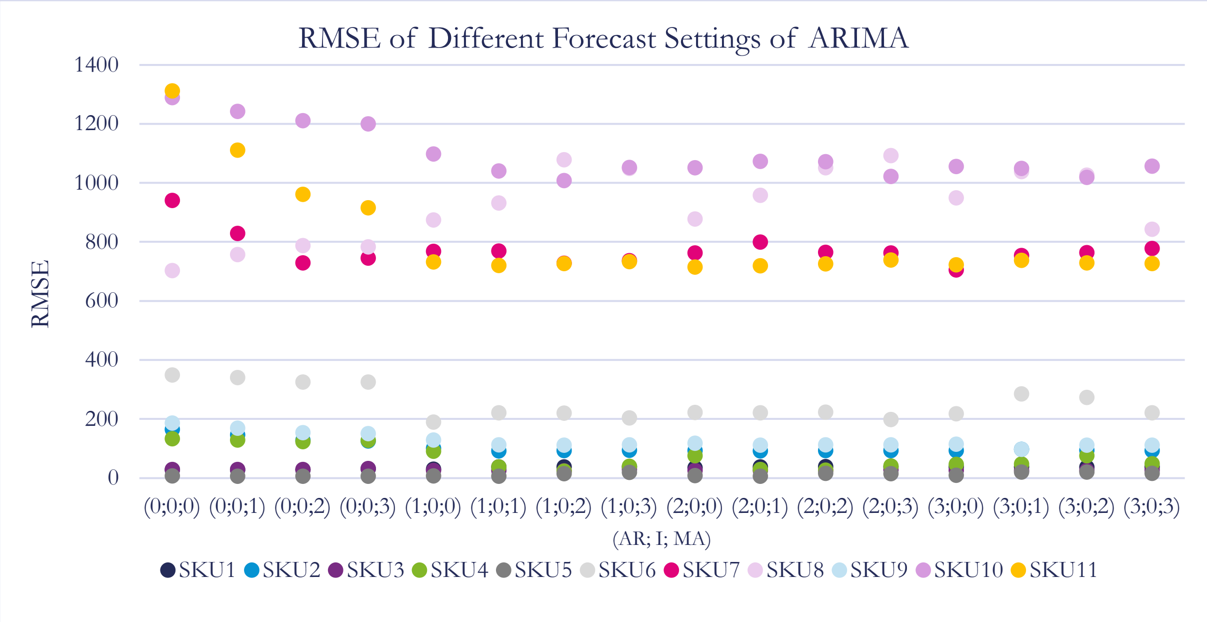

Figure 15: RMSE of the Forecasts Using Different Configurations for the ARIMA Forecasting Method

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
(AR;I:MA)	(0;0;1)	(2;0;1)	(1;0;2)	(1;0;2)	(0;0;1)	(1;0;0)	(3;0;0)	(0;0;0)	(3;0;1)	(1;0;2)	(2;0;0)
RMSE	28.7	92.7	25.9	23.9	6.5	190.3	706.3	704.2	97.9	1008.2	715.9
MAE	161.6	66.6	88.9	41.3	47.5	305.0	258.3	1216.9	58.3	300.8	62.0
bias	-17.7	0.2	-15.3	-12.2	-2.2	-31.8	430.3	621.9	-19.4	-679.9	-21.0

Table 20: Forecasting Results ARIMA Method

5.2 Insights in Forecast Performance

In this section, the subquestion "Which insights can be gained regarding the forecasts?" will be answered. As previously described in Section 4.1, forecasts are generated using 7 different approaches, which are Naïve, sNaïve, MA, SES, Holt, Holt Winters and ARIMA. For the methods MA, SES, Holt, Holt Winters and ARIMA numerous forecasts have been generated to find the settings which lead to the best performing forecast. Figure 16 gives an overview of these best performances for each of the methods.

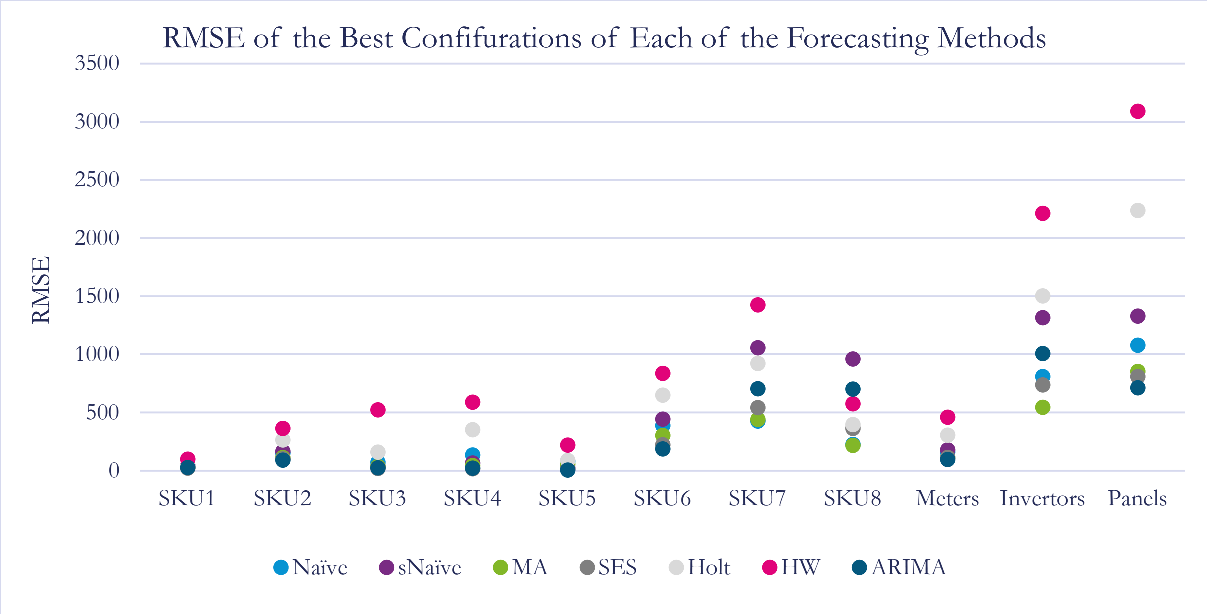

Figure 16: MAE% of the Best Performing Forecast for Each Forecasting Method

As can be seen in Figure 16, the best performing methods are SES and ARIMA and MA has the second-best performance. For most SKUs the difference between ARIMA and MA is small, however, for the SKU groups this difference increases. As mentioned in Section 3.2.1, the ARIMA model is a combination of the AR and MA models, which explains the similar performance of MA and ARIMA. However, when forecasting the panels as a SKU group, good forecasting results can be obtained. Table 15 summarizes the best RMSE observed for each of these methods, the overall best method is illustrated by green. The settings of these best performing method per SKU and SKU group are shown in Table 22.

	Naïve	sNaïve	MA	SES	Holt	HW	ARIMA
SKU1	34.8	46.8	32.0	24.2	39.6	101.0	28.7
SKU2	153.4	168.5	113.7	104.2	266.1	365.9	92.7
SKU3	74.3	38.0	38.1	21.0	161.9	524.9	25.9
SKU4	137.3	67.6	46.1	19.2	353.3	590.7	23.9
SKU5	82.2	27.5	55.8	18.2	91.3	222.3	6.5
SKU6	389.0	445.5	305.1	201.6	650.9	838.2	190.3
SKU7	429.1	1058.5	441.3	543.8	923.6	1427.4	706.3
SKU8	227.7	963.3	220.5	365.2	397.1	577.0	704.2
Meters	162.8	180.4	114.1	115.9	306.3	462.3	97.9
Invertors	810.1	1316.9	547.6	738.7	1503.7	2214.2	1008.2
Panels	1080.4	1331.9	854.2	811.4	2240.2	3093.3	715.9

Table 21: RMSE of the Best Performing Forecast for Each Forecasting method

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
Method	SES	ARIMA	SES	SES	ARIMA	ARIMA	Naïve	MA	ARIMA	MA	ARIMA
Parameters	(0.05)	(2; 0; 1)	(0.05)	(0.05)	(0; 0; 1)	(1; 0; 0)	N/A	2	(3; 0; 1)	10	(2; 0; 0)

Table 22: Methods and Parameters of the Best Performing Forecasts

Using these methods, the forecasts for the SKUs as shown in Figure 17 and the forecasts for the SKU groups as shown in Figure 18 can be obtained.

