University of Twente
Faculty of Behavioural, Management and Social Sciences

Industrial Engineering and Management

Modelling uncertainty in a configurable environment: a case study at Pan Oston.

Author:
Dhr. D.M. Overmars

Supervisors University of Twente:
Dr. E. Topan
Dr. I. Seyran Topan

Supervisor Pan Oston:
Dhr. D. Broekhof - Supply Chain Manager

Date:
August 20, 2021

UNIVERSITY OF TWENTE.

Pan Oston®

HEART OF STEEL / SENSE OF DESIGN

Management Summary

Introduction

This study is conducted at Pan Oston, which is a company specialized in designing and producing check-outs, self-service desks and kiosk solutions. Changing consumer demands force Pan Oston to find new and innovative solutions, but on the other hand there is a need to remain competitive and serve demand within the requested date of customers. In order to achieve this goal, finished goods need to be produced from components that are available from stock which means decisions on which components to stock are required. But, Pan Oston feels that inventory levels are high because of the need to meet demand of their customers.

Problem statement

We use a problem cluster to identify the core problem from the problem Pan Oston perceived. We identified that the perceived problem is caused by three main core problems: an unreliable demand forecast, uncertainty in product configuration requirements of customers and inventory control policy parameters determined by gut feelings. Based on the formulated main problems, we define the following research question:

How can Pan Oston use available historical and future sales data to improve their inventory levels of components, while maintaining their current service levels?

Approach

First of all, we analyse the current situation with data collection methods such as interviews and data analysis. As Pan Oston is a project-based company, selling their products to other companies only, the sales forecast is based on employees that are heavily involved with customers to derive their future plans. The result is a sales forecast aggregated on product groups, where we can distinguish three distinct time zones during the forecast horizon: the first consists of customer orders only, the second is a combination of customer orders and forecast and the third consists of forecasts only. Because of the order configuration uncertainty, it is hard to translate the demand on product group level to component requirements, especially when customer orders are not yet final.

The literature describes multiple ways to classify demand patterns, forecast demand and handle order configuration uncertainty. First, we describe Supply Chain Management in general and introduce the different types of uncertainty in Assemble-To-Order (ATO), Configure-To-Order (CTO) and Forecast-To-Order (FTO) strategies. After this, we focus on the demand planning process and find out that both quantitative and qualitative methods can be used for this purpose. Furthermore, we investigate how others have been able to model the uncertainty in the supply planning phase. Finally, we find a classification where products are divided into four categories based on their demand patterns, namely smooth, erratic, intermittent and lumpy. The classification indicates how hard it is to manage and forecast a component, which we use during our solution test.

Next, we develop our own method. From our analysis of the current situation, we conclude that the aggregate forecast on product group level is important in determining the component requirements. We choose to calculate product configuration probabilities from delivered configurations retrieved from past production and sales orders and use the aggregate forecast as input to determine the total component demand. In literature, most authors proposed base-stock levels and assumed stationary demand for tractability. We agree on the reasoning that base-stock levels are sufficient for the supply planning phase and that operational issues such as order sizes and Minimum Order Quantities (MOQs) can be explored later. However, we identify that demand at Pan Oston is non-stationary, since demand on product level often comes from small or large projects. We decide to include this in our model and generate base-stock policies per month. We use general expressions to calculate the component and finished good fill rates and propose a greedy heuristic to determine base-stock levels.

Results

We use different classifications to show our results and measure the performance of our forecast using the Mean Square Error (MSE), Mean Absolute Deviation (MAD), bias and the Weighted Absolute Percent Error (WAPE). First, we can distinguish common and customer-specific components and see that our proposed solution achieves an equal accuracy for both types. As described, we used the classification found in the literature to classify each component as smooth, erratic, intermittent or lumpy. From the analysis of these patterns, we conclude that our model is able to forecast component demand of all demand patterns, but it has the general tendency to have a negative bias which we can explain as the result of lacking demand forecasts for some product groups. From Table 1 we can conclude that a better accuracy is achieved when we only use the actual component consumption of the product groups forecasted. However, human knowledge and interaction are still needed since it is impossible for our model to capture component introductions or the end-of-life phase of components.

Table 1: Performance forecasting models

	MAD	RMSE	Bias	WAPE
Smooth	38	48	-17	18%
Erratic	51	60	-44	30%
Intermittent	12	20	2	*
Lumpy	13	18	-10	*

*** Undefined, dividing by zero**

After validation of the forecast performance, we perform a simulation with the actual consumption of components to study the performance of our proposed base-stock levels. In this analysis, we only include the components that are already controlled with a base-stock policy. The simulation shows that the proposed base-stock levels result in lower inventory costs. We achieve a decrease of 48.9% for components with a smooth demand pattern, 39.1% for components classified as erratic, 34.2% for intermittent components and 42.2% for components with a lumpy pattern. With our simulation study, we are also able to compare the expected fill rate and the simulated fill rate of all components. As a result of the unpredictable demand, we can see that some components have very large differences. Based on earlier configurations and the forecast of demand on product level, our model expects that demand for some components is low, which means a low base-stock level can guarantee a high fill rate. When demand for a component is higher than expected, the actual fill rate will be way lower. From this, we can conclude that our quantitative inventory model provides a good basis for parameter settings, but qualitative information should still be taken into account. The figure below presents the differences between the simulated and expected fill rate and the inventory savings for each of the components.

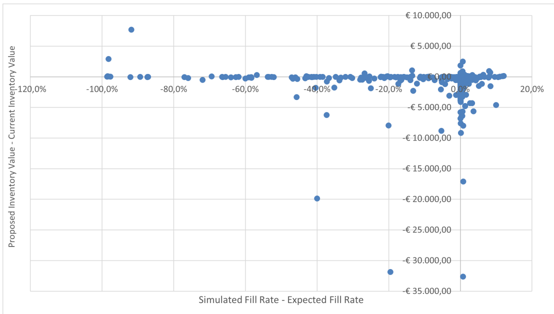

Figure 1: Fill Rate difference versus inventory decrease

We end our solution test with a sensitivity analysis. In the validation above we assume that the sales forecast on product group level has no uncertainty and modelled it as fixed but we are also interested in how our method behaves with small deviations in the aggregated forecast on product group level. We evaluate 12 scenarios in which we increase and decrease the forecast for six customers with 10%. Finally, we explore how marginal increases in the target fill rates influenced the average inventory value. From this analysis, we can conclude that a trade-off between the desired fill rate and average inventory value should be made. When trying to approach a fill rate close to 100%, it will yield a high increase in the average inventory value, but will almost have no yield in material availability in comparison with lower fill rates.

Recommendations

We recommend Pan Oston to implement our proposed method by using the tool developed. Despite the worse performance for some components with unpredictable demand, the proposed base-stock levels can always be used as reference values. Before this tool can be put into practice, a project team should agree on the product groups and make guidelines to make sure newly introduced products are assigned to the correct product group in the future. The aggregated sales forecast should contain exactly the groups as agreed and the groups should also be implemented in the ERP system of Pan Oston. As our tool uses target fill rates for each product group to determine the base stock levels, the team should also agree on these. Finally, it is important that the used historical order data remains available, also after implementation of the new ERP-system. Further research can focus on the effect and incorporation of demand uncertainty of finished goods on product group level and the lead time uncertainty of components. Besides, the effect of including assembly lead times can be investigated.

Preface

Dear reader,

With great pleasure I present you my research conducted at Pan Oston to obtain the Master's degree in Industrial Engineering & Management, with a specialization in Production & Logistics Management, at the University of Twente. This graduation assignment is the final chapter of my study time. I want to express my gratitude to some people who helped and guided me during the thesis and my studies.

I would like to thank all the people from Pan Oston for welcoming me, but also for their contribution to this research. A special thanks to Davy Broekhof for creating this thesis opportunity and his guidance during the internship.

Furthermore, I want to thank my supervisors from the University of Twente, Engin Topan and Ipek Topan. Engin was involved in this research from the very beginning and provided valuable feedback and input during this research. Ipek joined later and helped me constructing the thesis in the right way and lift the research to a higher level. I am thankful for their online supervision during the COVID-19 pandemic.

Lastly, I want to thank my family, friends and colleagues from the university who supported me during my studies. You helped me to keep the most important goals insight and were always willing to discuss my thesis and concerns whenever I needed it.

All that remains for me is to wish you a pleasant read!

Daimen Overmars
Raalte, August 20, 2021

List of Abbreviations

Abbreviation	Definition	Introduced on page
ADI	Average Demand Interval	20
AX	Microsoft Dynamics AX	4
BOM	Bill of Material	12
CODP	Customer Order Decoupling Point	17
CTO	Configure-To-Order	1
CV^2	Coefficient of Variation	20
ERP	Enterprise Resource Planning AX	3
ETO	Engineer-To-Order	1
FTO	Forecast-To-Order	1
MOQ	Minimum Order Quantity	4
MTO	Make-To-Order	11
MTS	Make-To-Stock	11
ZZV	Stock-keeping sub-assembly	8

List of Figures

Contents

1 Introduction

This chapter introduces the Master Thesis assignment performed at Pan Oston to complete my master in Industrial Engineering and Management. In Section 1.1, the company Pan Oston is shortly introduced. Section 1.2 describes the motivation of this research. Next, Section 1.3 describes the problem statement and identifies the core problems. Section 1.4 defines the research objective. Finally, Section 1.5 presents the research design.

1.1 Company introduction

Pan Oston is a company specialized in designing and producing check-outs, self-service desks and kiosk solutions. Besides the production facility in Raalte, where the headquarters are located, Pan Oston also outsources their activities to a partner in Slovakia. At their own facilities, around 170 people are employed. The company Pan Oston has its origin in Finland as Seppo Halttunen founded the company Halton in 1969. In 2005, the production facility and office moved to the Netherlands. Since then, the company Pan Oston continued as an independent company. Nowadays, Pan Oston is active on the European and Scandinavian market with a focus on the Netherlands, Germany, Belgium, the United Kingdom and Denmark (Pan Oston B.V., 2020).

The company produces their products using three main types of manufacturing principles, which they classified as Configure-To-Order (CTO), Engineer-To-Order (ETO) and Forecast-To-Order (FTO). In all policies, some activities are performed when they are requested by the customer. However, the policies differ in terms of variations and customizations. In the CTO policy, products are configured from sub-assemblies and components available from stock. In an ETO policy, a new product is made entirely based on the desired specifications of a customer, which means products should be engineered first. In the FTO policy forecasts of customers are used to match sub-assembly inventory with anticipated customer demand. This way, the final products are finished on time. Within Pan Oston, a customer is classified as FTO when products are standardized for this customer. When products are configured from standardized components, resulting in an unique finished good, the customer is classified as CTO. The sub-assembly of a FTO customer is already customized for this specific customer, where the sub-assembly of a CTO product is still generic.

Most of the production activities of CTO and ETO are performed in Raalte. As the ETO policy requires engineering activities, the demo, prototype- and zero series are always produced in Raalte. Besides, Pan Oston promises a maximum lead time of 6 weeks for other products, which is only achieved when production is performed locally. If the requested delivery date is feasible for production in Slovakia, the decision is made to outsource these activities. As FTO customers provide forecasts of their demand, the required sub-assemblies are produced in Slovakia in advance. The product is finished in Raalte or Slovakia, depending on the requested lead time and customer location.

1.2 Research motivation

The mission of Pan Oston is to find new and innovative solutions together with their customers. In executing the daily activities, they focus on giving their customers the best service possible. In the market where Pan Oston is active, it is critical to serve demand within the requested date of customers. Currently, Pan Oston experiences difficulties to fulfill the demand of their customers since it is hard to predict the demand of their products.

Customers do not provide long-term forecasts on the number of stores and the required products. Also concepts of stores are changed frequently, which could lead to a change in required products per store. Besides the lack of a long-term forecast, shifted requested delivery dates occur frequently. As a customer-oriented company, Pan Oston uses safety stock to avoid stoppages in the production process caused by stock-outs. However, stock-outs still occur. Extra costs are made to solve these shortages, for example by the arrangement of expedited deliveries. Table 1.1 shows the increased inventory and the number of stock-outs over the last five years.

Table 1.1: Turnover, inventory value (million euros) and number of stock-outs

Year	2015	2016	2017	2018	2019	2020
Turnover	17.83	22.34	30.10	31.94	35.25	30.78
Inventory value	2.06	2.61	3.44	4.15	5.14	6.04
Number of stock-outs	*	*	*	835	1159	990

*** Data not collected**

Recently, Pan Oston formulated their plans and one of their goals is to achieve a higher inventory turnover while maintaining customer service levels. To achieve this goal, they feel that more improvements are required. They currently experience that inventory levels are high and that not always the right components or sub-assemblies are produced or purchased. Because of this, they are interested in methods to improve their supply chain to handle uncertain demand better. With an improved forecast of the components, Pan Oston should be able to respond to customer demand without having high inventory levels and the risk of potential obsoletes or stock-outs.

1.3 Problem statement

We use a problem cluster to identify the core problem. In the problem cluster, the causal relations are given. The problem cluster is depicted in Figure 1.1. The observed action problem at Pan Oston is the feeling that inventory levels are high, which is marked in green. This is caused by the problems colored in blue (core problem) and red (non-core problem). In this section, we describe both types of problems and their effects.

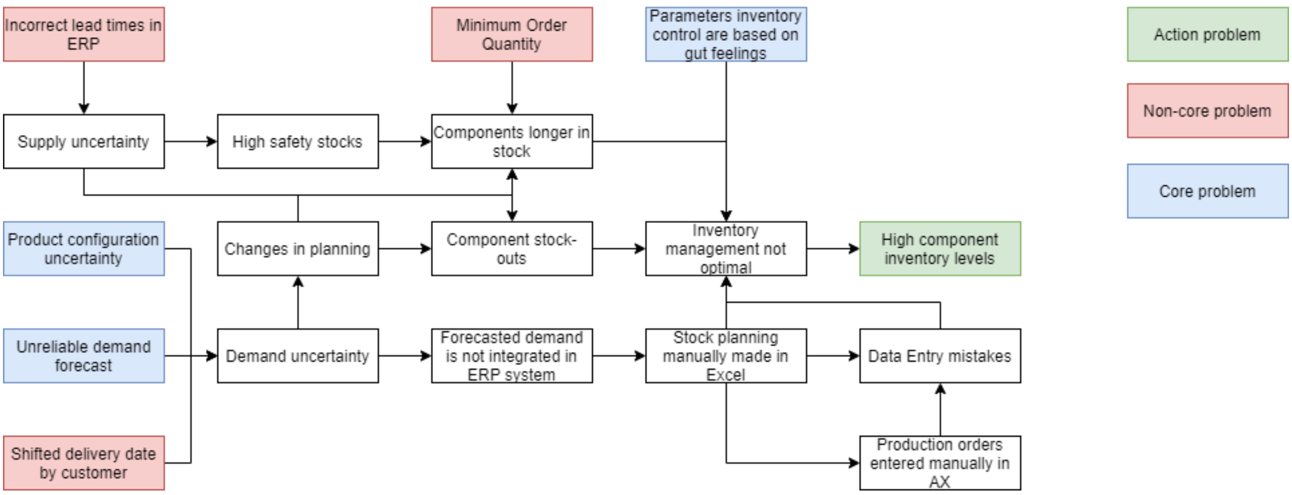

Figure 1.1: Problem Cluster

Unreliable demand forecast

Pan Oston experiences difficulties in predicting the demand of customers. Customers do not provide reliable long-term forecasts on the number of stores to serve or on the products needed. Besides, Pan Oston continuously tries to improve their products together with their customers. As a result, new products are introduced and products are changed or phased out. It is important that the departments within Pan Oston communicate about the life cycle of their products. As historical data of changed or new products is not available, the inventory planner can only use demand of similar products. This means both historical data and future sales data are needed to provide a good estimate on the customer demand.