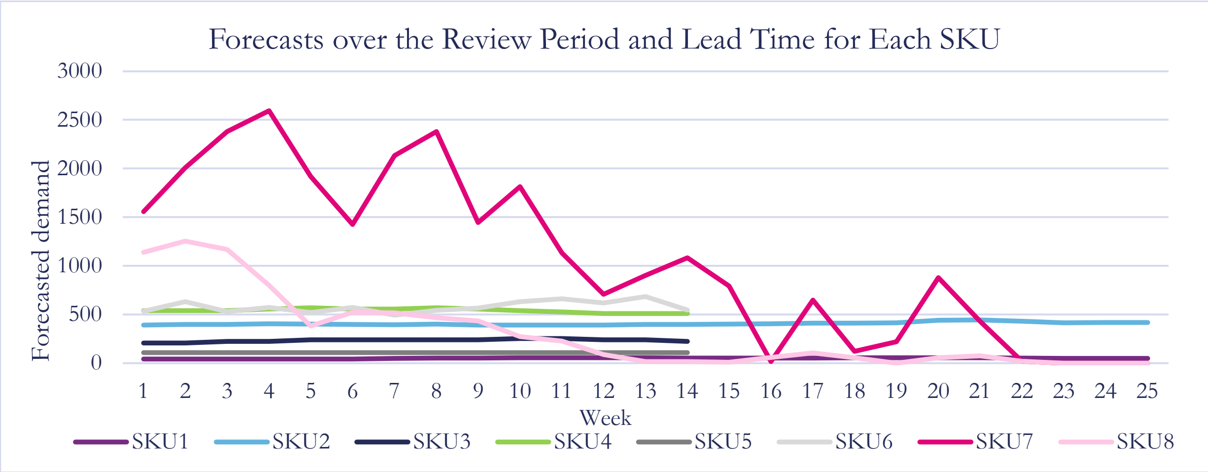

Figure 17: Most Accurate Forecasts for the SKUs

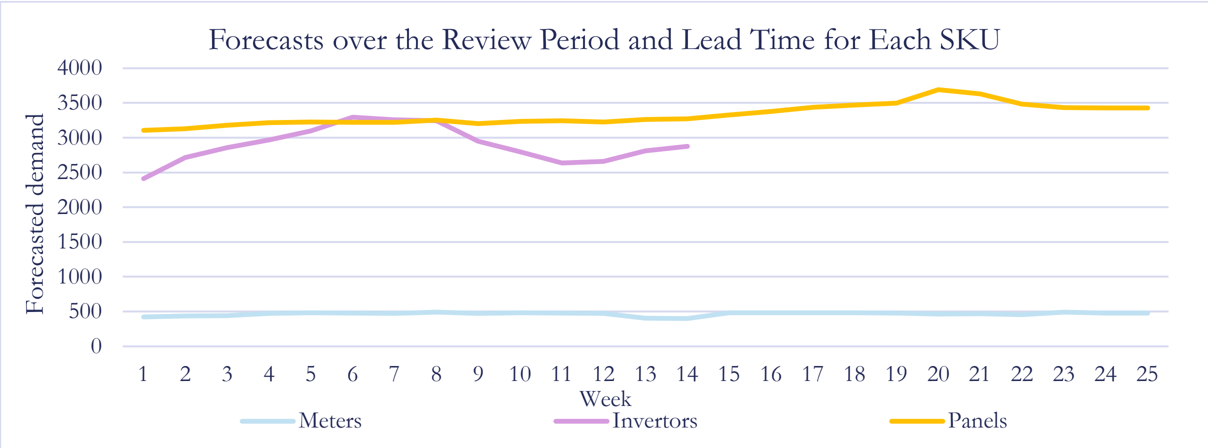

Figure 18: Most Accurate Forecasts for the SKU groups

5.3 Analysis of Inventory Policy Results

In this section, the subquestion "What is the performance of the inventory management policy for Ecorus?" will be discussed. As described in Section 4.2.2, there are several parameters that have been determined prior to generating the forecasts, as summarized in Table 13. However, there are also some parameters that can only be determined after the forecasts are generated, which will be done in 5.3.1. Furthermore, in Section 5.3.2, a complete inventory policy is created for one of the SKUs. Next, in Section 5.3.3, a sensitivity analysis is performed which can be used to compare the total costs to the fillrates.

5.3.1. Further Determination of Parameters Required for the (s, Q) Policy

As discussed in Section 4.2.2, the standard deviation over lead time demand, the safety stock, and the reorder point can only be determined after the forecasts are created. Using the best performing forecasts, as described in Section 5.2, the forecast errors can be determined when comparing the forecast to the demand over the leadtime. The forecast errors over the leadtime for SKU 1 are shown in Figure 19. When determining the mean and the standard deviation of these errors for all SKUs, the results as shown in Table 23 can be obtained.

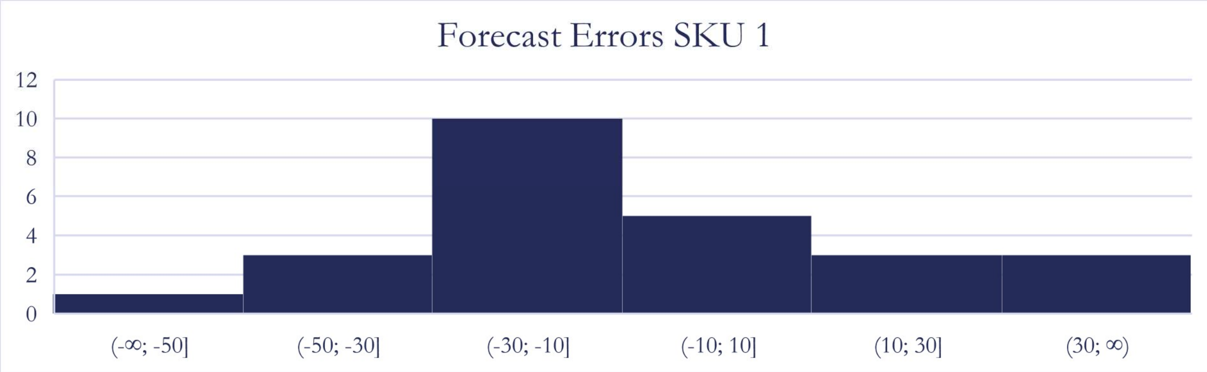

Figure 19: Histogram of Forecast Errors for SKU 1

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
μ	-5.8	0.3	-7.6	10.9	-4.7	-68.1	94.4	34.8	-23.3	-150.0	-25.2
σ_L	25.9	101.6	29.7	25.9	8.2	270.1	428	200.4	104.7	283.1	783.8

Table 23: Distribution of Forecast Errors

Once these distributions are known, the safety stock can be determined. This can be done by using Equation (16), as described in Section 5.2. The resulting values for the safety stocks are shown in Table 24.

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
k	2.19	3.13	3.28	3.47	2.73	3.45	3.32	3.65	3.21	4.01	3.73
σ_L	25.9	101.6	29.7	25.9	8.2	270.1	428	200.4	104.7	283.1	783.8
SS	7	14	6	6	3	20	41	29	14	35	53

Table 24: Safety Stock for each SKU and the SKU groups

Once the safety stock is known, the reorder point at each t can be determined by using Equation (15). An overview of the reorder points at in the different weeks can be seen in Figure 20 for the SKUs and in Figure 21 For the SKU groups.