Because of the unreliable demand forecast, forecasts are currently aggregated on a product group level. As an Enterprise Resource Planning (ERP) operates at product level, these forecasts are currently not integrated in the ERP-system of Pan Oston. The unreliable demand forecast results in differences between the forecasted and actual demand. The less accurate a forecast is, the more inventory is needed to handle the uncertain demand and achieve the desired service level. At the end of the life cycle, you do not want excess stock, which is the overstock without customer defined demand. Excess stock will result in obsolete stock, which is stock that has no demand at all.

Purchasing and production activities are performed after agreement of the project team. The FTO department uses an Excel file where the inventory levels of end-products are predicted based on current inventory levels and predicted customer demand. This information is manually updated, since the Excel file has no connection with the ERP-system. Based on the stock planning, production or purchasing orders are created. As a result, these orders are entered manually in the ERP-system. Both manual activities can result in data entry mistakes.

Shifted delivery date by customer

Currently, customers shift their requested delivery date frequently. A shifted delivery date can be seen as demand uncertainty, which can cause changes in the operational planning. This can lead to two things: stock-outs or products that are longer in stock. When an order is postponed, products are held in stock longer. When a customer asks for an earlier delivery date, stock-outs can occur. Since Pan Oston wants to serve their customers, they are trying to avoid stock-outs with higher safety stocks. Based on this, both problems lead to higher inventory levels. As the delivery dates are shifted by customers, this problem is not influenceable and classified as non-core problem.

Product configuration uncertainty

As described, the products of Pan Oston are highly customized. The bill of materials of the finished goods are not fixed, it is uncertain which configuration a customer wants. A customer can choose from multiple components and sub-assemblies which are purchased or produced to stock or when an order is placed. At time of the sales quotation or sales order, the configuration is known and component requirements can be determined. Because of both the internal and external lead-times, it is not always possible to postpone customer-driven activities until the customer order is arrived. As a result, the sales and project teams try to forecast this demand as accurately as possible by using historical data and information from customers. Based on this information, the purchasing department translates the demand on product group level to component requirements. But, as explained, the order configurations are not completely known before the quotations or orders, which makes it hard to determine component requirements. The process of translating demand on product group level to component requirements is mainly executed by human judgment and is not supported by data in a structured way.

Incorrect lead times in ERP

Lead times of components are entered in the ERP-system of Pan Oston. It occurs that these lead times are incorrect or outdated. Unfortunately, the current ERP-system has no feature to update the lead times at once based on actual performance or information from suppliers. The lead times of suppliers are manually entered and updated when a deviation is noticed, which is not always the case. It also occurs that agreements on the lead time are made with suppliers, but that the supplier is unable to meet this agreed lead time.

As a result of incorrect lead times, the ERP-system calculates the wrong required order date, which means that products arrive earlier or later. When products arrive earlier than requested, these products are held in stock longer. On the other hand, when products arrive later than needed, stock-outs occur. As a consequence, more inventory is held to make sure out of stocks do not occur. Unfortunately, the current safety stocks do not avoid all stock-outs.

Parameters inventory control are determined by gut feelings
The purchasing department is responsible for updating the required parameters of the inventory control policy within the ERP-system. These parameters are determined based on experience. Currently, there is no procedure to update parameters for all products at once. Parameters can only be updated manually for each SKU. Besides, no specific procedure for determining these parameters exists. When the parameters of the inventory control policy are based on experience or not updated frequently, it could lead to a non-optimal inventory control policy.

Minimum Order Quantity
Some suppliers use a Minimum Order Quantity (MOQ). This means that there is a minimum amount that Pan Oston should order. Besides the MOQ, suppliers do also give quantity discounts, which means that products become cheaper when larger quantities are purchased. When the MOQ is higher than the amount the purchaser wants to order, extra products are ordered. These products are stored in the warehouse. So, an MOQ results in higher inventory levels since products are longer in stock. As the MOQ is determined by the suppliers, it is not possible to influence this. This means that this problem is a non-core problem.

Based on the problem cluster, the core problem can be selected. An important property of the core problem is that it must be influencable (Heerkens & van Winden, 2012). Based on this, the red marked problems are eliminated. We identify that the core problems of this research are:

- Unreliable demand forecast
- Product configuration uncertainty
- Parameters inventory control are based on gut feelings

Summarizing, the unreliable demand forecast makes it hard to serve customer demand without having high inventory levels. Because of this, a lot of manual actions are performed and decisions are mainly based on experience. The manual actions are very time-consuming and prone to errors.

1.4 Research objective

Based on the problem description, the main goal of the research is defined as:

Determine how Pan Oston can forecast monthly component demand better to reduce the inventory value of components while maintaining their service level.

Due to the available time, the scope of this research should be determined. First, this research only focuses on the effect and usage of the monthly forecast in finding a better inventory control policy for components. The impact of the forecast on the production planning is not investigated. Second, only component demand from CTO and FTO is predicted. As ETO orders are engineered for customers, component demand is not predictable and thus out of scope. We also do not include the demand of the spare parts used for after-sales services. Last, the proposed solution should be applicable to the future ERP-system. The current ERP-system Microsoft Dynamics AX (hereafter referred to as AX) does not support the import of external data sources. Pan Oston has plans to implement a new ERP-system in the near future, which has the ability to import files.

1.5 Research design

In the previous section we identified the objective of this research. Based on this, the following research question is defined:

How can Pan Oston use available historical and future sales data to improve their inventory levels of components, while maintaining their current service levels?

This research question will be answered by answering the sub-questions below.

1. **How is the current situation regarding the forecasting and inventory of Pan Oston?**
   1.1. *What are the characteristics of the products produced?*
   1.2. *How is the current demand and supply process organized?*
   1.3. *Which inventory control policy is used?*
   1.4. *How does Pan Oston currently deal with supply and demand uncertainty?*
   1.5. *What is the performance of the current inventory control policy and forecasting approach?*

Chapter 2 gives an answer on the first research question and the sub-questions below. To be able to improve the situation at Pan Oston, it is required to have a clear view on the current situation at Pan Oston. The current situation includes a description of the products, planning, production, purchasing department and current performance of the inventory control and forecasting approach. The data is collected by interviews, documents and information available in the ERP system of Pan Oston.

2. **What methods to handle component demand uncertainty are available in the literature?**
   2.1. *Which methods for component demand pattern classification are available?*
   2.2. *How can the demand of components be modeled using forecasts?*
   2.3. *What are possible inventory control policies?*
   2.4. *How should the parameters of the inventory control policy be determined?*

Chapter 3 presents the answer to research question 2. Literature research is conducted to find available methods to classify components based on their demand patterns. Besides, the chapter discusses methods to model the demand of companies similar to Pan Oston and how these forecasts are used in practice. Finally, the chapter discusses how the inventory control parameters should be determined based on the forecasted demand.

3. **How can the methods found in the literature be applied at Pan Oston?**
   3.1. *What elements from the literature are suitable for the situation at Pan Oston?*
   3.2. *How can components used in finished goods be forecasted at Pan Oston?*
   3.3. *How can demand forecasting be used to determine the control parameters of the inventory control policy?*

Chapter 4 describes the method that can solve the problem of Pan Oston. This method is based on the information found in the literature, which is the answer on research question 2. First, it describes the suitable elements of the literature based on the situation at Pan Oston. Then, we describe how these elements can be used to forecast the component requirements. Finally, the chapter describes how to manage the inventory based on this forecasted demand.

4. **What is the effect and improvement of the proposed method?**
   4.1. *How is the performance of the proposed method compared with the current method?*
   4.2. *How does the proposed method handle uncertainty?*

4.3. *How can the results found be validated and verified?*

Chapter 5 describes the performance of the proposed method. The chapter starts with a description of the expected results of the new method. Next, the chapter describes how the method handles the demand and supply uncertainty. Finally, we elaborate on the validation and verification of the results found. The information of this chapter is mainly based on literature study.

5. **How can the proposed forecasting method be implemented at Pan Oston?**

   5.1. *What activities should be performed for a successful implementation of the proposed model?*

   5.2. *How do these activities have to be carried out?*

Chapter 6 explains how Pan Oston can implement the proposed method. For this, we first explain which activities are required to implement our proposed model. Besides, we also explain how these activities have to be carried out and what the possible risks are. The implementation plan is created by literature search and interviews within the company.

The deliverables of this research are:

- A user-friendly model/tool to gain insight in future component demand and support decision making regarding the inventory levels based on these predictions
- A master thesis report which includes answers on the above-mentioned questions
- Evaluation and implementation plan for integration in the current practice

2 Current Situation

This chapter describes the current situation at Pan Oston, which is the answer to research question 1. The chapter starts with Section 2.1, which briefly describes the products of Pan Oston and how these products are built from components and sub-assemblies. Then, Section 2.2 gives information on the types of inventory and the production process. Sections 2.3 and 2.4 describe how the sales and operational forecasts are created and how the production is planned based on this information. The chapter ends with a description of the current performance in Section 2.5.

2.1 Product characteristics

As described in the previous chapter, Pan Oston produces different finished products. These can be divided into three main categories, namely checkout and counters, self-checkout and self-service kiosks. These types are depicted in Figure 2.1. Figure 2.1a depicts the checkouts, which are belted checkout and express counters for (non) food, do-it-yourself (DIY) stores and Cash & Carry. A product which became very popular are the self-checkouts. An example is depicted in Figure 2.1b, which enables the customer to complete their own transaction without the need of a traditional checkout. The final product group is the self service kiosk, which is used to display information or facilitate customer actions. An example of a kiosk is displayed in Figure 2.1c.

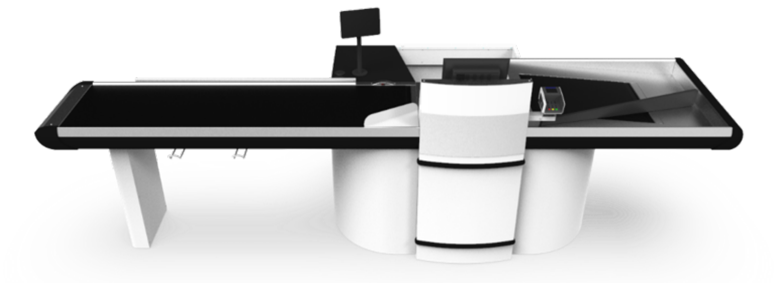

(a) Checkout & counter

(b) Self-checkout

(c) Kiosk

Figure 2.1: Product categories

As described in the introduction, the finished goods are configured based on customer wishes, which results in a lot of unique finished goods. Since the beginning of 2019, 937 different configurable products are manufactured. Based on the sales order information from 03-2018 to 03-2021, the number of sales per product category is calculated. This information is shown in Table 2.1.

Table 2.1: Number of sales per product category

Product category	Number of sales
Checkout & counter	11989
Kiosk	5065
Self-checkout	4675

The finished goods given above require a unique number of components and sub-assemblies. Some of the components or sub-assemblies are used in different types of finished goods, which is known as component commonality. On the other hand, family or customer-specific components exist, which are only used in the finished goods of a product family or for finisheds goods of a customer only. Figure 2.2 shows the generic Bill-Of-Material of a configured finished good, which describes how a finished good is built from components and sub-assemblies.

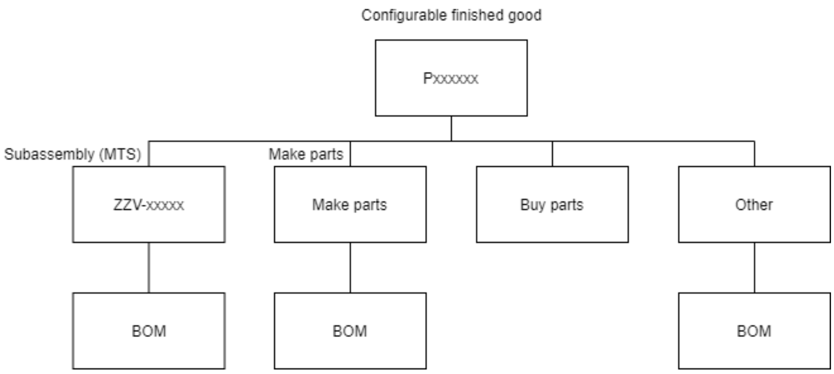

Figure 2.2: Product Structure

Sub-assemblies

A lot of configurable finished goods are built from a sub-assembly, which is an assembled part but designed to be a part of a larger assembly. Pan Oston classifies sub-assemblies as 'ZZV' in their ERP-system. The sub-assemblies are produced using a Make-To-Stock production strategy to make sure customer demand is met. Generally, the ZZV sub-assembly is finished for about 80-90%. We analysed data from the ERP-system to check in which and how many finished goods the sub-assemblies are used. From 2019, 277 ZZVs are used in the production of 260 configurable finished goods. Other finished goods are not built from a sub-assembly. Most of the sub-assemblies (65%) are only used in one finished good, from which we can conclude that the sub-assembly demand is dependent on the demand for this finished good. The other sub-assemblies are used in more finished goods, but as Figure 2.3 shows, the most sub-assemblies are only used in a few finished goods. From this, we conclude that the sub-assembly commonality is low.

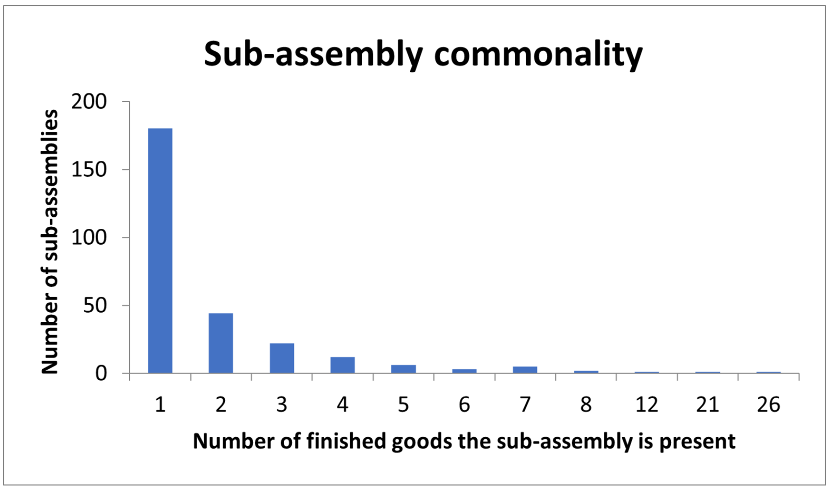

Figure 2.3: Sub-assembly commonality

Make parts

The make parts are built from other components represented in a Bill-Of-Material. This category consists of parts with direct demand (customer order) or indirect demand (part of finished good) and are produced to stock or when customer demand occurs. Since the beginning of 2019, 517 different make parts were used in finished goods. As these parts can be generic and customer specific, some parts are used more frequently in finished goods than others. We used the same ERP data to find the make parts commonality. From this, we see that in all finished goods 727 times a make part is used. 80% of the make parts are used in less than 10 different finished goods. However, the analysis also shows that some make parts are used in more finished goods, which means that the make part demand is dependent on the demand of these finished goods combined. Figure 2.4 shows the results of the analysis of the make parts in detail.

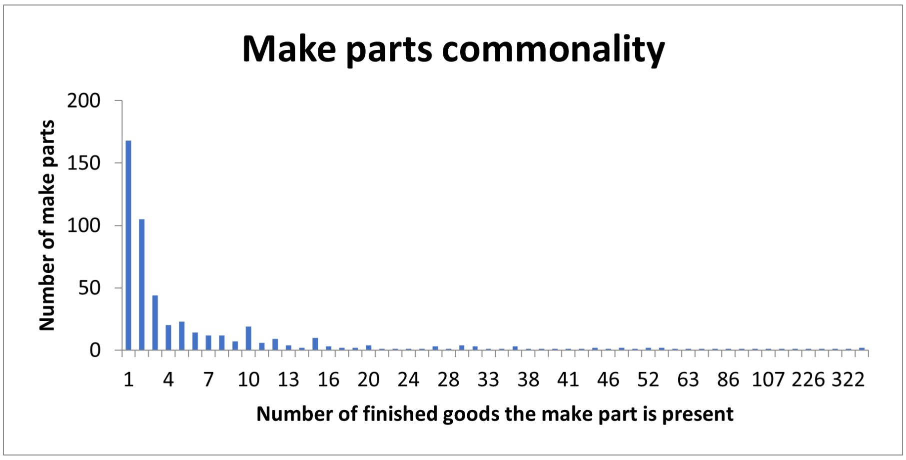

Figure 2.4: Make-parts commonality

Buy parts

The buy parts are parts purchased directly from suppliers, which is performed by the purchasing department. As the BOM shows, these parts are directly used in configurable finished goods but can also be part of a sub-assembly or make part. Pan Oston classified two types of buy parts, namely items that are purchased to stock or purchased after the receipt of a customer order. Since 2019, 3165 unique buy parts are used in the production of the finished goods. Figure 2.5 visualizes the commonality of these parts. More than half (1611 buy parts) are only used in one finished good, which means that their demand is completely dependent on the demand for this finished good. As visualized, other buy parts are used in more finished goods.

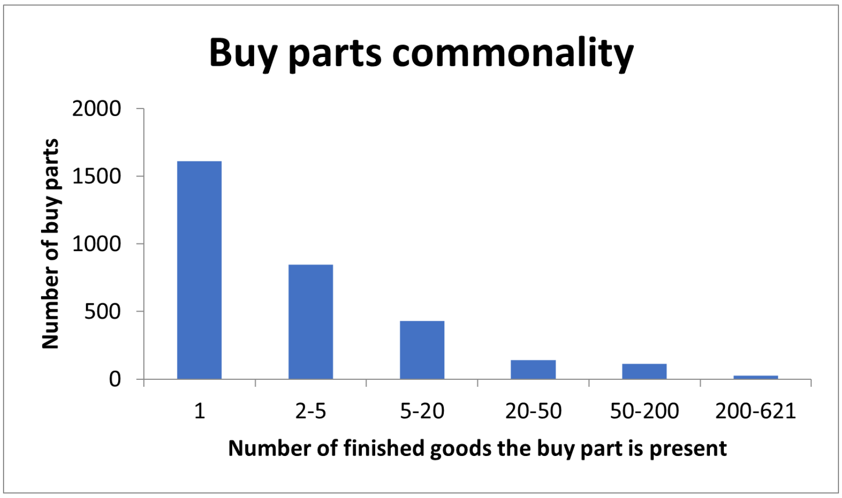

Figure 2.5: Buy parts commonality

Other

The configured products are also built from customer-specific metal parts. As these items have short internal lead times and are produced based on customer orders, there is no need to forecast these parts. Because of this, they are beyond the scope of this study.

2.2 Production and inventory

As mentioned before, the production process of Pan Oston consists of different stages. Based on the previous chapter, we can distinguish three types of inventory that are used during this process. Figure 2.6 globally shows the material flows of these inventory types at Pan Oston. The final product category, the finished goods, are delivered directly to the customers. In theory, the finished goods are not held in inventory, as they are produced after customer orders.

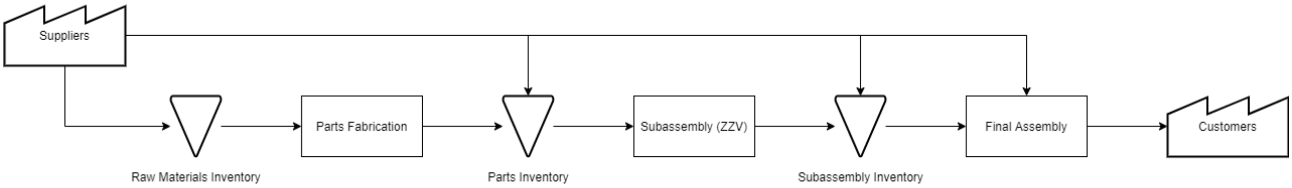

Figure 2.6: Overview Production and Inventory

The raw materials are purchased from suppliers and need to be machined before being assembled. Pan Oston has not added any value to the items in stock. The parts inventory are items in stock where value is already added. These parts are produced in-house or externally manufactured. The sub-assemblies are produced in the production facility in Raalte or production is outsourced to their partner in Slovakia. When the decision is made to outsource the production activities, the required raw materials and parts are sent from the warehouse in Raalte to the production facility in Slovakia. Before finishing the product, the final configuration is needed. This production step is performed in Raalte and Slovakia, depending on the customer location. The final configuration depends on the configuration ordered by the customer.

2.3 Forecasting

In this section, all forecasting activities are described. First, in Section 2.3.1, we describe how the sales team creates the long term forecast. After that, Section 2.3.2 describes how the project teams try to forecast finished good demand on a daily basis.

2.3.1 Long term forecast

The sales team creates a long term forecast focused on sales, which enables Pan Oston to continuously project what they plan to sell in the future. Each account manager is responsible for determining the expected sales of his customers, expressed in project value and total number of finished products. The sales data is aggregated on customer, product group and project. On a rolling basis, the forecast is updated with new information and new projects are forecasted. The forecast is monthly updated and based on confirmed sales orders (confirmation 6-8 weeks before delivery), information from customers and information from the project managers. After updating, the sum of all expected project turnovers is calculated to find the total expected turnover.

Within the planning-horizon of the long term forecast, we can distinguish three distinct time zones where different information is known. In the first time zone, all customers orders are known exactly. The second time zone has both customer orders and forecasts. In the final time zone, only sales forecasts are present. To determine the demand during the final two time zones, the sales team are heavily involved with customers to try to derive their future plans. More specific, the sales forecast can also contain unconfirmed orders. The sales representative tries to estimate the probability that the customer proceeds with the project and the probability that this project is ordered at Pan Oston. With this probability, the expected turnover of this project is calculated. Together with their past experiences and the behaviour of the customer in previous years, the sales forecast is generated.

2.3.2 Operational forecast

Both the CTO and FTO department are responsible for fulfilling customer requests. For both departments, Pan Oston uses a hybrid of Make-To-Stock (MTS) and Make-To-Order (MTO). This means that a set of components (sub-assemblies) are built to stock and configured or finished to order. The MTS policy requires that Pan Oston has an accurate forecast of future demand to be able to determine how many sub-assemblies should be produced. As the sales forecast only expresses demand in amount of euros, this does not serve the needs of the project teams. They are interested in information on the amount of individual end-products needed to be able to release production orders for the sub-assemblies and components produced to stock. However, since products are different, both departments use a different approach to determine the required stock levels of the components and sub-assemblies in time.

The FTO department creates their own stock planning for their customers, which is manually performed in Excel. The sales forecast is not used in this process, the production forecast is entirely based on information the FTO department receives from the customer representative. The stock planning has three columns for each product, showing the expected or confirmed demand, inventory levels and planned production quantities. The demand and planned production orders of products are entered to be able to calculate the expected inventory levels. Figure 2.7 shows a simplified example of a stock planning. In reality, the spreadsheet has more columns, since all products a customer can order are listed. Each row represents customer demand or production. In case of a customer order, information on the store, delivery date and expected number of finished goods required are entered. If a row is a production order, the availability date of the produced finished good is estimated.

In the spreadsheet, different row colors are used to show the status of an order. A green colored row means that a customer order is placed and that the exact content of the order is known. The remaining inventory levels after this order are exact and not estimated. A white colored row means that it is known that the customer plans to place an order, but the exact content of the order is not known. For most customers, the average number of finished goods is calculated to get an expectation of the remaining inventory levels after this order is placed. A red colored row means that Pan Oston sent a quotation including a specification of the quantities.

					Demand				Inventory				Production			
					P804157	P804162	P804163	P804164R	ZZV-80095	ZZV-80360	ZZV-80361	ZZV-80362R	ZZV-80095	ZZV-80360	ZZV-80361	ZZV-80362R
Store number	Store location	Order	Delivery date	Week	Product 1	Product 2	Product 3	Product 4	Product 1	Product 2	Product 3	Product 4	Product 1	Product 2	Product 3	Product 4
Minimum inventory level									10	5	0	0				
Inventory level end of 2020									10	9	3	3				
1	Customer order	OC-20-169641	05-01-21	wk 1	2				8	9	3	3				
2	Customer order	OC-20-170243	08-01-21	wk 1		3			8	6	3	3				
3	Customer order	OC-21-170621	12-01-21	wk 2	2				6	6	3	3				
TZB	Production order		12-01-21	wk 2					18	6	3	3	12			
4	Forecast		13-01-21	wk 2	1,8	0,0	0,1	0,0	16	6	3	3				
5	Customer order	OC-20-170066	14-01-20	wk 2	1				15	6	3	3				
6	Customer order	OC-20-168512	15-01-21	wk 2	4				11	6	3	3				
7	Forecast		17-01-21	wk 3	1,8	0,0	0,1	0,0	9	6	3	3				
8	Customer order	OC-20-168263	18-01-21	wk 3	2				7	6	3	3				
TZB	Production order		26-01-21	wk 4					19	6	3	3	12			
9	Quotation	Q-21-007045	27-01-21	wk 4	2				17	6	3	3				
10	Forecast		28-01-21	wk 4	1,8	0,0	0,1	0,0	16	6	3	3				

Figure 2.7: Example stock planning FTO

CTO customers do not give information on the predicted demand. Because of this, the CTO department mostly acts on actual customer demand (order confirmed by customer) or open quotations. No other forecast is made to predict customer demand. As the CTO customers order configurable products, the final configuration is known when a customer places an order.

Both departments are facing problems to handle the demand uncertainty. Customers do not provide the same information, since some FTO customers are unable to provide accurate forecasts. Besides, the information that the CTO department receives is not reliable and because of this, no stock planning is made.

2.4 Material planning

In Section 2.3, the sales and operational forecast methods are explained. This section describes how the operational forecast is processed into a production planning. This planning represents when and what specific configurations Pan Oston plans to produce to make sure demand is met. Besides, this section also describes how the dependent demand is generated and fulfilled by the production and purchasing departments by translating the long term forecast, operational forecast and the use of an MRP.

2.4.1 Production planning

Based on the input of the FTO and CTO department, production decisions are made. As described, the FTO department uses the stock planning to project the individual product demand in time. As the inventory levels are also entered in the Excel sheet, the sheet checks when the reorder point is reached by highlighting the cell in pink. During a meeting organized every two weeks, the involved departments check which products need replenishment. Then, the production order for these products is entered in the ERP-system AX. As there is no forecasting approach for the CTO customers, production decisions are purely based on customer orders or open quotations. Within this process, an automated report with the current inventory levels is used to project future inventory levels.

The generated master schedule, which only describes the production of the independent demand, is converted to a detailed schedule. This process is performed by the Material Requirements Planning (MRP). For each finished good on a sales order, a Bill-Of-Material (BOM) is generated in AX. This BOM contains information about the parts, components and sub-assemblies needed to produce the final product. Based on the BOM of the finished good, the dependent demand is determined. Because of this, the MRP is able to generate a schedule which describes when both production departments (metal and assembly) should start producing. Based on the production plan, material requirements that are needed for the planned production jobs can be derived. Available data such as the MOQ, lead times, processing times and available inventory are used to generate these requirements in time. Here, backward scheduling is used to determine the required start date of all activities.

As changes in the demand of final products occur, the master plan is updated daily. After updating, all calculations are performed again, which means the requirements in time are also changed.

2.4.2 Purchasing

Supply Planning

The purchasing department tries to translate the long term forecast of sales on product group level into a supply plan. In close contact with sales, they try to determine how the demand for finished goods translates into component requirements. After this, the expected demand for components is communicated with the suppliers. Arrangements with the suppliers are made to make sure Pan Oston has all components available when requested. Based on the conversations with suppliers, decisions are made regarding the inventory control policy of the components. Depending on the flexibility of suppliers, the decision is made to use an order up-to level to make sure no stock-outs occur.

The translation from product group level forecasts into component requirements is manually performed by the purchasing department based on human judgments. The process is not supported by data. The inventory control parameters are not updated on a regular basis, they are only updated manually after judgments of the purchasing department and the sales team.

Material Requirements Planning

Based on the MRP, AX automatically generates drafts of purchasing orders when required components are not available from stock. Besides, it also checks if the components with a fixed reorder point reached their reorder point. On a daily basis, the purchasing department checks if one or more purchase orders drafts are in the system. If this is the case, the draft is checked and sent to the supplier.

In the draft, the quantity to order is already filled out based on the lot size in AX. This quantity is determined by both experience and available historical data. However, there is no standardized procedure and the quantity is not updated automatically. Some suppliers work with an MOQ, which is the minimum order size a supplier is willing to accept. If a product has an MOQ, at least the MOQ is ordered. For stock keeping units, order quantities are set as a parameter and these quantities are used. However, the purchasing department uses a so-called coverage group. Each product is assigned to a coverage group with a predefined coverage time, which is the time period where demand is bundled to generate the purchase order quantity. This means that the purchasing order covers the sum of the required demand during the coverage period. This technique is used to avoid ordering costs. Besides, Pan Oston uses safety stock, which acts as a buffer stock to avoid stock-outs. This safety stock should be enough to absorb the demand and lead time uncertainty. The safety stock determination is based on experience and historical data.