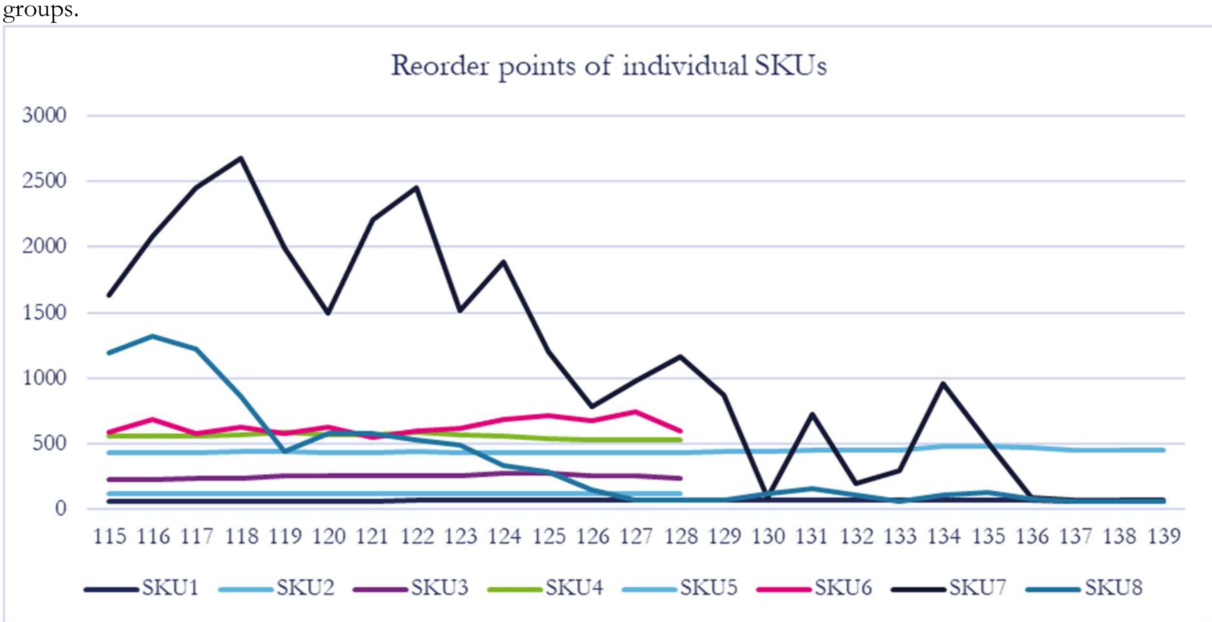

Figure 20: Reorder points of SKUs

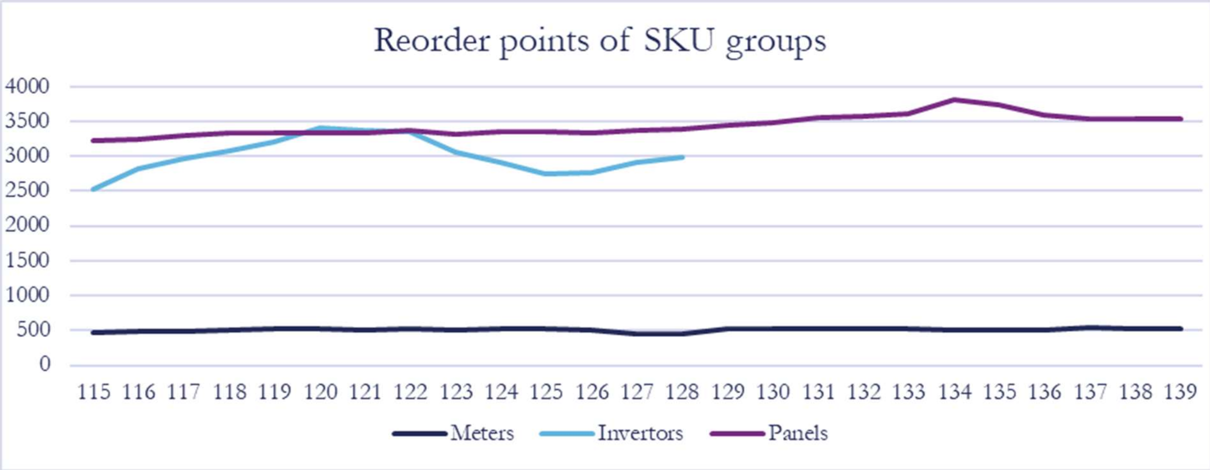

Figure 21: Reorder points of SKU Groups

Using these reorder points and Equation (17), the total expected yearly costs can be determined. An overview of these costs is shown in Table 24.

	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
TC (€)											

Table 25: Expected Total Annual Costs

5.3.2. Implementation of (s, Q) Policy

When the forecasts over the review period and lead time, as shown in Section 5.2, are combined with the parameters described in Section 4.2, the (s, Q) policy can be created for the SKUs. When using the Equations in Section 3.6, the results as shown in Table 26 can be obtained for SKU 1. A visual representation of this data is shown in Figure 22.

t	Beginning on-hand inventory	Received replenishment order	Realized demand	Demand fullfilled from stock	Demand backordered	Ending on-hand inventory	Total pipeline	Inventory position	Forecasted demand during	Reorderpoint	Replenishment order
115	64	0	9	9	0	55	0	55	44	63	120
116	55	0	13	13	0	42	120	162	44	63	0
117	42	0	14	14	0	28	120	148	44	63	0
118	28	120	18	18	0	130	0	130	44	63	0
119	130	0	13	13	0	117	0	117	44	63	0
120	117	0	20	20	0	97	0	97	44	63	0
121	97	0	25	25	0	72	0	72	48	67	0
122	72	0	19	19	0	53	0	53	52	71	120
123	53	0	29	29	0	24	120	144	52	71	0
124	24	0	26	24	2	0	120	118	56	75	0
125	0	120	6	6	0	114	0	114	56	75	0
126	114	0	12	12	0	102	0	102	56	75	0
127	102	0	5	5	0	97	0	97	56	75	0
128	97	0	9	9	0	88	0	88	52	71	0
129	88	0	20	20	0	68	0	68	52	71	120
130	68	0	12	12	0	56	120	176	52	71	0
131	56	0	18	18	0	38	120	158	52	71	0
132	38	120	26	26	0	132	0	132	56	75	0
133	132	0	12	12	0	120	0	120	56	75	0
134	120	0	6	6	0	114	0	114	56	75	0
135	114	0	0	0	0	114	0	114	56	75	0
136	114	0	0	0	0	114	0	114	52	71	0
137	114	0	0	0	0	114	0	114	48	67	0
138	114	0	16	16	0	98	0	98	48	67	0
139	98	0	19	19	0	79	0	79	48	67	0

Table 26: Inventory positions, on-hand inventory, reorder point, safety stock, and replenishment orders for SKU 1

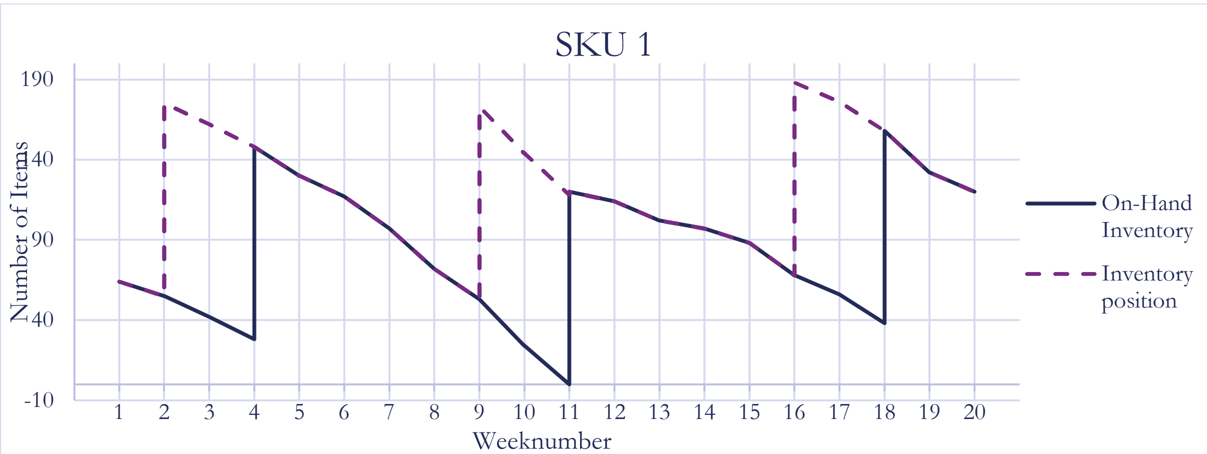

Figure 22: On-hand Inventory and Inventory Position for SKU 1

The actual costs of these policies can also be calculated, in addition to the total expected costs as shown in Table 24. Since these costs do not cover an entire year, the costs are also divided by the number of periods and multiplied by 52 to determine the expected annual costs, notated as TAC. An overview of the total costs as well as the achieved fillrate are shown in Table 27.

	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
n	25	25	15	15	15	15	25	25	25	15	25
TC (€)											
TAC (€)											
Fillrate (%)	99.0	90.0	100.0	100.0	100.0	17.0	83.0	94.0	100.0	100.0	100.0

Table 27: Total Costs and Fillrates for the (s, Q) policy

When comparing Table 24 to Table 27, there are some SKUs for which the observed costs, as in Table 27, differ significantly from the expected costs. Especially for SKU 6, the costs are much more than expected. This can also be seen in the corresponding fillrate of only 17%. Further analysis of this SKU shows that insufficient SKUs are present at the start. Figure 23 shows the on-hand inventory of SKU 6 and the inventory of position. Since this item has a lead time of 14 weeks, the orders arrive too late. Ordering more items would not change the low fillrate. Due to the low fillrate, the costs per unit short are significantly higher than expected.

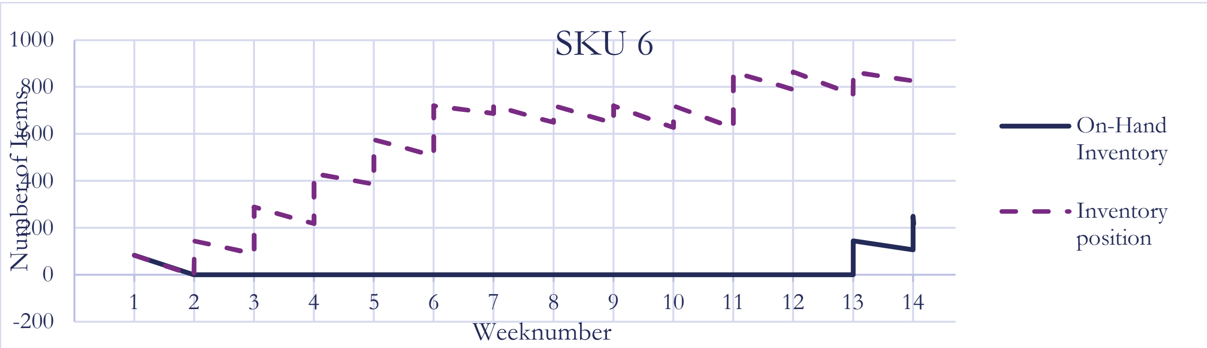

Figure 23: On-hand Inventory and Inventory Position of SKU 6

SKU 7 also has a quite low fillrate. When further analyzing this SKU, it can be seen that there are some instances at which the inventory position and the on-hand inventory drop to or below 0, as can be seen in Figure 24. For this SKU it might be beneficial to adjust k or Q.