2.5 Current performance

In this section, the performance of the current strategy of Pan Oston is measured. In Section 2.5.1 we provide the current inventory levels. In Section 2.5.2 we describe the stock-outs of the past three years and find out the major causes.

2.5.1 Current inventory levels

Figure 2.8 shows the inventory value based on the inventory classification of Section 2.2. Here, we classified raw materials and parts inventory, which is colored blue in this figure. As the figure shows, the raw materials and parts are responsible for about 30% of the total inventory value. Besides, we see that the total inventory value increased over the last years, which is mainly because of the increase in inventory of sub-assemblies and final assemblies in 2020 (colored in red). This increase can be explained by the COVID-19 pandemic, since production orders were completed as planned but a lot of customer orders were postponed. In this period, the inventory value of the raw materials increased only slightly.

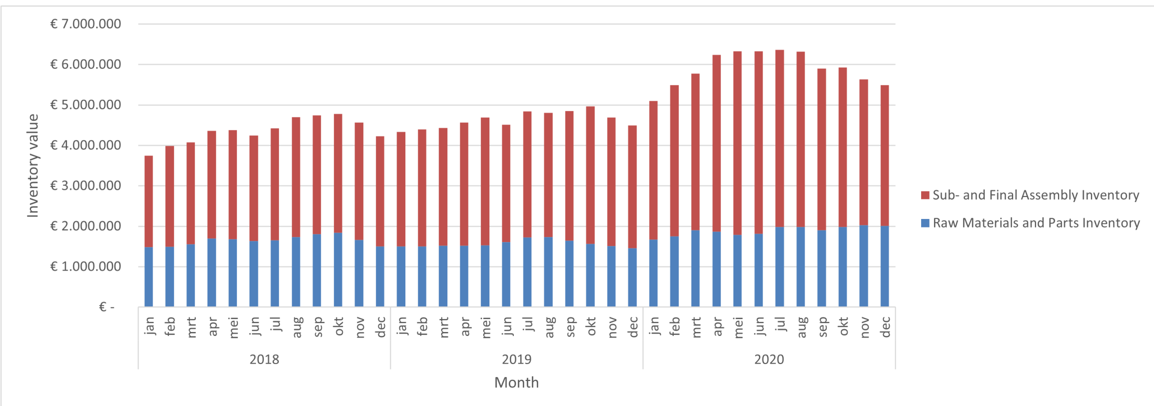

Figure 2.8: Inventory value

2.5.2 Current stock-outs

An appropriate measure for the performance of the current inventory control policy is the number of stock-outs. In 2017, Pan Oston started tracking stock-outs for the internal and outsourced assembly. A stock-out including the cause of the stock-out is recorded. At Pan Oston, a stock-out is defined as an event when an item that is on production order is not in stock when picked for assembly. From 2017 until now, 3354 stock-outs occurred. However, sometimes components are picked too early, which can also result in a stock-out. Figure 2.9 shows all causes of the stock-outs. We explain these in the remainder of this section.

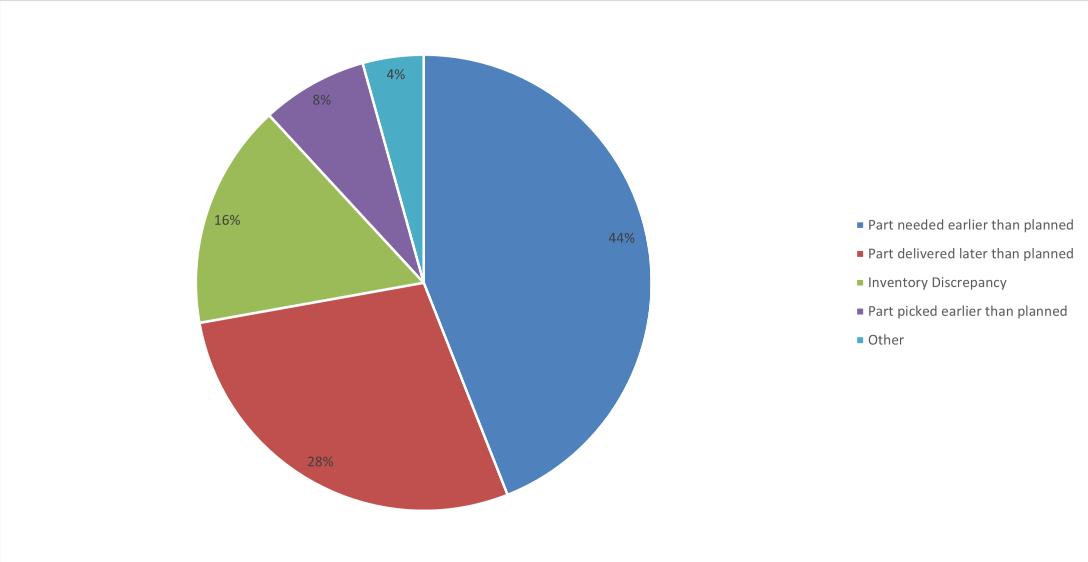

Figure 2.9: Picking stock-outs

Part needed earlier than planned

The main cause of stock-outs is that parts are needed earlier than planned. This means that these parts should be picked to start production, but are missing because the requested delivery date is later than needed. This can have multiple reasons. The first reason could be that the part requirement arrived too late at the purchasing or production department. The purchasing department is responsible for the buy parts. As products contain components that are also produced at Pan Oston, it also occurs that this department is not informed about the part demand. Besides this, it also occurs that the part demand is known, but due to an incorrect planning from the purchasing or production department the part arrived too late.

Part delivered later than planned

As discussed in the previous section, the component lead times are used to generate the requirements in time. However, the time needed to receive a component from a supplier may differ from expected, which may lead to a component that is delivered later than planned. This deviation may be caused by two reasons: a supplier does not meet the agreed lead time or the lead time in AX differs from the actual lead time.

To determine the reliability of the component deliveries, the number of working days a supply is late is checked of 22302 supplies. These supplies have information on the confirmed delivery date (delivery date confirmed by supplier) and the actual delivery date (goods receipt). In the analysis a component is considered as too late when the delivery date in the ERP system is at least two days late, since the goods receipts are sometimes posted one day after actual receipt of the goods. Besides, the purchasing department reserved two days to perform the logistics tasks needed before production. This extra time can be seen as a buffer, since a later delivery does not directly mean that parts are missing in the production process. In 3.76% of the cases, this buffer could be used. In 2.15% of the deliveries production may be delayed when no inventory is held. However, this does not directly mean that the production is delayed. The purchasing department tries to order products earlier than AX suggest to make sure that the products arrive on time. In this process, also other possible benefits are taking into account, for example lower delivery costs by bundling products of the same supplier. There is no strict procedure in this: on a daily basis the planned purchasing orders are checked and confirmed based on experience.

Table 2.2: Working days too late

Working days too late	Number of deliveries	Percentage	Cumulative percentage
Between 0 and 1	20984	94.09%	94.09%
Between 1 and 3	839	3.76%	97.85%
More than 3	479	2.15%	100%

However, Section 1.3 also describes that the lead time are often incorrect or outdated since these are manually updated by the purchasing department. Based on this description, we can conclude that incorrect component lead times result in purchasing orders that are placed too early or too late. Because of this, components do not arrive on the planned date.

Inventory discrepancy

Some stock-outs are caused by differences between the physical stock and the stock registered in AX. These discrepancies are caused by multiple reasons, for example by incorrectly or misplaced inventory. These causes are not investigated further, since they are not within the scope of this research.

Part picked earlier than planned

Sometimes parts are picked earlier than planned, which is the reason of 8% of all stocks-outs. This happens when the production department decided to work ahead of schedule. The responsible department performed all required actions to make sure the part is in stock when planned. However, when the plan is changed and parts are picked earlier, it occurs that the part is still not in stock yet.

Other causes

4% of all stock-outs are caused by reasons not directly related to the flow of goods. The most important causes are parts rejected because of quality issues, incorrect technical specifications or an incorrect Bill-Of-Material. As these causes sum up to 4% of all stock-outs, we do not investigate these reasons any further.

2.6 Conclusion

This chapter provides insight in the current situation regarding the supply and demand process at Pan Oston by answering the question: *How is the current situation regarding the demand forecasting and inventory of Pan Oston?*

Pan Oston produces a lot of different products, which are configurable from many different components and sub-assemblies. The sub-assemblies (classified as 'ZZV') are produced to stock to be able to guarantee short lead times. Besides, components can be divided into buy and make parts. Our analysis shows that differences exist in the commonality of the components and sub-assemblies. Some components are used in a lot of finished goods, while other components are product-specific.

The sales team is responsible for forecasting the future sales on an aggregated level. This forecast is not used in operational planning, since they are interested in specific finished good demand to determine when and what to produce to stock. For FTO customers, a stock planning is made to project the inventory levels of sub-assemblies. Since the CTO department does not receive reliable forecasts from customers, this department completely acts on sales quotations or confirmed orders. Based on the input of both departments, supply is planned by the purchasing department on a component level. For example decisions are made to set base-stock levels for components or to only purchase on order. So, dependent demand is derived with the use of an MRP or fixed reorder points are used.

The performance analysis showed that the inventory levels increased over the last years. From the categorization of the stock-outs we can conclude that most are caused by parts needed earlier than planned. Literature research will be conducted to find methods to forecast these components requirements better.

3 Literature Review

This chapter elaborates on the questions related to the literature, which are all part of research question 2. Section 3.1 discusses the relevant Supply Chain Management concepts in relation to the problem context. Then, the chapter continues with Section 3.2, where we describe different techniques to generate component requirements from the demand plan. In Section 3.3 we describe how we can identify different demand patterns, which statistical forecasting methods are available in literature and how the performance of forecasts can be measured. Finally, Section 3.4 summarizes the findings of the literature review.

3.1 Supply chain management

3.1.1 Customer Order Decoupling Point

The customer order decoupling point (CODP) is a term introduced to identify the point where the flow changes from a push to pull strategy. The chosen strategy determines how processes are organized. The most frequently used strategies are Make-To-Stock (MTS) and Make-To-Order (MTO). In an MTS environment, products are produced based on forecasts in stock. Customers orders are delivered from inventory. In an MTO manufacturing strategy, production begins after a confirmed customer order is received. In literature, strategies as Assemble-To-Order (ATO), Configure-To-Order (CTO) and Finish-To-Order (FTO) are described as hybrid between the MTO and MTS strategy. In these strategies, components are produced to stock, but final configuration is performed after a customer order is confirmed.

In an ATO strategy, a standard product is finished from components held in stock. In literature, the difference with the CTO strategy is not described consistently. However, the most important distinction is found in the different use of standard products. In the ATO strategy the standard product is delivered to the customer without an option for customization. In a CTO strategy a limited number of standard products exist, but customers are allowed to configure the standard product to make sure their needs are served. As shown in Figure 3.1, manufacturing the parts and sub-assemblies is forecast-driven, with the final assembly being order-driven. The CODP is not positioned before the assembly phase, since this strategy assumes that certain assembly steps can also be performed without customer order. The FTO strategy represents the latest CODP, since only a nearly finished product is processed. This strategy does not change the functionalities or product structure, but only changes in the appearance of the product are made.

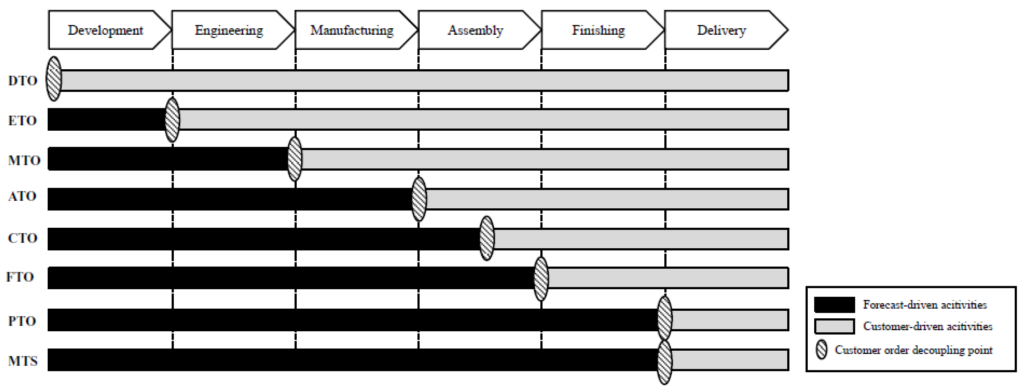

Figure 3.1: Customer Order Decoupling Point (Bogner et al., 2017)

3.1.2 Demand planning

As Figure 3.1 shows, in CTO and FTO a lot of activities are forecast-driven. Components are purchased and kept in inventory. At Pan Oston, also sub-assemblies are produced to make sure finished goods can be delivered within the promised lead time. Demand planning refers to the process of the sales team where demand is forecasted for each period in the upcoming planning horizon. During this process, marketing plans, product introductions and obsolescence should be taken into account. Silver et al. (2016) developed a forecasting framework where forecasting is based on historical data and informed judgement. This framework is depicted in Figure 3.2. Other literature typically distinguishes three types of forecasting methods: time-series analysis, causal models and human judgment. The first two methods are quantitative methods based on mathematical models. We will elaborate on the quantitative methods in section 3.3.

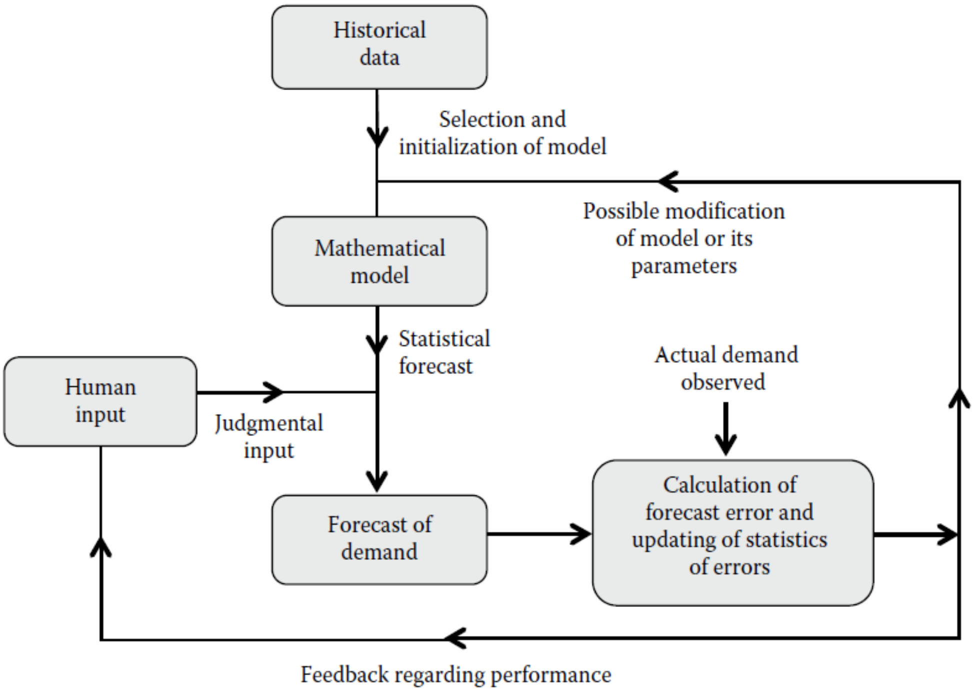

Figure 3.2: Forecasting Framework (Silver et al., 2016)

In some cases, it is also useful to incorporate human judgement in the forecasts besides the quantitative forecasting methods. In these cases, the quantitative methods do not capture the business context. As a result, human judgement is crucial to improve the forecasts. For example when a new product is introduced, no historical data is available. With human judgement, the demand of this new product can be derived when this introduction substitutes other products. Fildes, Goodwin, Lawrence, and Nikolopoulos (2009) concluded that adjusting forecasts might improve the forecasting accuracy. However, they also found out that these judgments can be biased and inefficient.

Literature discusses multiple principles to avoid possible limitations of incorporating human judgement. First, Fildes and Goodwin (2007) describe that judgmental adjustment should only be applied when information is not available to the statistical method. When it is possible to incorporate the information by adding factors to the quantitative forecast, this should always be preferred. Besides, they describe that reasons and assumptions of adjustments should always be documented and that appropriate error measures should be used to assess the performance frequently. Silver et al. (2016) add to this that a small committee, consisting of different department representatives, should be responsible for the judgmental input. The generated statistical forecast should always be the starting point, aggregated across different groups of items. Then, adjustments should be made individually before talking as a group. The committee can take the average and adjust the aggregated demand, which results in an adjustment for item level demand. They also describe that it is crucial to provide feedback on the performance of the judgmental input by comparing the actual demand with both the the original statistical forecast and the revised forecast.

3.1.3 Material Requirements Planning

A crucial part of production systems is inventory control. When a company uses an inefficient inventory control policy, it can lead to shortages or excess stock. Hautaniemi and Pirttilä (1999) created a classification method to be able to select an appropriate inventory control policy for each product in an MRP environment. MRP is a frequently used method to find inventory requirements for dependent demand. In deterministic environments, the MRP logic gives an optimal just-in-time schedule. However, for a stochastic environment, proper parameters are needed to deal with uncertainties (Dolgui & Prodhon, 2007). The goal of the MRP is to find a requirements schedule for each time period in a predefined time-horizon. This schedule is based on the final assembly plan (MPS), the Bill of Material (BOM), current inventory levels and lead times. In the production environment of Pan Oston, the final assembly plan is only known a few weeks in advance. This causes problems in the MRP system, since long lead time components are ordered too late. An accurate planning method based on forecasts is required to deal with this problem. Additionally, the BOM structure has to be up-to-date to make sure component requirements can be planned successfully. Most MRP-systems assume a fixed BOM, which is not the case for CTO systems.

A drawback of an MRP system is what in literature is defined as nervousness. Steele (1975) defined a nervous MRP system as a system that causes major rescheduling actions to lower level requirements when minor changes to the master schedule are made. There are different causes for changes in the MPS, which are mainly related to changes in quantity or timing of planned orders or scheduled receipts (Jacobs, Berry, Whybark, & Vollmann, 2011). Ho and Ireland (1998) investigated the occurrence of system nervousness at lower levels of the production system as a result of forecast error. They concluded that the existence of forecast errors resulted in scheduling instability in MRP systems. In literature, a lot of solutions are proposed to reduce nervousness of MRP systems. First, Jacobs et al. (2011) described that causes of changes to the MRP should be reduced. This can be done by using freezing or time fences. Freezing the MPS for a specified period resulted in a less nervous MRP, but it can also result in more stock-outs for MPS planned items because of a low responsiveness to demand changes (Zhao & Lee, 1993). Other ways to reduce the nervousness are incorporating spare parts forecasts or controlling the introduction of parameter changes. When nervousness still exists, Jacobs et al. (2011) described that different lot-sizing procedures for different product structure levels can contribute towards reducing the nervousness. Finally, the use of firm planned orders in the MRP or MPS can help stabilizing the lower level requirements, since capacity and material is reserved at an early stage.

3.2 Inventory control and supply planning in CTO environments

ATO and CTO environments are extensively described due to their complex product structure. Demands occur only for the products, but inventory is kept for components or sub-assemblies only. In ATO and CTO systems, demand can only be fulfilled if components are available simultaneously. Besides, a single component may be common to multiple products. As described, a CTO system is a type of an ATO system where a customer can choose from a set of options to order their finished good. This introduces an extra uncertainty in component demand, as demand on product group level does not require the same components each time. In literature, there are multiple examples available where companies try to cope with the difficulties of those systems. Most of the literature assume independent base-stock policies to make the analysis easier. We can divide the literature into two categories, since the first deals with periodic review systems and the other with continuous review systems.

Cheng, Ettl, Lin, and Yao (2002) developed an analytical model for a configure-to-order operation with periodic review and deterministic demand for product families. They developed an exact algorithm for the case where there is one component for each product family not used in any other product family demand. As this is an unrealistic assumption for many companies, they also developed a greedy heuristic for the non-unique component case. Agrawal and Cohen (2001) developed a multi-item inventory control problem with component commonality and correlated finished product demands. They minimized the total expected component inventory costs under order fill rate constraints with stationary demand and deterministic lead times. The results are used to determine base stock policies to manage component inventory levels. They found out that the contribution of a component to the finished product service level could vary by component. For example components with a high degree of commonality should have high fill rates.

Song (1998) considered finished product orders service levels in a continuous time frame with Poisson arrival rate of orders. FCFS policy for assigning components to orders. In their model, demand of a certain type requires a fixed set of components, but in a random amount. They calculated the order fill rate by a series of convolutions, which is a mathematical operation to calculate the probability density function of the sum of two independent random variables.

Srinivasan, Jayaraman, Rappold, Roundy, and Tayur (1998) investigated the procurement of common components in a multi-period stochastic environment. The components to procure are common to several products and involve no uncertainty, which means that order configurations are fixed. They analysed the case where the uncertainty came from demand on product level. Their objective was to determine the component quantities that satisfy a pre-specified service level for products. Chen-Ritzo, Ervolina, Harrison, and Gupta (2010) investigated the S&OP problem in CTO systems with order configuration uncertainty. They propose a sample average approximation and test their approach on data obtained from IBM. They assume a deterministic demand planning, which means that the expected sales are met exactly. This is exactly the opposite problem as introduced by Srinivasan et al. (1998).

van den Bogaert (2017) presented a general algorithm to optimize replenishment policies in CTO systems. Instead of focusing on customer demand streams, the capacity of the factory is used to find component requirements. Besides, he assumed a FCFS allocation method and stationary demand. As product level fill rates in ATO are computationally intractable, he used an approximation assuming that component inventory levels behave independently. The problem is solved using a Langrangian Relaxation, where both base-stock policies and (s,S) policies are explored.

3.3 Statistical forecasts

In this section, we first focus on demand characteristics. Then, we will discuss different statistical forecasting techniques.

3.3.1 Demand characteristics

Demand can have different patterns, which require different forecasting techniques. A lot of models assume demand to be normally distributed. However, this assumption is not always the case, which can result in high forecast errors. Syntetos, Boylan, and Croston (2005) created a categorization framework which distinguished four non-normal demand patterns: erratic, smooth, lumpy and intermittent. They operationalized these categories by using two parameters: the Average inter-Demand Interval (ADI) and the Squared Coefficient of Variation (CV^2) .

The ADI is the average time between two demand occurrences and can be calculated by:

$$ADI = \frac{\sum_{n=1}^{N} t_n}{N} \tag{3.1}$$

where t_n is the time between two consecutive demand periods and N is the number of periods with non-zero demand.

The CV^2 is the standard deviation of the demand divided by the average demand. The formula is as follows:

$$CV^2 = (\frac{\sqrt{\frac{1}{N}\sum_{n=1}^{N}(d_n - d)^2}}{d})^2 \tag{3.2}$$

where d_n is the demand of period n and the average demand d can be calculated by $d = \frac{\sum_{n=1}^{N} d_n}{N}$.

Based on these parameters, the demand patterns can be categorized:

- Smooth demand: ADI $<$ 1.32, CV^2 $<$ 0.49. Characterized by a frequent and stable demand pattern.
- Erratic demand: ADI $<$ 1.32, CV^2 $>$ 0.49. Characterized by a frequent but irregular demand pattern.
- Intermittent demand: ADI $>$ 1.32, CV^2 $<$ 0.49. Characterized by a sporadic, but stable demand pattern.
- Lumpy demand: ADI $>$ 1.32, CV^2 $>$ 0.49. Characterized by a sporadic and irregular demand pattern.

Table 3.1: Overview categorization demand patterns

CV^2 $>$0.49	Erratic	Lumpy
CV^2 $<$0.49	Smooth	Intermittent
	ADI $<$1.32	ADI $>$1.32

The classification of Table 3.1 can help in determining the difficulty of forecasting the demand of the products. For products with sporadic demand you have to determine when the next demand will occur. When demand is irregular, you have to determine what the volume of the demand is. Forecasting smooth demand is the easiest, since demand occurs frequently and stable. For lumpy demand, you have to determine when the next demand will occur and what the volume of this demand is.

3.3.2 Time-series analysis

In time-series analysis, historical data of a variable is used to forecast the future of the same variable. Time Series Analysis look for time patterns such as a trend, seasonality and cycles in historical data. In contrast to causal models, which are explained in subsection 3.3.3, time-series analysis do not take external influences into account. The most frequently used methods are explained in this section.

Simple and Weighted Moving Average

Moving Average is a simple, but frequently used method where the average of a specified number of observations is used to forecast the future. The chosen number of observations depends on the data to be forecasted.

The simplest method is called Simple Moving Average and assumes that all observations are equally important. This method is useful for products with stable demand. In the formula below, we define $\hat{x}_{t,t+\tau}$ as the demand of period $t+\tau$ forecasted after the observation of period t arrived. This forecast is based on the observed demand D_{t-i} of the chosen number of time periods N. The corresponding formula of Simple Moving Average is:

$$\hat{x}_{t,t+\tau} = \frac{1}{N}\sum_{i=0}^{N} D_{t-i} \tag{3.3}$$

Another method is the Weighted Moving Average. This method assumes that the latest observations are most important. In contrast to the simple moving average, where all weights are equal, weighted moving average uses different weights for each observation. Because of this, this method is more suitable for products with some seasonality. The formula of Weighted Moving Average is:

$$\hat{x}_{t,t+\tau} = \frac{1}{N}\sum_{i=0}^{N} W_{t-i} * D_{t-i} \tag{3.4}$$

In this formula, W_{t-i} is the weight assigned to the demand D_{t-i}.

Simple Exponential smoothing

Another forecasting method is simple exponential smoothing. This method gradually gives declining weights to historical data. The method assumes the demand follows a level pattern, since a trend or seasonality are not included. After each new observation, the level $\hat{a}_t$ is updated by:

$$\hat{a}_t = \alpha D_{t-1} + (1-\alpha)\hat{a}_{t-1} \tag{3.5}$$

The parameter α is called the smoothing factor, which controls the rate at which the weight of the historical data declines. This value should be set between 0 and 1. A value close to 1 gives the most weight to the most recent observation. A value close to 0 means that the earlier observations are more important, which results in a more stable level.

The forecast of the demand in period t can be calculated by:

$$\hat{x}_{t,t+\tau} = \hat{a}_t \tag{3.6}$$

Holt's Exponential Smoothing

The Simple Exponential smoothing assumes the demand has no trend or seasonality. In Holt's Exponential Smoothing, also known as Linear Exponential Smoothing, the formula of simple exponential smoothing is extended with a trend variable. (Holt, 1957) After each observation, both components are updated. In addition to the smoothing parameter of the level, the smoothing of the trend can be controlled by smoothing factor β.

The level and trend can be updated by:

$$\hat{a}_t = \alpha D_{t-1} + (1-\alpha)\hat{a}_{t-1} \tag{3.7}$$

$$\hat{b}_t = \beta(\hat{a}_t - \hat{a}_{t-1}) + (1-\beta)\hat{b}_{t-1} \tag{3.8}$$

The forecast of the demand in period t is a combination of the level and the trend, which can be calculated by:

$$\hat{x}_{t,t+\tau} = \hat{a}_t + \hat{b}_t\tau \tag{3.9}$$

Holt-Winters Exponential Smoothing

Winters (1960) introduced an extension to the procedure of Holt, which also includes seasonality. Products can contain significant seasonality, which are periods in which events are recurrent. The seasonality of period t is expressed with seasonal index $\hat{F}_t$. First, the number of seasons P should be determined based on the demand pattern. For example, $P = 52$ assumes that each week shows a seasonal pattern.

The level, trend and seasonal factors are updated by the following three equations:

$$\hat{a}_t = \alpha D_{t-1} + (1-\alpha)\hat{a}_{t-1} \tag{3.10}$$

$$\hat{b}_t = \beta(\hat{a}_t - \hat{a}_{t-1}) + (1-\beta)\hat{b}_{t-1} \tag{3.11}$$

$$\hat{F}_t = \gamma(D_t/\hat{a}_t) + (1-\gamma)\hat{F}_{t-P} \tag{3.12}$$

After updating the parameters, the demand of period $t+\tau$ is estimated with:

$$\hat{x}_{t,t+\tau} = (\hat{a}_t + \hat{b}_t\tau)\hat{F}_t \tag{3.13}$$

ARIMA

Another widely used time series analysis method is ARIMA, often called the Box-Jenkins method. This method is often notated as ARIMA(p,d,q) and consists of three different components, which we explain below:

- AR: Autoregressing
- I: Differencing
- MA: Moving Average

The first AR part uses regression to forecast the variable based on past values of the variable, which is called auto regression. The parameter p is the number of periods of our time series (lags) used. Formula 3.14 describes how the autoregressive model can be written. Here, x_{t-i} is the lag of time period $t-i$ and φ_i is the regression coefficient of predictor i estimated by the model. Parameter c is the constant term indicating the intercept value.

$$\hat{x}_t = c + \sum_{i=1}^{p} \varphi_i x_{t-i} + \varepsilon_t \tag{3.14}$$

The MA part uses past forecast errors of the autoregressive model to perform linear regression. The model is shown in the formula below.

$$\hat{x}_t = c + \varepsilon_t + \sum_{i=1}^{q} \theta_i \varepsilon_{t=i} \tag{3.15}$$

Both the AR and MA method require stationary demand. As this is often not the case, an integrating step indicated by I is required. The parameter d is the number of times that the data values are replaced by the difference between the data values and the previous data values. The first order differencing step is shown in Formula 3.16. Depending on the data, it may be necessary to apply more differencing steps to create stationary demand.

$$x_t^{'} = x_t - x_{t-1} \tag{3.16}$$

The model described above can be extended to include seasonal patterns or multivariate time series. Seasonality can be included using SARIMA. In this method, $(P, D, Q)_m$ describe the seasonality of period m. When it is preferred to predict an outcome based on multiple time-dependent variables, the extension VARIMA (Vector Autoregressive Moving Average Model) can be useful.