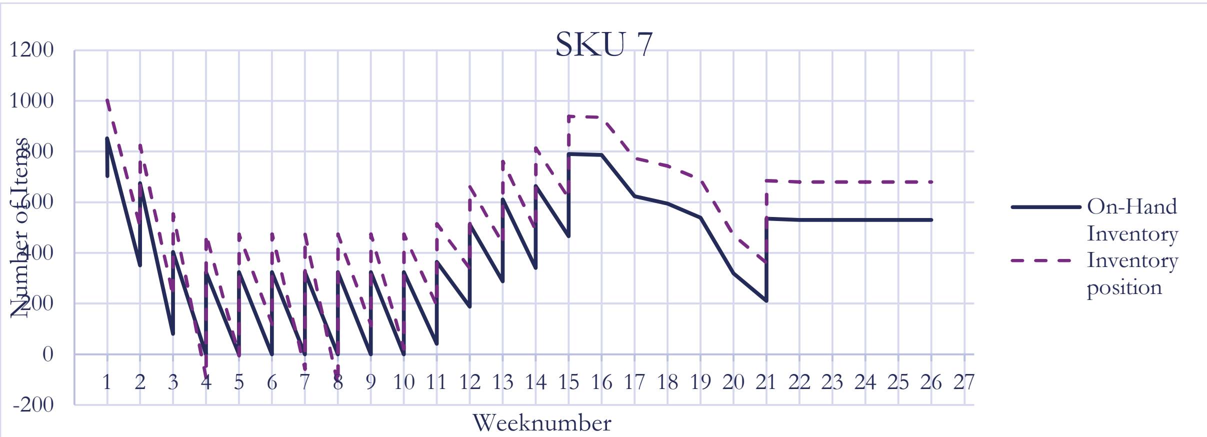

Figure 24: On-hand Inventory and Inventory Position of SKU 7

5.3.3. Sensitivity Analysis Comparing the Total Costs to Fillrate

As mentioned in Section 4.2.3, a sensitivity analysis can be performed to evaluate the influence of several parameters. Figure 25 to Figure 28 show the effects on the total costs and the fill rate if the order quantity and the safety factor are adjusted. Each datapoint has a label indicating the adjustment compared to the base policy, as described in Section 5.3.2. The definition of these labels is listed in Table 28.

Label	Adjustment
Base	Base policy as described in Section 5.3.2.
Q+1	An increase of 1 pallet size to Q
Q+2	An increase of 2 pallet sizes to Q
Q-1	A decrease of 1 pallet size to Q
Q-2	A decrease of 2 pallet sizes to Q
k+0.5	An increase of 0.5 to k
k+1	An increase of 1.0 to k
k-0.5	A decrease of 0.5 to k
k-1	A decrease of 1.0 to k

Table 28: Adjustments to Perform a Sensitivity Analysis Compared to the Base Policy

As can be seen in Figure 25, any adjustment for SKU 1 decreases the costs. When ordering two pallets extra each time, the costs are reduced by 13%. This can be explained by the relatively low unit price. With each order placed fixed costs are charged, which are compared to the variable unit price high. By increasing the order size, the total costs will therefore be lowered. Since the fillrate is also increasing, it is advised to adjust the order size to 360 items instead of 120 items. For SKU 2 the fillrate and the total costs do not change by adjusting the safety factor. If the order quantity is increased, the fillrate will remain equal but at a lower cost. However, since there is only a small cost difference and Ecorus has a maximum capacity at the warehouse, as described in 2.3, it is advised to keep the order size at 360 SKUs, as previously determined in Section 4.2.2.

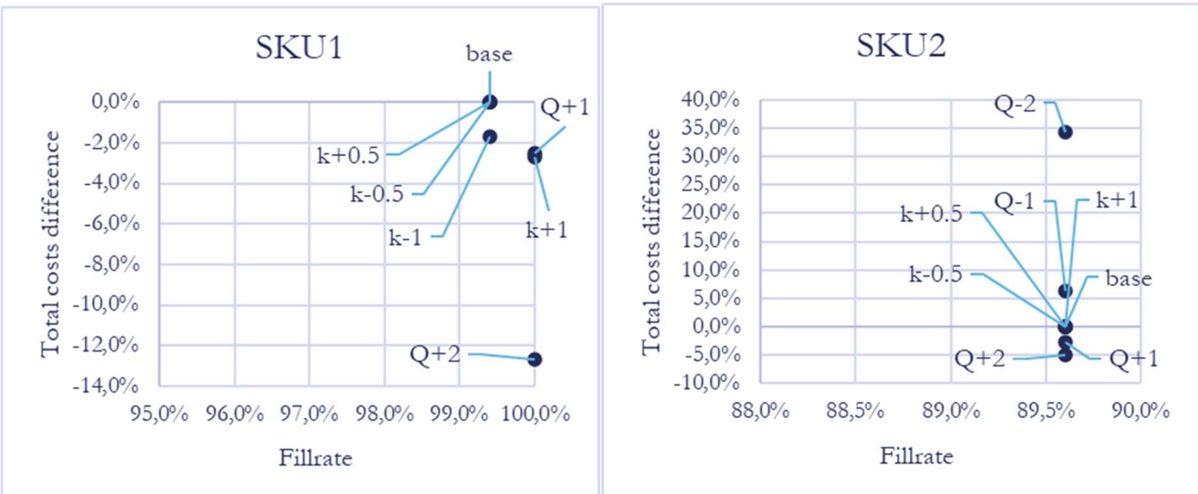

Figure 25: Sensitivity Analysis of SKUs 1 and 2

As can be seen in Figure 26, there is no change in both the fillrate as the costs when adjusting the safety factor or the order quantity for SKU 3. This can be explained by the absence of replenishment orders since the beginning on-hand inventory is sufficient to fulfill all demand during the period in which this policy is tested. Since there is no indication that adjustments of the parameters will be beneficial, it is advised to keep the base policy for SKU 3. When looking at SKU 4, it can be concluded that the base policy is most efficient in terms of costs, while the fillrate is also at the maximum of 100%. Therefore, the safety stock and the order quantity of SKU 4 should not be adjusted.

When analyzing Figure 27, it can be concluded that adjusting the order quantity for SKU 5 can decrease the costs by approximately 7%, without lowering the fillrate. The same reasoning as SKU 1 can be applied to SKU 5. Since the variable unit price is relatively low, financial benefits can be gained by ordering more pallets. It is therefore advised to adjust the order quantity of SKU 5 from 360 to 600. When looking at the costs compared to the fillrate for SKU 6, this same pattern can also be seen. Without changing the fillrate, lower costs can be obtained by ordering 1 or 2 more pallets at once. Since Ecorus has limited warehouse space and there is no cost difference between 1 and 2 extra pallets, it is advised to order only 1 pallet extra.