Croston's method

Croston (1972) found out that the traditional time series analyses have trouble with demand patterns with irregular demand and developed an own forecasting method that performed better under these conditions. In this method, the forecast is determined by using the demand size and demand interval. Both are forecasted with separate exponential smoothing estimates.

The parameters of Croston's method are only updated when the most recent demand x_t is not equal to zero. The demand size $\hat{z}_t$ and demand interval $\hat{n}_t$ are updated with the use of smoothing factor α. n_t represents the inter-arrival time the previous non-zero demand and the current demand at period t. If demand occured ($x_t > 0$), the formulas below are used to update the demand size and demand interval.

$$\hat{z}_t = \alpha x_t + (1 - \alpha)\hat{z}_{t-1} \tag{3.17}$$

$$\hat{n}_t = \alpha n_t + (1 - \alpha)\hat{n}_{t-1} \tag{3.18}$$

After updating both parameters, the forecast of period $t + \tau$ can be calculated by:

$$\hat{x}_{t,t+\tau} = \frac{\hat{z}_t}{\hat{n}_t} \tag{3.19}$$

3.3.3 Causal models

Causal models are very useful when the forecast of demand for an item is based on the demand of another item or variable (Axsäter, 2006). An example is when demand of an item depends on the demand of a final product. Causal models assume that the variable that is forecasted has a causal relation with other independent variables.

Simple and multiple regression
Literature typically distinguishes between two types of regression, namely simple and multiple regression. In the simplest case, the regression model assumes a linear relationship between the two variables: the outcome variable $\hat{x}_t$ and the predictor variable x. Formula 3.21 describes how the forecasted variable can be computed using simple linear regression, where $\hat{a}$ represents the intercept and $\hat{b}$ is the slope. Both can be estimated using the least squares criterion.

$$\hat{x}_t = \hat{a} + \hat{b}x \tag{3.20}$$

With multiple regression, two or more independent variables can be used to predict the dependent variable. The parameters $\hat{a}$ and $\hat{b}_i$ can be estimated using the same formula. In the formula below, we assume that the dependent variable has a causal relation with k independent variables.

$$\hat{x}_t = \hat{a} + \sum_{i=1}^{k} \hat{b}_i x_i \tag{3.21}$$

3.3.4 Measuring forecasting performance

It is important to measure the forecasting performance because of multiple reasons. First, the standard deviation of forecast error in lead time demand should be estimated to determine required safety stock levels. Second, it is important to monitor the validity of the forecasting model and the chosen parameters. Finally, forecast errors can be used to get feedback on the impact of subjective input to the forecasts.

A common approach in forecasting is dividing the available historical data into a training and a test set. The training set is used to build the model and forecast the test set data. Then, the test set is used to evaluate the forecast accuracy by using different indicators, each having their advantages and disadvantages.

One measure of the accuracy is the mean square error (MSE). The main disadvantage of the MSE is that outliers are heavily weighted, as each error value is squared. The MSE can be computed by:

$$MSE = \frac{1}{n}\sum_{t=1}^{n}(D_t - \hat{x}_t)^2 \tag{3.22}$$

Another measure is the mean absolute deviation (MAD). This measure is more robust than the MSE, since it is not sensitive to outliers. The estimate of the MAD can be calculated by:

$$MAD = \frac{1}{n} \sum_{t=1}^{n} \mid D_t - \hat{x}_t \mid \tag{3.23}$$

The third measure is the bias. This measure is useful to identify over- and under-forecasting. The bias is positive if the forecast is greater than actual demand. When the bias is negative, the forecast is lower than the actual demand. The formula of the bias is:

$$Bias = \frac{1}{n} \sum_{t=1}^{n} (\hat{x_t} - D_t) \tag{3.24}$$

The fourth well-known measure is the MAPE, which is the mean absolute percentage error. The advantage of this measure is that it is unit-free, which means forecast performances of different data sets can be compared. The biggest disadvantage is that this measure is not useful when there are periods with low or zero demand, as this results in an high MAPE or undefined outcome. The MAPE can be calculated by:

$$MAPE = \frac{1}{n} \sum_{t=1}^{n} \mid \frac{D_t - \hat{x}_t}{D_t} \mid \tag{3.25}$$

The final measure is the WAPE. In the situation were the number of sales is low or the product analysed has intermittent sales, WAPE can give better results than MAPE. The main reason for this is that we divide our forecast error by the demand during all periods. The formula of the WAPE is:

$$WAPE = \frac{\sum_{t=1}^{n} \mid D_t - \hat{x}_t \mid}{\sum_{t=1}^{n} D_t} \tag{3.26}$$

3.4 Conclusion

This chapter presents the literature review and gives an answer to sub question 2: *What methods to handle component demand uncertainty are available in the literature?*.

Section 3.1 introduced production environments and their CODP. In ATO, CTO and FTO policies, components and sub-assemblies are produced to stock, but the final step is performed after arrival of a customer order. This means that these systems require forecast-driven activities. During the demand planning phase, the sales team forecasts demand for each period in the upcoming planning horizon. Statistical forecasts can help during this process, which is also proposed by Silver et al. (2016).

With the classification of Syntetos et al. (2005) four demand types can be identified, which require different forecasting techniques. This classification uses the Average inter-Demand Interval (ADI) and Squared Coefficient of Variation (CV^2). The four demand types are smooth, erratic, intermittent and lumpy.

Quantitative forecasts can be divided into two categories, namely time-series analysis and causal models. In time-series analysis, the demand is forecasted based on earlier observations of this demand. We found out that there are different types of time-series analysis, each one having its own advantages and disadvantages. In causal models also other independent variables can be used to forecast the demand. In addition to these techniques, human judgement might be useful to improve the forecasting accuracy. Since qualitative data can be biased and inefficient, it is important to only use this data when it is not possible to include in the statistical analysis. Besides, the adjustment should be properly documented and the performance should be measured frequently.

Based on the literature review, we conclude that different techniques can be used to derive component requirements in the supply planning phase for systems with order configuration uncertainty. Traditional time-series analysis can be used, but only use past demand and do not incorporate available sales information. Other techniques use information on product group level to determine component requirements and optimize inventory levels, but often assume stationary demand for tractability. This assumption does not hold for the environment of Pan Oston, since demand on product level often comes from small or large projects. We have to incorporate the non-stationary demand, which we discuss in the next chapter. To test the accuracy of our forecast, we can use a traditional time-series analysis as benchmark.

4 Solution Design

This section describes the creation of the model. We develop the model by selecting methods from the literature that are suitable for the supply chain characteristics and the formulated model requirements. After this, we describe the different stages of our developed method. Section 4.1 provides an overview of the model and motivates our decision to incorporate forecasts on product group level. In Section 4.2 we describe the design of our model and explain all model steps in short. In the next section, we introduce the variables we use in the remainder of the thesis and describe the steps of our model in more detail by providing the required calculations and examples of the in- and outputs. Here, we also describe how we evaluate given base-stock levels and how we propose the base-stock levels based on our optimization approach. Section 4.6 finalizes this chapter with the conclusions.

4.1 Model overview and motivation

We can distinguish two types of approaches to determine component requirements. The first consists of using time-series forecasting models based on historical consumption data. With this method, patterns are explored to predict future component demand. This strategy is limited to historical data and does not consider available sales information. The second approach uses the BOM of the finished good to derive component requirements based on the finished good forecasts. For tactical purposes, the finished good forecasts are combined on a more aggregate level. As we want to include the projected sales information, the first approach is not suitable. In the second approach, all finished goods should be forecasted. Given the numerous product configurations at Pan Oston, forecasting at the SKU level is inefficient and highly inaccurate. Instead of forecasting each finished good separately, we choose an approach where we use the sales forecasts on group level. These aggregate forecasts are more accurate and less effort is needed to forecast these groups. At the time purchasing decisions should be made, the customer demand on product group level can be known, but the customer's selection of components can still be uncertain. To deal with this, we calculate the probability how many units of a certain component are required in a finished good of the corresponding group based on historical data. With this, we can compute the monthly component demand. In the next section, we introduce the steps of our model in more detail.

4.2 Model design

4.2.1 Model steps

Figure 4.1 visualizes the model design, which consists of four steps. Our model consists of two main parts: first we generate component requirements, then we determine base stock levels for each component that satisfy the required service levels at lowest possible costs. The steps are the following:

1. Calculate product configuration probabilities from past production and sales orders
2. Create judgmental sales forecasts of product groups
3. Calculate component demand from judgmental sales forecasts
4. Determine near-optimal base stock levels using a greedy heuristic

The planning horizon for our model is one year (12 months) with monthly periods. For the entire planning horizon, we determine component requirements and optimize base stock levels. We choose this horizon because of two main reasons. First, the current sales forecast can plan a maximum of 12 months ahead since customers also provide their plans at maximum one year ahead. Second, product life cycles are relatively short and new components are introduced frequently. With a longer planning horizon, less information from customers is known and components are replaced frequently during our time horizon.

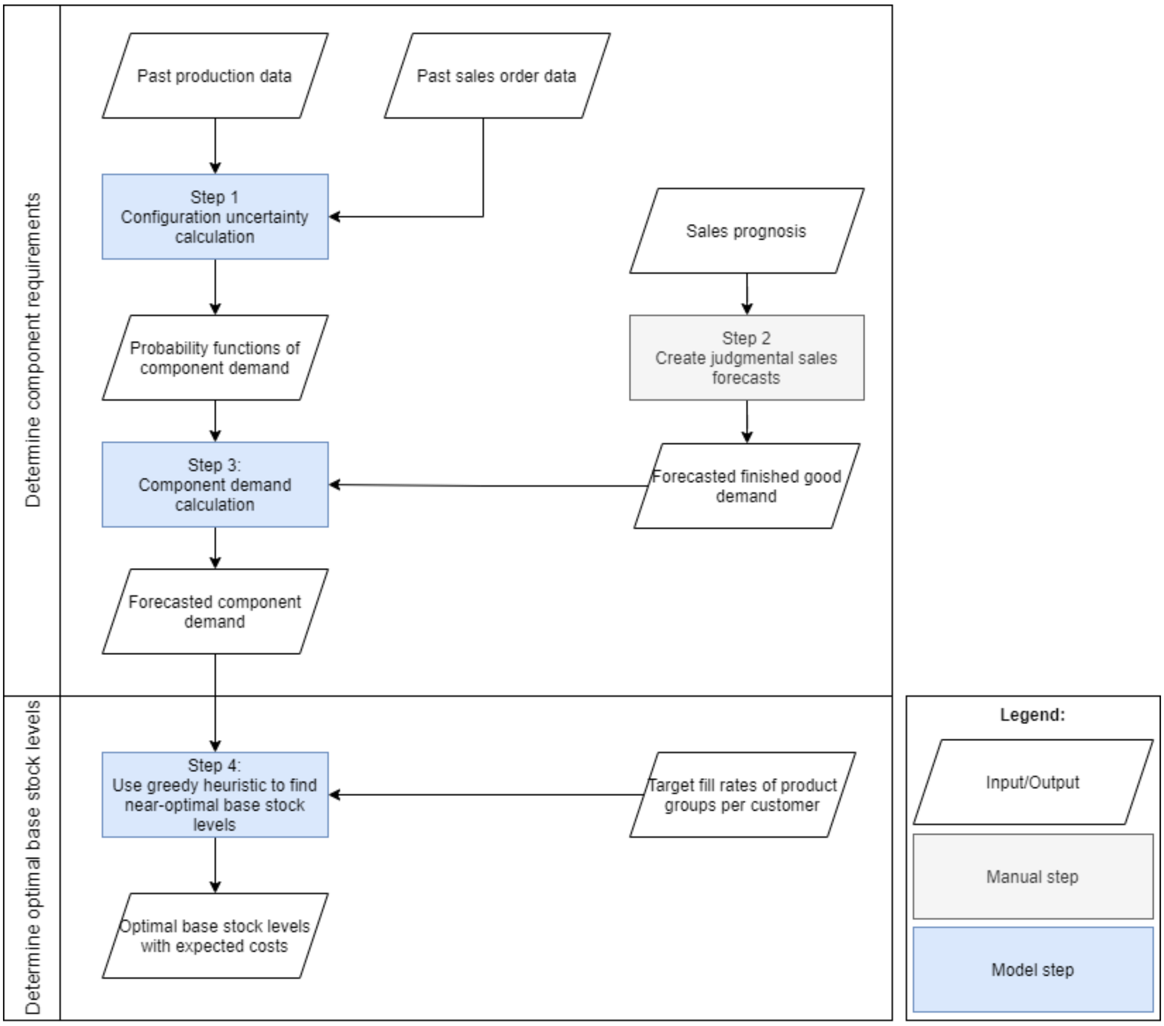

Figure 4.1: Overview of proposed model

4.2.2 Model introduction

The first step in our model is to calculate the probability that a certain amount of a component is required in a finished good of a product group. As explained in Chapter 2 the products of Pan Oston can be divided into separate categories: checkouts & counters, self-checkouts, kiosks and others. Based on this classification, each product on a sales order can be categorized into a combination of a customer and product-group. Because of the uncertainty in configuration, each requested product of this group can require a certain amount of a component. We use both past sales and production orders to find the probability of this occurrence. The finished good chosen by the customer is retrieved from the sales order. Since products are highly configurable, we can distinguish two main components: customer-specific options (chosen by a customer) and fixed elements (sub-assemblies or components that are always included). The configurable part in a sales order is linked to a specific production order. However, the fixed elements are produced anonymously and in advance, which means they are not linked to a specific sales order. So, we use the production order information of this generic part to derive the component demand of the particular order. We assume that the data from AX is up-to-date. Sales orders are sent to customers and accurately describe what is delivered. Production orders are used for various processes within the company and we already showed that stock-outs are rarely caused by inventory discrepancies.

The second step is to create judgmental sales forecasts, based on the prognoses provided by sales, project teams and customers. As explained in Chapter 2, the judgmental forecasts can currently not be accurately translated to component requirements. In our analysis of the current situation, we found out that the current sales forecast is obtained from employees heavily involved with customers to get insights in their future demand. Each month, the sales forecast is updated with new information obtained. Changes in the forecast exist, but over the chosen planning horizon, the forecast on product group level is quite accurate due to this close customer contact. Because of this, we assume deterministic demand on product group level. We propose some changes to create input that can be used in step 3, since product groups in the ERP-system AX should be in line with the forecasted groups.