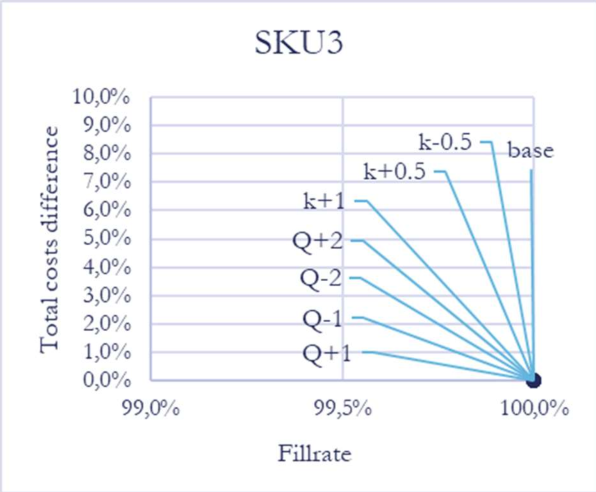

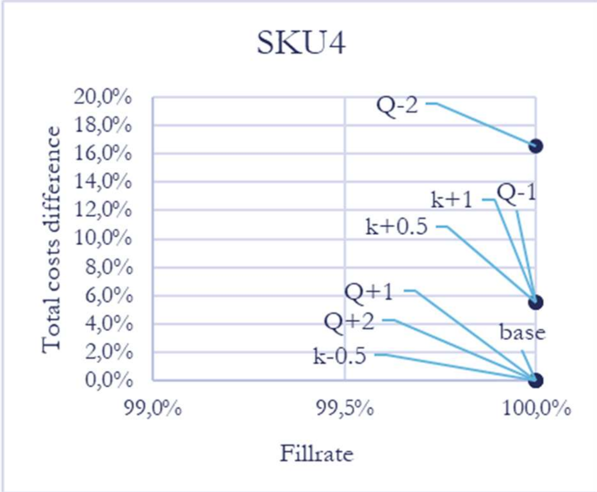

Figure 26: Sensitivity Analysis of SKUs 3 and 4

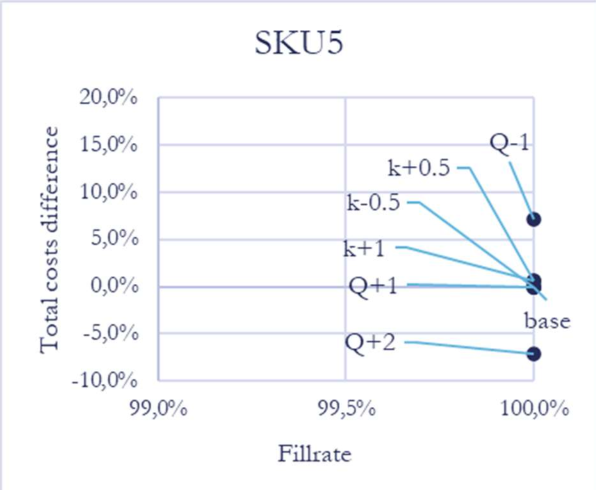

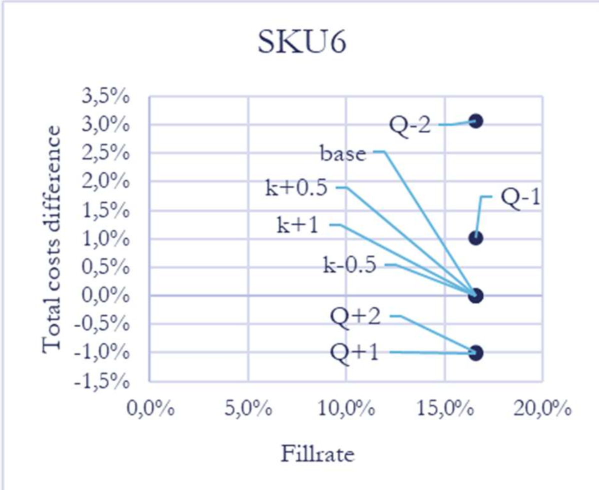

Figure 27: Sensitivity Analysis of SKUs 5 and 6

When analyzing Figure 28, it can be seen that financial gains as well as an increase in the fillrate can be obtained by ordering extra pallets of SKU 7. When ordering two additional pallets, a decrease of 26% of the costs can be achieved while the fillrate increases with 9,2%. As previously shown in Figure 24, with the base policy, there occur several stock-outs, combined with ordering new pallets most of the weeks. When ordering two more pallets at once, stock-outs occur less often in the first weeks, and the order costs are lower in the last weeks compared to the base policy as shown in Figure 24. It can be concluded that increasing the order size from 324 to 396 is extremely beneficial for SKU 7.

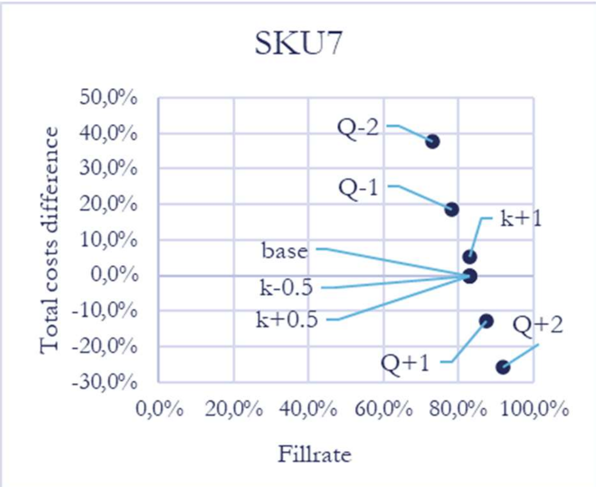

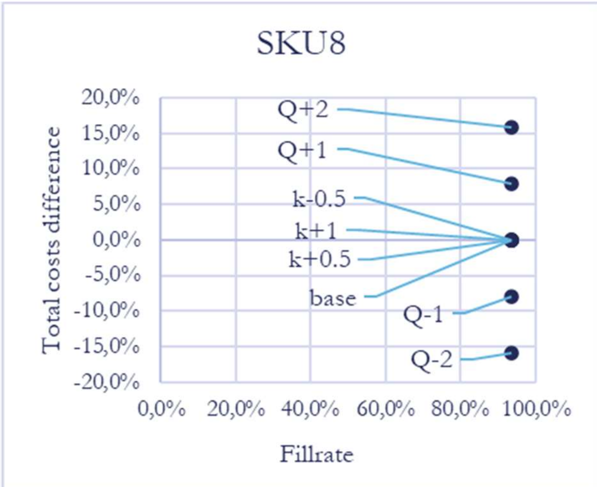

Figure 28: Sensitivity Analysis of SKUs 7 and 8

Looking at Figure 29, a different pattern than previously observed can be seen. For the Meters, the total costs will decrease if the number of pallets is also lowered by 1. However, since there is only a slight decrease of costs when this is lowered and higher prices occur when the number of pallets is lowered by 2, it might be risky to adjust the number of pallets ordered. It is therefore advised to keep the order quantity equal to the base policy. When looking at the invertors, it can be concluded that ordering 2 additional pallets result in an approximate decrease of 4% in costs, while the fillrate maintains at 100%. Regarding the invertors, it is advised to order 264 SKUs instead of 216 SKUs at once.

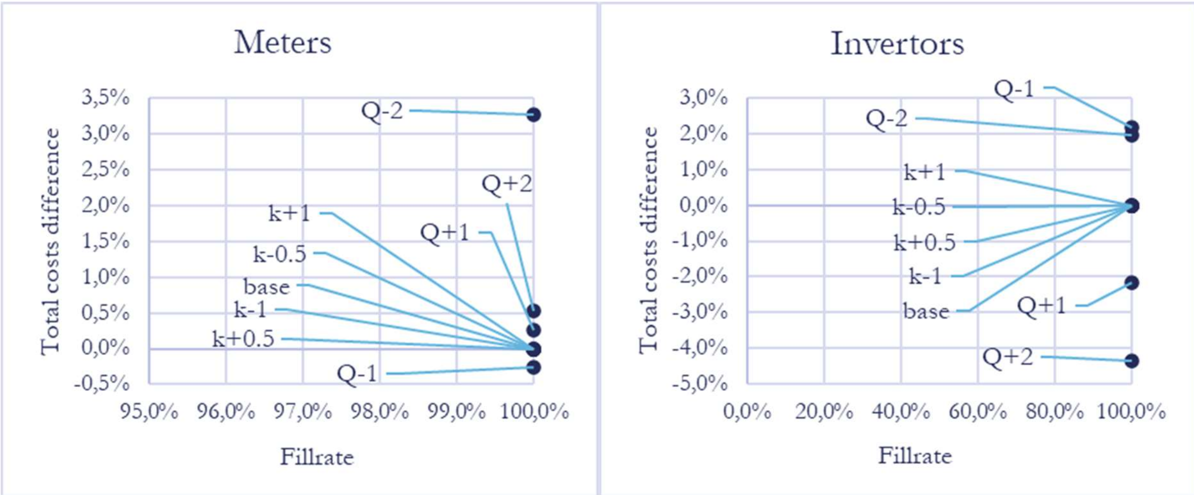

Figure 29: Sensitivity Analysis of Meters and Invertors

Figure 30 shows the influence of the safety factor and the order quantity on the Panels group. Neither of these adjustments influences the fillrate, however, ordering 1 pallet less results in approximately 6% cost reduction. Therefore, when ordering panels 930 panels should be ordered instead of 961.

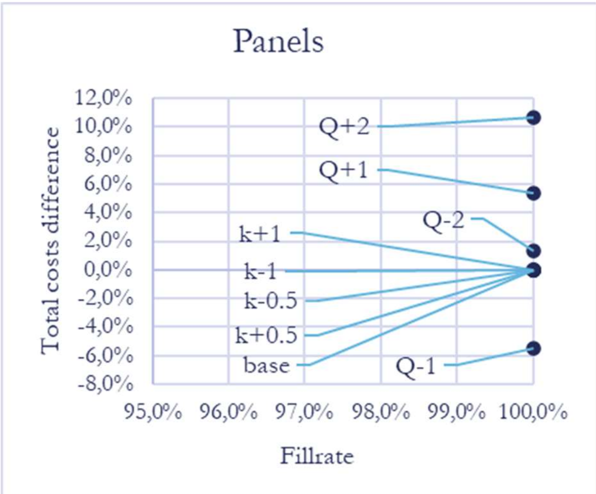

Figure 30: Sensitivity Analysis of Panels

5.4 Conclusion

In this section, the performance of various forecasting methods is evaluated to address the subquestion "What is the performance of the forecasts?" Seven different forecasting methods were analyzed, including Naïve, sNaïve, MA, SES, Holt, Holt-Winters, and ARIMA. The results of each method, as illustrated in Tables 14 to 20 and Figures 10 to 15, provide insights into their effectiveness across different SKUs and product groups. Among the methods evaluated, SES and ARIMA consistently emerged as the best performers, demonstrating robust performance across various SKUs and product groups. Notably, MA exhibited the second-best performance, particularly for SKU groups. While some methods like Naïve and sNaïve showed competitive performance for certain SKUs, their effectiveness varied depending on the strength of trend and seasonality.

Further analysis revealed that while parameter optimization significantly influenced the forecasting accuracy of methods such as MA and SES, Holt's and Holt-Winters exhibited suboptimal performance across various SKUs and SKU groups. This outcome underscores the sensitivity of these methods to parameter settings and highlights the challenges in fine-tuning parameters to achieve optimal forecasting outcomes. Despite efforts to adjust parameters, Holt and Holt-Winters methods consistently lagged behind SES and ARIMA in terms of forecast accuracy, indicating limitations in their ability to capture the patterns observed in Ecorus demand data.

Beginning with a discussion on predetermined parameters and the necessity of post-forecast adjustments, a comprehensive inventory policy for the SKUs was developed. A sensitivity analysis was conducted to assess the impact of parameter variations on total costs and fill rates across SKUs. The determination of safety stock, reorder points, and subsequent calculation of total expected yearly costs provided valuable insights into inventory management. However, discrepancies between observed and expected costs were uncovered for certain SKUs, notably SKU 6 and SKU 7, highlighting the need for parameter adjustments to improve performance.

Sensitivity analysis shed light on the effectiveness of parameter variations in optimizing costs and fill rates across SKUs. While adjustments yielded favorable outcomes in terms of decreased costs or increased fillrates for some SKUs, others remained relatively unaffected, underscoring the complexity of inventory management optimization.

Chapter 6 Conclusions, Recommendations, and Future Research

In this chapter, this research will be concluded in Section 6.1. Next, the research question "What recommendations can be given to Ecorus regarding the solution and its implementation?" will be answered in Section 0. Next, in Section 6.3, the last research question "What limitations are there in this research and what recommendations can be given to Ecorus for future research?" will be answered.

6.1 Conclusions

This research aimed to improve the forecasting and inventory management practices at Ecorus Home to better match their capacity with increasing demand. Through the analysis of various forecasting methods, several key findings emerged:

Best Performing Methods

The study revealed that SES and ARIMA consistently outperformed other forecasting methods. These methods demonstrated robust performance across different SKUs and product groups. ARIMA models can adapt to various patterns, including linear trends, seasonal fluctuations, and irregular variations, making them implementable for many different demand patterns. For SKUs that have a minimal difference in the performance of ARIMA and SES, for example SKU 6, it is advised to use SES, since SES is a straightforward and computationally efficient forecasting method that can provide reliable short-term predictions. Despite its popularity, the Holt-Winters method exhibited suboptimal performance in this context. The limited number of seasons available for analysis likely contributed to this outcome, aligning with literature recommendations to use at least 4 seasons worth of data to effectively use Holt-Winters. Table 29 gives an overview of the recommended forecasting method and the corresponding parameters.

KPI	SKU1	SKU2	SKU3	SKU4	SKU5	SKU6	SKU7	SKU8	Meters	Invertors	Panels
Method	SES	ARIMA	SES	SES	ARIMA	SES	Naïve	MA	ARIMA	MA	ARIMA
Parameters	(0.05)	(2; 0; 1)	(0.05)	(0.05)	(0; 0; 1)	(0.05)	N/A	2	(3; 0; 1)	10	(2; 0; 0)

Table 29: Advised Forecasting method per SKU

Insights into Parameter Optimization

Detailed analysis of parameter settings for methods like Moving Average, Holt's method, and Holt-Winters provided valuable insights. For SKUs characterized by low strength of trend, such as meters and panels, increasing the moving average window or smoothing parameters resulted in improved forecast accuracy. This observation can be attributed to the reduced capture of trend variations with larger window sizes, aligning with the nature of the demand patterns of the SKUs. The performance of Holt's linear trend model varied based on the α and β parameters. SKUs with a high strength of trend, like inverters and high-trend SKUs, benefitted from higher values of α, indicating the importance of capturing recent observations in reactive models. Conversely, SKUs with lower strength of trend exhibited better performance with lower α values, suggesting a balance between responsiveness and stability in trend modeling.

Implications for SKU Groups
Forecasting at the group level is a well-established practice known for its ability to reduce forecasting deviations through risk pooling. By aggregating demand data across related SKUs within a group, this approach capitalizes on the smoothing effect of combined demand patterns and trends. While this principle is widely recognized, this research highlights the tangible benefits of applying it in the context of Ecorus Home's inventory management. Interestingly, a notable disparity in performance between ARIMA and MA models when applied to SKU groups is observed. ARIMA consistently demonstrated superior performance, particularly in capturing and forecasting nuanced demand patterns, such as seasonal fluctuations and trend changes. While the difference between ARIMA and MA performance was marginal for individual SKUs, the advantages of forecasting at the group level became evident. This includes improved forecast accuracy, stability, and reliability by pooling demand patterns and trends across related SKUs within a group.

Sensitivity Analysis of Inventory Policy
Through examination of various parameters and adjustments of the inventory policies, Ecorus can gain invaluable insights into the responsiveness and effectiveness of its inventory management strategies. Sensitivity analysis empowers decision-makers to explore the impact of parameter variations on key performance metrics, such as total costs and fill rates. By systematically adjusting parameters, including order quantities and safety factors, Ecorus can assess the magnitude of these changes on inventory costs and operational performance. Furthermore, sensitivity analysis facilitates the identification of optimal parameter configurations that strike a balance between cost efficiency and service level attainment. By iteratively refining inventory policies based on sensitivity analysis results, Ecorus can fine-tune its inventory management practices to align with evolving market dynamics and demand patterns.

6.2 Recommendations

This section discussed the limitations of the current research and offers recommendations for future investigations aimed at advancing forecasting and inventory management practices at Ecorus Home. By identifying areas of improvement and outlining potential topics for future exploration, this chapter aims to guide future research.

Implement the Proposed Forecasting Methods
It is crucial to approach the implementation of forecasting methods with care, ensuring that the selected methods, such as SES and ARIMA, are performing as expected. Careful monitoring of forecast performances post-implementation is essential, allowing for the refinement of parameters used in the selected forecasting methods. This approach ensures ongoing optimization of forecasting accuracy and reliability. Regular reviews and post-implementation evaluations are recommended to identify areas for improvement and optimization. Furthermore, as mentioned in Section 6.1, there is insufficient data available to accurately initialize the parameters for Holt-Winter's. It is therefore recommended to analyze the different forecasting methods again once 4 years of data has been collected.

Continued Utilization and Adjustment of the Excel Tool
The Excel tool developed for this research serves as a valuable asset in streamlining forecasting and inventory management processes at Ecorus. Its user-friendly interface and robust functionality enable efficient data analysis, forecasting, and parameter optimization. Therefore, it is highly recommended to continue utilizing this tool as an integral part of the inventory management toolkit. Moreover, leveraging the insights gained from the tool's analyses, it is essential to implement a systematic approach to adjust parameters in response to evolving demand patterns and market dynamics. Regular reviews of forecast performances and inventory levels, facilitated by the Excel tool, should inform parameter adjustments aimed at optimizing forecasting accuracy, inventory turnover, and cost-effectiveness.

Forecast Demand for Other Departments
While the research focuses on Ecorus Home, extending the application of forecasting methods to forecast demand for shared SKUs across multiple departments is recommended. Forecasting demand for shared SKUs across departments reduces variability and aligns ordering processes, resulting in potential cost savings and operational efficiencies.

Align Orders Across Departments
The coordination of orders for shared SKUs across departments minimizes variability and optimizes inventory management practices. Therefore, aligning orders reduces ordering costs, optimizes inventory levels, and ensures consistent supply chain operations.

Negotiate Pallet Size
For some SKUs, the Economic Order Quantity is significantly lower than the minimum order quantity, due to the number of items on a pallet. This is for example the case for SKU 1, where the EOQ is 16 compared to a pallet size of 160. In this research lower values of Q are only analyzed when Q was bigger than the pallet size, therefore it is recommended to calculate the expected difference in costs when the EOQ is used compared to when Q is used. When this difference is significant, it is recommended to negotiate if partial pallets can be ordered, or alternatively, find another supplier that can deliver less items.

By implementing these recommendations, Ecorus Home can enhance its forecasting accuracy, optimize inventory management practices, and improve overall supply chain performance, ultimately leading to increased operational efficiency.