The third step is to calculate the total demand for all components, using the probability functions of individual component demand and the forecasts on product group level as input data. Since all products are forecasted on a product group level, we first calculate the probability functions of component demand from a specific product group. In this calculation, we assume that the time needed to produce a finished good from the set of needed components is negligible. This is because Pan Oston is able to react quickly to customer demand when components are in stock. Longer outbound lead times are mostly caused by the low availability of components.

The final, fourth step is to determine the optimal base stock levels given a set of parameters. We use the total component demand we calculated in the previous step and other relevant parameters related to the components, such as component prices and lead times. We assume that the parameters from the ERP-system are valid and fixed, with no uncertainty involved. We decide to use a greedy heuristic where we evaluate the service level of the forecasted finished goods based on component availability. We set a target service level for each product group and try to find the minimum costs given these targets.

4.3 Model description

In the previous section we introduced our model and explained each step briefly. This section describes the four main steps of our model in more detail.

4.3.1 Notation

In the remainder of the thesis, we use the notation below.

Notation	**Description**
G	set of customer- & product groups ($g \in G$)
C	set of components used in finished goods ($c \in C$)
T	set of time-periods (months) ($t \in T$)
p_c	purchase price of component c
l_c	supply lead-time of component c
$F_g(t)$	final forecast of group g in period t
X_{cg}	number of units required of component c to produce one finished good of group g
Y_{cgt}	number of units required of component c to produce demand of group g in month t
FR_g^{obj}	target fill rate of group g
D_{ct}	total demand of component c in month t
$I_c(t)$	on-hand inventory of component c at start of period t
$s_c(t)$	base stock level of component c at time t

4.3.2 Calculate product configuration probabilities

The first step is to calculate the probability functions based on the historical sales orders and production orders. As we are only interested in the demand for components, data should be pre-processed to create the flattened

BOM for each produced finished good where all intermediate levels are removed. In this BOM, the lowest level components are directly linked to the finished good. After this first step, we exclude all non-critical components. Non-critical components, such as bolts and nuts, are often low-value and controlled by two-bin systems. Besides these components, we also exclude all components that are purchased or produced on order (such as metal parts). These components can be made available after receiving a customer order, which means there is no need to forecast future requirements. The resulting BOM consists of the most important components of a particular finished goods and is displayed in Table 4.1. With this modified BOM, we can start calculating component requirements.

Table 4.1: Example of historical sales and production order input

Sales ID	FG ID	FG quantity	Customer- & productgroup	Component	Component quantity
OC-19-130459	P812020R	1	Cust1Checkouts	700005	1

As described, customers can configure their finished goods by selecting the components they want. We model this uncertainty in product configuration by a discrete empirical distribution. When the sales forecast on a product group level is expressed in number of units, we are interested in the probability that demand for one unit of this product group requires x units of component c. We denote these probabilities by the component probability mass function in Equation 4.1.

$$p_{X_{cg}}(x) = P(X_{cg} = x) \tag{4.1}$$

The probability function above is computed for each possible combination of component and customer- & productgroup. For this, we retrieve a file from the ERP system which shows the BOM data of all finished goods delivered in the last 12 months in the format of Table 4.1. Then, we count for each product group the number of times each quantity of a component occurs in the BOM of the finished good. Finally, we divide this by the total number of finished goods of this product group to find the corresponding probability. The output of the first step is a probability function for each component for each product group. Table 4.2 shows an example of the output for one combination of component and product group. This output is used in step 3.

Table 4.2: Example of component probability mass function output

Number of units required	0	1	2	3	4	5	6	7
Probability	0.011	0.147	0.230	0.062	0.356	0.140	0.0517	0.002

4.3.3 Create judgmental sales forecasts of product groups

The second step is to create the judgmental sales forecast for each product group. As we express our probabilities per unit of finished goods, the expected demand from the sales forecast should also be expressed in number of units per group. As explained in Chapter 2 this step is already part of the current way of working. As customers do not provide the same type of information, judgmental forecasts are needed to forecast the demand on product group level. It is important that all forecasted product groups are also available in the ERP-system and all finished goods should be assigned to these groups. If this is not the case, demand on product group level will not translate to component level.

The output of this step is the forecast on product group level for the coming 12 months. This demand is not stable and can fluctuate during the year. An example of the output for one customer/productgroup is given in Table 4.3.

Table 4.3: Example of customer/productgroup forecast

Month	1	2	3	4	5	6	7	8	9	10	11	12
Forecasted units	9	13	30	91	73	127	81	125	37	64	90	117

4.3.4 Calculate component demand from judgmental sales forecasts

In step 3 we use the output of step 1 (probability mass functions) and step 2 (forecasted demand on product group level) as input. When the demand of product group g is known, we can calculate the probability that x units of component c are required for fulfilling the total demand of product group g. In our case, this demand is not known. As explained in Chapter 2, the product group demand at time t is forecasted by the sales department. $F_g(t)$ indicates the number of units requested of product group g. In the calculation below, we assume that the forecasted demand is met exactly. The random variable X_1 can be described by the the probability mass function in Equation 4.1, which gives the probability that x_i units of components c are required in a single demand (expressed as unit) of group g. An extra unit or order of this same group g can require the same random number of units (expressed by X_2) with the same probabilities, which means that their probability mass function is the same. Each extra component demand can be expressed by $X_3, X_4, ..., X_{F_g}$. Then, the probability that the total forecasted demand $F_g(t)$ requires x units of component c equals the summation from $i = 1$ to $F_g(t)$. Denote by $p_{Y_{cgt}}(x)$ this probability mass function:

$$p_{Y_{cgt}}(x) = P(\sum_{i=1}^{F_g(t)} X_i = x) \tag{4.2}$$

We determine the summation in Equation 4.2 by convolution. We assume that each component requirement occurs independent, i.e. $X_1, X_2, ..., X_{F_g}$ are independent random variables. As explained, the $F_g(t)$ are obtained from the judgmental forecasts.

After obtaining the distribution function of component demand on product group level, we want to obtain the distribution of component demand of all product groups combined. We use the same convolution technique as explained before, where the random variable X_g has the probability mass function given in Equation 4.2.

$$p_{D_{ct}}(x) = P(\sum_{g \in G} Y_{cgt} = x) \tag{4.3}$$

4.3.5 Determine near-optimal base stock levels using a greedy heuristic

The total demand of a component D_c is distributed according to the probability mass function calculated in Equation 4.3. Given this demand, we are interested in how much inventory of each component is needed to be able to fulfill the product level demand. We will assume a base stock policy and perform fill rate calculations based on the chosen base stock level. A base stock policy means that a replenishment order is placed when the inventory position at the beginning of period t is lower than s_c to restore the inventory position to s_c. A base stock policy is not optimal due to economies of scale (MOQs and quantity discounts are not incorporated). As we focus on the tactical trade-off between inventories and service levels, operational problems such as order sizes are not relevant at this stage and could be explored later.

Base stock performance evaluation

The total component demand in month t can be filled from stock if the on-hand inventory I_c is higher than or equal to this demand. In a base-stock policy, the on-hand inventory before fulfilling the demand of period t is equal to $S-n$ with probability $P(D_L = n)$. Then, the total demand of that period can be delivered immediately when the demand is lower than the on-hand inventory, which occurs with probability $P(D_c \leq S - n)$. Thus, the probability that all component demand can be filled from shelf, which we call the in-stock probability (ISP) equals:

$$ISP(s_c) = \sum_{n=0}^{s_c} P(D_L = n)P(D_c \leq s_c - n) = \sum_{n=0}^{s_c} P(D_L \leq s_c - n)P(D_c = n) \tag{4.4}$$

More interesting is the fraction of component demand that can be be filled directly from shelf without shortages, which is called the Volume Fill Rate and is known as one of the most commonly used performance measures in inventory control. In literature, this is mostly denoted as the fill rate, Axsäter described this as S_2 (Axsäter, 2006). This measure is also widely known as P_2, introduced by (Silver et al., 2016). When component demand arrives, the on-hand inventory is $s_c - n$ with probability $(D_L = n)$. Then, we can calculate the fraction of components delivered directly from shelf with:

$$FR(s_c) = \frac{\sum_{n=0}^{s_c}(P(D_L = n)(\sum_{x=0}^{s_c-n} P(D_c = x)(x - s_c + n) + s_c - n)}{\sum P(D_c = x)x} \tag{4.5}$$

The first summation is from $N = 0$ to s_c. When lead time demand equals the base stock level, the current periods demand is filled with probability 0. The summation over the current periods demand is from 0 to $s_c - n$, as $s_c - n$ is the maximum amount of the current periods demand that can be filled directly from shelf. We multiply the probability of each demand in this range with $x - s_c + n$, which can be seen as the number of units short. After summation of the multiplications, we again add $s_c - n$ (which is the amount of inventory left) to calculate the expected amount of this periods demand that can be filled directly from shelf, given the demand lead time is n. Dividing this by the expectation of the current periods demand results in the volume fill rate.

In the ISP and FR calculations, the distribution of demand during lead time should be known. The lead times are loaded from the ERP system and assumed to be fixed. Since our demand is non-stationary, the distribution of demand during lead time is dependent on the month. With a given value of ℓ, which is the replenishment lead time, the component demand during lead time is equal to the sum of the component demand from time $t-\ell$ until t. Since our component demand is distributed according to an empirically observed probability distribution, we find the probability distribution of demand during lead time using ℓ-fold convolution of the single demand distributions $X_{t-\ell}, X_{t-\ell+1}, ..., X_{t-1}$. Note that the calculation below assumes that at the beginning of the first period s_c units are on hand with no backorders remaining.

$$p_{D_L}(x) = P(\sum_{i=t-\ell}^{t-1} X_i = x) \tag{4.6}$$

Now we know the probability that a component is available when demand for this component arrives, we can calculate the probability that all required units of component c for the finished good demand of product group g are available from shelf. In the formula below, O is the set of all possible values of X_{cg}. We denote the probability by the PFR (Product Fill Rate) and is equal to:

$$PFR_{cg}(s_c) = \sum_{x \in O} P(X_{cg} = x)FR(s_c)^x \tag{4.7}$$

Consequently, the probability that finished good of product group g can be filled completely from the shelf equals the product of the PFR for each possible component in this group. So, the fill rate of product group g is:

$$FR_g = \prod_{c \in C} PFR_{cg}(s_c) \tag{4.8}$$

Optimization model

Prior to the realization of the configurations, decisions regarding the replenishment control system should be made. The goal of these decisions is to minimize the total costs of purchasing (setup and variable costs), inventory holding and shortage costs over the entire planning horizon, while meeting component demand to make sure all finished goods can be produced. Since we assume that there are no economies of scale, we do not take purchasing costs into account. Given the chosen base stock policy, the expected cost of component c in period t is given by the average inventory of month t multiplied by the price of component c:

$$C_t(s_c) = p_c * \frac{I_c(t) + I_c(t+1)}{2} \tag{4.9}$$

Our objective is to minimize $C_t(s_c)$. In the optimization problem, $S = \{s_1, ..., s_C\}$ denotes the base-stock policy for all components. For each month from $t = 1, t = 2, ..., t = 12$, the optimization problem is as follows:

$$\begin{array}{llll} \min & \sum_{c \in C} C_t(s_c) & & \\ \text{s.t.} & FR_{tg}(S) \geq FR_g^{obj} & \forall t \in T & \forall g \in G \\ & s_c \geq 0 & \forall c \in C & \end{array}$$

4.4 Greedy heuristic

This problem can be solved using different approaches, such as complete enumeration, relaxation or other heuristics. As the base-stock level can theoretically go to infinity, bounds should be defined when using complete enumeration. Then, optimal base-stock levels are found within these bounds. Relaxation or other heuristics can result in near-optimal base-stock levels. In this section, we propose a greedy heuristic to obtain a feasible solution to the problem formulated in the previous chapter. In each greedy step, the base-stock level of one component s_j is increased by one. This results in an increase in the total costs and an increase in the fill rate of the finished goods. This increase in total cost is equal to the price of one extra unit. The increase in the fill rate is equal to the Equation below. In this formula, we take the maximum of $FR_g(s_c + 1)$ and FR_g^{obj} because the increase above FR_g^{obj} does not contribute in achieving the target fill rate of this group as it is already achieved.

$$\Delta W_c = \sum_{g \in G} \max\left(FR_g(s_c + 1), FR_g^{obj}\right) - FR_g(s_c) \tag{4.10}$$

For the greedy algorithm we propose, the fill rate we evaluate should be increasing on its whole domain. The finished good fill rate is the product of the component fill rates that are possibly part of this finished good. When we start with $s_c := 0$ for all components, all component fill rates are equal to 0. When we increase the base-stock level of one component, the fill rate of this component increases to a non-zero number, but the other base-stock levels remain zero. Because of our multiplication, the product of the fill rates remains zero too. As a result, our evaluated fill rate is not increasing on its whole domain. To solve this, we propose a lower bound for each s_c.

We calculated the fill rate of a component by determining the on-hand inventory first. This is equal to the base-stock level minus the demand during lead time. We start fulfilling the demand for a component when the total demand during lead time is completely fulfilled before. Given this, we propose to set the lower bound to be the maximum value of demand during lead time plus 1 unit. This way, the fill rate of a component is not equal to zero, which means the evaluated fill rate is increasing on its whole domain. So, the lower bound is equal to:

$$s_c^{lb} = \max_{\forall t \in T} X_{ct} + 1 \ \forall c \in C \tag{4.11}$$

Finally, define ratio R_c as $\Delta W_c / p_c$. This ratio is the increase in aggregated fill rate per euro. The component k with the highest ratio is 'the biggest bang for the buck' and the base-stock level of this component is increased to $s_k + 1$. The steps are executed until the target fill rate of all finished good groups met.

In Algorithm 1, we describe our heuristic. Computational time is saved by only updating the results that are possibly updated because of the increase in base-stock level. For example, when we increase the base-stock level of a certain component, we only update the fill rate of the product groups that possibly use this component.

Algorithm 1: Greedy Heuristic

Input: s_c^{lb}, p_c, $p_{X_{cg}}(x)$, $p_{D_{ct}}(x)$
Output: Set S with s_c for all $c \in C$

```
s_c = s_c^lb ∀c ∈ C;
S = {s_1, ..., s_c} ∀c ∈ C;
while FR_g < FR_g^obj ∀g ∈ G do
    Determine ΔW_c(S) ∀c ∈ C;
    Determine R_c = ΔW_c(S)/p_c ∀c ∈ C;
    k = arg max R_c;
    s_k = s_k + 1;
    S_k = s_k
end
return S
```

4.5 Verification and validation of developed model

Now we have discussed our solution design, we need to confirm that our method produces valid and feasible outcomes. We used multiple approaches to determine if our model is free of errors. All coding done in Python is checked by going through the code step-by-step. For this, we used a smaller instance to be able to manually check the results of our code. Because of the nature of our solution approach, we were able to implement each step one-by-one. After the verification of a single step, we continued with implementing and verifying the next step. After we verified all steps, we used actual company data. Some of the data is retrieved from the ERP-system AX using SQL. Several meetings took place to make sure the input data is correct. Besides, manual checks are performed to make sure no incorrect data is used as input.