6.3 Limitations and Future Research

This section discusses the limitations of the current research and offers recommendations for future investigations aimed at advancing forecasting and inventory management practices at Ecorus Home. By identifying areas of improvement and outlining potential topics for future exploration, this chapter aims to guide future research.

Limited Historical Data Availability
One of the primary limitations of the current research lies in the availability of historical data. The exclusion of numerous SKUs due to insufficient historical data may have introduced potential gaps in insights. Additionally, the analysis was constrained by a maximum of two years of available data, which could have limited the robustness of the findings.

Computational Constraints
The research encountered computational constraints, particularly in the analysis of forecasting methods. Due to limited computational availability, the exploration of ARIMA models was restricted to AR and MA orders up to 3. This limitation may have constrained the accuracy of forecasts. Furthermore, the utilization of machine learning models for pattern recognition was hindered by computational restrictions, which could have provided additional insights and accuracy. Another limitation was the significant computational effort required, particularly in initializing parameters for ARIMA models. Analyzing 114 weeks of data for parameter initialization consumed considerable time. To mitigate this, future research could explore reducing the time window from weeks to months. By forecasting on a monthly level, the time required for parameter initialization could be substantially reduced, potentially expediting the identification of the best-performing method. Alternatively, additional computer power can be temporarily rented.

Enhanced Data Integration
Future research should prioritize the integration of demand data and the collection of historical data to enhance forecasting accuracy. By examining SKUs that are phased in and out with each other and combining SKUs that replace each other, more stable trends and seasonality can be identified, leading to more accurate forecasts. Additionally, leveraging advanced demand information at the housing associations to disaggregate SKU group forecasts into individual SKU forecasts could provide more granular insights into demand patterns.

Improvement of Forecasting Methods
There is a need for further investigation into forecasting methods, particularly ARIMA models. Expanding the analysis to include higher orders of AR and MA and exploring the application of differencing could enhance forecast accuracy. Furthermore, easing computational restrictions to enable the use of machine learning models for pattern recognition could unlock additional insights and improve forecasting accuracy.

Exploration of Alternative Inventory Policies
Future research could explore alternative inventory policies beyond the (s, Q) policy. Investigating the costs and fill rates of an (R, s, Q) policy could lower personnel costs. However, for the current way of working at Ecorus the (s, Q) policy is most suitable. Therefore, investigating an (R, s, Q) policy should only be done if it would be acceptable to Ecorus to change their ordering approach.

Disaggregation of Forecasts

An essential avenue for future research is the development of methodologies to disaggregate forecasts at a more granular level. Specifically, exploring the use of weighted factors to disaggregate SKU group forecasts into individual SKU forecasts could enhance the accuracy and precision of demand predictions. By assigning appropriate weights based on historical factors such as sales volume, product characteristics, and market trends, the disaggregated forecasts can more accurately reflect the demand patterns of individual SKUs within a group. This approach would enable Ecorus Home to make more informed inventory management decisions tailored to the specific needs and characteristics of each SKU, ultimately improving overall supply chain efficiency and client satisfaction.

References

Abdel-Aal, M., El-Sharief, M., El-Deen, A., & Nassr, A.-B. (2010). A Framework For Evaluating and Comparing Inventory Control Policies in Supply Chains. *Journal of Engineering Sciences*, *38*, 449–465. https://doi.org/10.21608/jesaun.2010.124377

Ammar, O. Ben, Dolgui, A., Hnaien, F., & Louly, M. A. (2013). Supply planning and inventory control under lead time uncertainty: A review. *IFAC Proceedings Volumes*, *46*(9), 359–370. https://doi.org/10.3182/20130619-3-RU-3018.00592

Anderson, O. D. (1977). The Box-Jenkins approach to time series analysis. *RECHERCHE OPÉRATIONNELLE*, *11*(1), 3–29.

Athanasopoulos, G., Ahmed, R. A., & Hyndman, R. J. (2009). Hierarchical forecasts for Australian domestic tourism. *International Journal of Forecasting*, *25*(1), 146–166. https://doi.org/10.1016/j.ijforecast.2008.07.004

Axsäter, S. (2006). *Inventory Control*. Springer US. https://doi.org/10.1007/0-387-33331-2

Billah, B., King, M. L., Snyder, R. D., & Koehler, A. B. (2006). Exponential smoothing model selection for forecasting. *International Journal of Forecasting*, *22*(2), 239–247. https://doi.org/10.1016/j.ijforecast.2005.08.002

Brealey, R. A., Myers, S. S., Allen, F., & Edmans, A. (2022). *Principles of Corporate Finance* (14th ed.). McGraw Hill.

Breiman, L. (2001). Random Forests. *Machine Learning*, *45*(1), 5–32. https://doi.org/10.1023/A:1010933404324

Chai, T., & Draxler, R. R. (2014). Root mean square error (RMSE) or mean absolute error (MAE)? – Arguments against avoiding RMSE in the literature. *Geoscientific Model Development*, *7*(3), 1247–1250. https://doi.org/10.5194/gmd-7-1247-2014

Chicco, D., Warrens, M. J., & Jurman, G. (2021). The coefficient of determination R-squared is more informative than SMAPE, MAE, MAPE, MSE and RMSE in regression analysis evaluation. *PeerJ Computer Science*, *7*, e623. https://doi.org/10.7717/peerj-cs.623

Chopra, S., & Meindl, P. (2015). *Supply Chain Management : Strategy, Planning, and Operation* (6th ed.). Pearson.

Devabhaktuni, V., Alam, M., Shekara Sreenadh Reddy Depuru, S., Green, R. C., Nims, D., & Near, C. (2013). Solar energy: Trends and enabling technologies. *Renewable and Sustainable Energy Reviews*, *19*, 555–564. https://doi.org/10.1016/J.RSER.2012.11.024

Dhoka, D. K., & Choudary, Y. L. (2013). "XYZ" Inventory Classification & Challenges. *IOSR Journal of Economics and Finance*, *2*(2), 23–26. www.iosrjournals.org

Ecorus. (2022). *Ecorus*. https://ecorus.com/

Ernst, R., & Cohen, M. A. (1990). Operations related groups (ORGs): A clustering procedure for production/inventory systems. *Journal of Operations Management*, *9*(4), 574–598. https://doi.org/10.1016/0272-6963(90)90010-B

European Commission. (2022, September 13). *In focus: Solar energy – harnessing the power of the sun*. https://commission.europa.eu/news/focus-solar-energy-harnessing-power-sun-2022-09-13_en#:~:text=Solar%20power%20is%20a%20cheap,accessible%20one%20for%20European%20households

García, F. P., Pedregal, D. J., & Roberts, C. (2010). Time series methods applied to failure prediction and detection. *Reliability Engineering and System Safety*, *95*(6), 698–703. https://doi.org/10.1016/j.ress.2009.10.009

Gross, C. W., & Sohl, J. E. (1990). Disaggregation methods to expedite product line forecasting. *Journal of Forecasting*, *9*(3), 233–254. https://doi.org/10.1002/for.3980090304

Heerkens, H. (2017). *Validity*. ITC E-Learning. https://vimeo.com/showcase/2938606/video/117885780

Heerkens, H., & van Winden, A. (2021). *Solving Managerial Problems Systematically*. Routledge. https://doi.org/10.4324/9781003186038

Ho, S. L., & Xie, M. (1998). The use of ARIMA models for reliability forecasting and analysis. *Computers & Industrial Engineering*, *35*(1–2), 213–216. https://doi.org/10.1016/S0360-8352(98)00066-7

Holt, C. C. (2004). Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting*, *20*(1), 5–10. https://doi.org/10.1016/J.IJFORECAST.2003.09.015

Hyndman, R. J., Ahmed, R. A., Athanasopoulos, G., & Shang, H. L. (2011). Optimal combination forecasts for hierarchical time series. *Computational Statistics & Data Analysis*, *55*(9), 2579–2589. https://doi.org/10.1016/j.csda.2011.03.006

Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and Practice*. Otext.

Isidore, C. (2022, June 6). *Why US gas prices are at a record, and why they'll stay high for a long time*. CNN Business. https://edition.cnn.com/2022/06/06/energy/record-gas-prices-causes/index.html

Janssen, F., Heuts, R., & de Kok, T. (1996). *The value of information in an (R,s,Q) inventory model*. https://www.researchgate.net/publication/4783037_The_value_of_information_in_an_RsQ_inventory_model

Janssen, F., Heuts, R., & de Kok, T. (1998). On the (R, s, Q) inventory model when demand is modelled as a compound Bernoulli process. *European Journal of Operational Research*, *104*(3), 423–436. https://doi.org/10.1016/S0377-2217(97)00009-X

Jenkins, G., Box, G., & Reinsel, G. (1994). *Time Series Analysis: Forecasting and Control* (G. Box, G. Jenkins, & G. Reinsel, Eds.; 3rd ed.). Prentice Hall .

Kapgate, D. (2014). *Weighted Moving Average Forecast Model based Prediction Service Broker Algorithm for Cloud Computing*. https://api.semanticscholar.org/CorpusID:30396547

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems* (Vol. 30). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/6449f44a102fde848669bdd9eb6b76fa-Paper.pdf

Khair, U., Fahmi, H., Hakim, S. Al, & Rahim, R. (2017). Forecasting Error Calculation with Mean Absolute Deviation and Mean Absolute Percentage Error. *Journal of Physics: Conference Series*, *930*, 012002. https://doi.org/10.1088/1742-6596/930/1/012002

Lordahl, A. E., & Bookbinder, J. H. (1994). Order-statistic calculation, costs, and service in an (s, Q) inventory system. *Naval Research Logistics (NRL)*, *41*(1), 81–97. https://doi.org/https://doi.org/10.1002/1520-6750(199402)41:1<81::AID-NAV3220410106>3.0.CO;2-9

Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2020). The M4 Competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting*, *36*(1), 54–74. https://doi.org/10.1016/j.ijforecast.2019.04.014

Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). M5 accuracy competition: Results, findings, and conclusions. *International Journal of Forecasting*, *38*(4), 1346–1364. https://doi.org/10.1016/j.ijforecast.2021.11.013

McHugh, M. L. (2013). The Chi-square test of independence. *Biochemia Medica*, 143–149. https://doi.org/10.11613/BM.2013.018

Netherlands Enterprise Agency (RVO). (2020). *Next-generation solar power Dutch technology for the solar energy revolution*. https://www.rvo.nl/sites/default/files/2020/08/NL-Solar-Guide-2020.pdf

Nowotynska, I. (2013). An Application of XYZ Analysis In Company Stock Management. *Modern Management Review*, *XVIII*(20), 77–86.

Paldino, G. M., De Stefani, J., De Caro, F., & Bontempi, G. (2021). Does AutoML Outperform Naive Forecasting? *The 7th International Conference on Time Series and Forecasting*, 36. https://doi.org/10.3390/engproc2021005036

Ramos, P., Santos, N., & Rebelo, R. (2015). Performance of state space and ARIMA models for consumer retail sales forecasting. *Robotics and Computer-Integrated Manufacturing*, *34*, 151–163. https://doi.org/10.1016/j.rcim.2014.12.015

Rijksoverheid. (2022a, April 11). *Verhuurderheffing: wetsvoorstel dat afschaffing regelt open voor reactie*. https://www.rijksoverheid.nl/actueel/nieuws/2022/04/11/verhuurderheffing-wetsvoorstel-dat-afschaffing-regelt-open-voor-reactie

Rijksoverheid. (2022b, June 1). *Duurzaam wonen voor iedereen*. https://www.rijksoverheid.nl/actueel/nieuws/2022/06/01/duurzaam-wonen-voor-iedereen

Rokach, L., & Maimon, O. (2005). Decision Trees. In *Data Mining and Knowledge Discovery Handbook* (pp. 165–192). Springer-Verlag. https://doi.org/10.1007/0-387-25465-X_9

RTL Nieuws. (2023, January 27). *Historische groei zonnepanelen: "Gebruik overdag meer stroom."* https://www.rtlnieuws.nl/tech/artikel/5362036/zonnepanelen-nederland-zonnepaneel-elektriciteit-stroom-netwerk

Salkind, N. (2010). *Encyclopedia of Research Design*. SAGE Publications, Inc. https://doi.org/10.4135/9781412961288

Sévano, V. (2022, June 9). *The Netherlands has the most solar panels per capita in Europe*. https://www.iamexpat.nl/expat-info/dutch-expat-news/netherlands-has-most-solar-panels-capita-europe

Silveira Netto, C. F., Bahrami, M., Brei, V. A., Bozkaya, B., Balcisoy, S., & Pentland, A. P. (2023). Disaggregating sales prediction: A gravitational approach. *Expert Systems with Applications*, *217*, 119565. https://doi.org/10.1016/j.eswa.2023.119565

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). *Inventory and Production Management in Supply Chains*. CRC Press. https://doi.org/10.1201/9781315374406

Spiliotis, E., Makridakis, S., Semenoglou, A. A., & Assimakopoulos, V. (2022). Comparison of statistical and machine learning methods for daily SKU demand forecasting. *Operational Research*, *22*(3), 3037–3061. https://doi.org/10.1007/s12351-020-00605-2

Supply Value. (2022). *Supply Value*. https://www.supplyvalue.nl/

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

Tollefson, J. (2022). What the War in Ukraine Means for Energy, Climate and Food. *Springer Nature Limited*, *604*, 232–233. https://www.nature.com/articles/d41586-022-00969-9

UWV. (2022, February 24). *Cijfers arbeidsmarkt: Krapte op alle niveaus*. https://www.uwv.nl/nl/achtergrond/cijfers-arbeidsmarkt-krapte-op-alle-niveaus

van Gastel, E. (2023, January 17). *Netbeheerders: salderingsregeling zonnepanelen haaks op wens om energierekening betaalbaar te houden*. Solar Magazine. https://solarmagazine.nl/nieuws-zonne-energie/i28784/netbeheerders-salderingsregeling-zonnepanelen-haaks-op-wens-om-energierekening-betaalbaar-te-houden#:~:text='De%20salderingsregeling%20heeft%20gezorgd%20voor,ons%20land%20van%20zonnepanelen%20voorzien.

van Kampen, T. J., Akkerman, R., & Pieter van Donk, D. (2012). SKU classification: a literature review and conceptual framework. *International Journal of Operations & Production Management*, *32*(7), 850–876. https://doi.org/10.1108/01443571211250112

Vandeput, N. (2021). *Data Science for Supply Chain Forecasting*. De Gruyter. https://doi.org/10.1515/9783110671124

Waters, D. (2003). *Inventory Control and Management* (2nd ed.).

Williams, T. M. (1984). Stock Control with Sporadic and Slow-Moving Demand. *The Journal of the Operational Research Society*, *35*(10), 939. https://doi.org/10.2307/2582137

Winters, P. R. (1960). Forecasting Sales by Exponentially Weighted Moving Averages. *Management Science*, *6*(3), 324–342. https://doi.org/10.1287/mnsc.6.3.324

Appendix A Decoding Key

Suppliers

S.1. .
S.2. .
S.3. .
S.4. .
S.5. .
S.6. .
S.7. .
S.8. .
S.9. .
S.10. .
S.11. .
S.12. .

Appendix B Selection of SKUs

SKU	Number of weeks without demand	Weeks with demand in 2023	Total demand	First introduction of SKU	Include?	Name used in research
					Yes	SKU 1
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
					Yes	SKU 2
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
					Yes	SKU 3
					Yes	SKU 4
					Yes	SKU 5
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
					Yes	SKU 6
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
					No	SKU 7
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
					Yes	SKU 8
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	
[redacted]					No	

Table 30: Selection of SKUs based on the times demand occurred and the total demand

Appendix C Forecasting Models Formulation

Method	Model
Naïve	$\hat{y}_{T+h\|T} = y_T$
Seasonal naïve	$\hat{y}_{T+h\|T} = y_{T+h-m(k+1)}$
SMA	$\hat{a}_t = \hat{x}_{t.N} = (x_t + x_{t-1} + x_{t-2} + \ldots + x_{t-N+1})/N$
SES	$\hat{a}_t = \alpha x_t + (1-\alpha)\hat{a}_{t-1}$
Holt	$\hat{a}_t = \alpha\, x_t + (1-\alpha)(\hat{a}_{t-1} + \hat{b}_{t-1})$ $\hat{b}_t = \beta\,(\hat{a}_t - \hat{a}_{t-1}) + (1-\beta)\,\hat{b}_{t-1}$ $S = \sum_{t=-n+1}^{0} \left[x_t - (\hat{a}_0 + \hat{b}_0 t)\right]^2$ $\hat{a}_0 = \frac{6}{n(n+1)} \sum_t t x_t + \frac{2(2n-1)}{n(n+1)} \sum_t x_t$ $\hat{b}_0 = \frac{12}{n(n^2-1)} \sum_t t x_t + \frac{6}{n(n+1)} \sum_t x_t$
Holt-winters	$\hat{a}_t = \alpha(x_t/\hat{F}_{t-P}) + (1-\alpha)(\hat{a}_{t-1} + \hat{b}_{t-1})$ $\hat{b}_t = \beta(\hat{a}_t - \hat{a}_{t-1}) + (1-\beta)\,\hat{b}_{t-1}$ $\hat{F}_t = \gamma(x_t/\hat{a}_t) + (1-\gamma)\hat{F}_{t-P}$
ARIMA	$x_t = \varphi_1 x_{t-1} + \varphi_2 x_{t-2} + \ldots + \varphi_p x_{t-p} + \epsilon_t + \theta_1 \epsilon_{t-1} + \theta_2 \epsilon_{t-2} + \ldots + \theta_q \epsilon_{t-q}$

Table 31: Used Forecasting Methods and their Models