Since everything above is checked and found to be correct, we conclude that our model works correctly and that the output is correct. This means we can validate our model, which we describe in Chapter 5. First, the performance of our forecast is determined. Next to this, we validate our base stock policies by a simulation.

4.6 Conclusion

In this chapter, we designed a model that can calculate the probability distributions of component demand and optimizes base-stock levels based on the found distribution functions. This model first calculates the probability function for each component when a single demand of a product group occurs. Based on the sales forecast input, the model calculates the probability function for the total product group demand. Our method contributes to theory since our model provides a way to find component requirements based on future sales forecasts of finished goods. Usually, component requirements are only determined by forecasting the components directly, but as explained, this method does not capture available future demand. Our approach is extremely useful at companies with short product life cycles, which is the case at Pan Oston.

5 Solution Test

In this chapter, we will elaborate on the results obtained from the solution approach described in Chapter 4. In Section 5.1 we validate our forecasting method by comparing our forecasts of component demand with actual component consumption. After this, we analyse our inventory control policy by using the relevant measures in Section 5.2. Section 5.3 describes our sensitivity analysis, which is performed to check how our outcomes changes with differences in inputs. The chapter is finalized with the conclusions.

5.1 Validation of component forecast

The first step in our proposed model is determining the component demand distribution functions. These distribution functions are input for the fill rate calculations on product level, which is where we apply our marginal analysis on to determine the component base-stock levels. The more accurate these distribution functions, the better our proposed base-stock levels are. Therefore, we are interested in the performance of our component demand forecasts.

To analyse the forecast, we compare the predicted component demand with the actual component consumption. We use different measures, which we explain in Section 5.1.1. First, we compare the performance of common and customer-specific components. After this, we test the performance of the forecast on the demand patterns introduced in Chapter 2.5 to determine the differences in forecastability. For relevant measures on aggregated level, we weight all our individual measures by multiplying with the value of demand and dividing by the total value of demand.

During all comparisons, we use a benchmark forecast to be able to provide a comparison between our proposed method and a more traditional approach. We explain the creation of our benchmark forecast in Section 5.1.2.

5.1.1 Input and output parameters

To be able to determine the performance of our component forecast, we need different inputs. We first compute the component forecasts as explained in step 2 of the model introduced in Chapter 4. To generate the component forecast, we train our model by using all historical data from 2019 to compute the component probabilities on product group level. The other required input for our model is the forecast on product group level. Unfortunately, the product groups in the ERP-system are not in line with the product groups currently forecasted by the sales department. As we compute the probability mass functions based on the product groups in AX, the forecast on product group level should contain the same product groups. Because of this, we decided to use actual sales data of the last 12 months to validate our approach which means forecast errors of on group level demand are not transferred to our component demand forecasts. However, we assume that forecast errors of groups are small due to close customer contact, which means that our method can be validated accurately.

To test the performance of our forecasts, we use a test set. As the actual component consumption of 2020 is already known, we retrieve this data from the ERP-system. To analyse the forecast, we use four forecast accuracy measures for each of the components: MAD, MSE, bias and the WAPE. We calculate these measures by using the expected component demand and the actual consumption, which are both input parameters. The formulas of these output parameters can be found in Chapter 3.

5.1.2 Creating a benchmark forecast

We are interested how our approach performs compared to traditional approaches described in literature, such as time-series analysis. Because of this, we also measure the performance of our method by performing a benchmark against the ARIMA time-series based forecasted method where only historical consumption of components is used. This method does not benefit from sales information. The ARIMA model is trained using component consumption data from 2016 till 2019 and also tested on 2020 to compare the results. Usually, the ARIMA model requires multiple parameters to achieve the best possible forecasting performance. We use the auto.arima function in Python to automatically find the required parameters by using the Akaike Information Critera (AIC), which is a widely used measure to find the best ARIMA(p,d,q) model for each component.

As both ARIMA and our model are unable to forecast components introduced in 2020, we only fit the ARIMA model on the components forecasted in our model to be able to benchmark the forecasting performance of our model.

5.1.3 Performance of forecast on SKU level

All components used in production are classified in a certain group to indicate for which customers the components are used. If a component is common in many different finished goods for different customers, the components are assigned to the 'common components' group, which is the case for 46% of the components. In this section, we will investigate the difference in forecasting performance between these two categories. First, we only distinguish common and customer-specific components.

Table 5.1: Performance forecasting models for a single product category

	Time-series analysis				Proposed method			
	MAD	RMSE	Bias	WAPE	MAD	RMSE	Bias	WAPE
Common	153	182	106	111%	101	128	-76	58%
Customer-specific	447	589	418	509%	55	76	-5	80%

Table 5.1 shows that our proposed method has relatively low weighted errors for all performance measures, compared to the traditional time-series analysis. Contrary to the time-series analysis, our method has a negative bias for both component types, which means that the component demand is under-forecasted. In the time-series analysis, both component groups have a positive bias, which means that on average the forecast is higher than the actual demand. Additionally, the results show that our proposed method achieves a better performance based on the MAD, RMSE and Bias. Because of commonality, you would expect that demand for common components is more stable, which can make forecasting easier. We observe the result of this for the time-series analysis, since a low accuracy is achieved for the customer-specific components compared to common components. Erratic, intermittent or lumpy demand patterns may result in a higher WAPE, which is the case for the customer-specific parts. We will investigate the performance of different demand patterns in the next section.

Based on the aggregate performance, we can clearly see that our method performs well on both the common and customer-specific groups. Compared to the time-series analysis, we can see that the biggest improvement is achieved for customer-specific components. There is a much smaller difference in accuracy between common and customer-specific components in our proposed method than in the time-series analysis, which is the result of the use of sales forecast on product group level. Within the customer-specific components, we can also compare the results between different customers to determine how our model performs for the different customers. When uncertainty in configuration is higher, it can be more difficult to predict component demand. We will focus on the three customers with the most valuable customer-specific components. The performance measures are displayed in Table 5.2.

Table 5.2: Performance forecasting models for customer-specific components

	Proposed method			
	MAD	RMSE	Bias	WAPE
Customer 1	66	90	3	74%
Customer 2	71	111	-25	83%
Customer 3	60	80	-26	64%

Customer 1

The finished goods for this customer have a lot of expensive components not used in products for any other customers. Because of the high amount of customer-specific components, there is a high level of uncertainty on the product configuration. This is also what we see when we focus on the items assigned to this customer. Almost 70% of the customer-specific components are over-forecasted as components are frequently replaced by other components. We clearly see that a lot of components are used in 2019, but arrived at the end of their lifecycle in 2020. Fig 5.1 shows an example of one of these components.

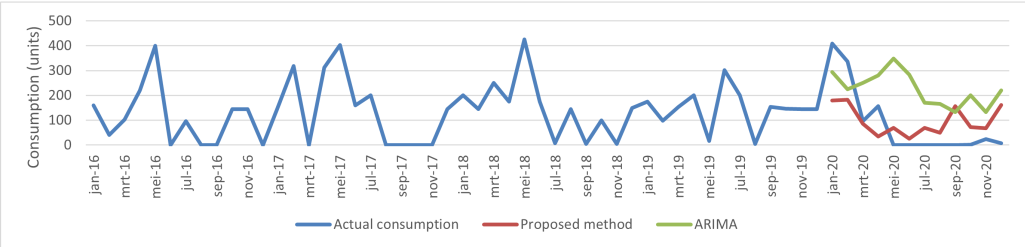

Figure 5.1: End of lifecycle phase

Customer 2

Contrary to customer 1, most of the components for customer 2 have a negative bias, which means that demand is under-forecasted. We do not see products that have a worse performance because they reached the EOL-phase. However, some components have an interesting pattern, which could be the result of an assumption introduced in Chapter 4. We see that components in this group perform well on a 12 months time-period, but have high monthly errors for some components. Figure 5.2 shows an example of a component with a good average performance, but the actual demand occurs one period before the forecasted demand (so our forecast predicts the demand one period late). One reason for this forecasting behaviour could be that we assumed that assembly lead times are negligible, which means demand for a component occurs in the same time period as the demand for a finished good. This is often the case, but Pan Oston often produces the sub-assemblies for this customer to stock. Because of this, actual consumption is earlier than expected.

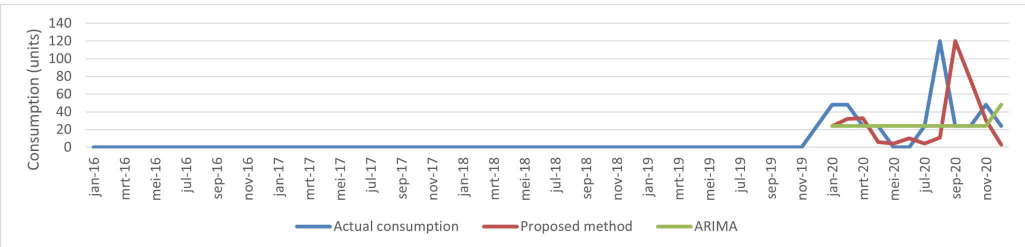

Figure 5.2: Component forecast delay

Customer 3

Table 5.2 shows that the weighted average bias of the components for customer 3 is negative. We see that some components for this customer are heavily under-forecasted, which could be the result of the input used for this analysis. As described in Section 5.1.1, we used the product groups available in the ERP-system as input for our analysis. Unfortunately, these product groups are not as detailed as our model needs, with a lot of finished goods classified in a group where various other products (with a lot of variation in components needed) are also included. Therefore, we chose to not include this group in our analysis. As a result of this choice, not all demand on the product group level is transferred to component demand, which can cause under-forecasting. This is also the case for the example in Figure 5.3. We showed that our model should be able to capture an increasing component demand when the increase on product group level is forecasted by the sales department. However, when not all product groups are forecasted that can possibly use this component, not all component demand is forecasted, which means under-forecasting occurs. Because of this, it is crucial to explore which groups are missing to be able to classify products in more detail.

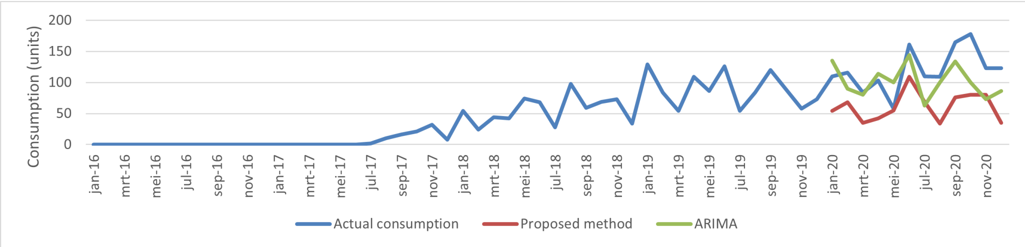

Figure 5.3: Component under-forecasted

5.1.4 Performance of forecast on different demand patterns

Besides our test on the SKU level, we are also interested how our approach performs given various demand patterns. We use the classification of Syntetos et al. (2005) for this. As described in Chapter 3, the demand pattern is an indication of the right forecasting approach. We classify all components into the four groups (smooth, erratic, intermittent and lumpy) based on their historical consumption from 2016 till 2020. Table 5.3 describes the number of components in each class.

Table 5.3: Number of components per demand pattern

Classification	Number of components
Smooth	116
Erratic	91
Intermittent	410
Lumpy	354

Table 5.4 gives the results of the ARIMA method and our proposed method. First, we see that our proposed method performs better on all demand patterns. Since our model uses the forecasts on product group level as input, it is able to outperform the traditional TSA on each pattern. The ARIMA model performs best for the smooth demand, as components with smooth demand have little variability in both demand interval and size. However, the ARIMA model has troubles in forecasting the other demand patterns. For these patterns, the benefit of using our model is significant. In the remainder of this section, we will discuss the performance of our method in more detail.

Table 5.4: Performance forecasting models for different demand patterns

	Time-series analysis				Proposed method			
	MAD	MSE	Bias	WAPE	MAD	MSE	Bias	WAPE
Smooth	167	194	121	93%	109	138	-66	50%
Erratic	193	222	154	157%	89	113	-51	71%
Intermittent	430	116	82	185%	43	54	3	77%
Lumpy	970	1326	933	1086%	51	84	-23	99%

Smooth

Both models achieve a high performance on the smooth demand patterns, compared to the other groups. Figure 5.4a displays a component with a repeating pattern every year. In 2019, the actual consumption already decreased compare to the years before. Due to COVID-19, the demand for finished goods dropped extremely, which also explains the decrease in component consumption. Since we use actual sales data on product group level (instead of the forecasted number of finished goods), the component forecast is not negatively affected. This means we can still measure the performance of our chosen approach in modelling the uncertainty in configuration. From this, we see that our model is able to capture the decrease in component demand when the forecast on product group level is accurate. Figure 5.4b shows another example where actual consumption is lower than expected, which is also captured by our proposed method.

Contrary to these examples, not all components are forecasted accurately. Since the component demand probabilities are calculated based on the finished good configuration data of the past year, components can still be over-forecasted. In Figure 5.4c & 5.4d we see the consumption and forecasted consumption of two components. Both components had relatively stable demand until 2020 and entered the EOL-phase after this. This pattern is not captured by our proposed method.

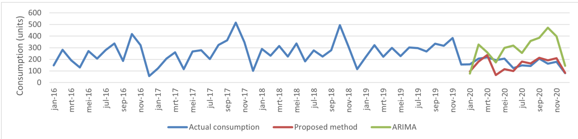

(a) 115995

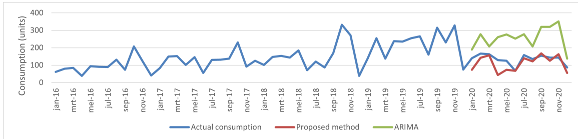

(b) 119706

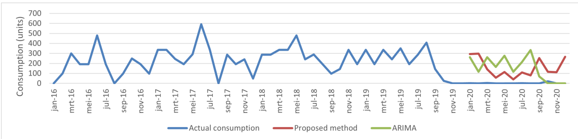

(c) 105005

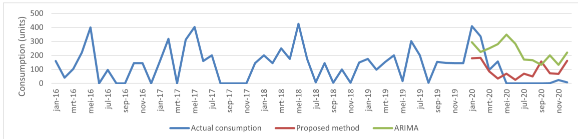

(d) 105531

Figure 5.4: Smooth performance

Erratic

Figure 5.5 shows four components classified as 'Erratic', which means they have regular occurrences in time, but also have high quantity variations. The components on Figure 5.5a and 5.5b are both used for multiple customers, but most of the consumption comes from demand on product level of one customer. Our model is able to capture the uncertainty in configuration.

On the other hand, we also see components with an erratic pattern where our method does not perform extremely well as demand is under- or over-forecasted. On Figure 5.5c a demand peak is expected in the second month, but the actual demand is equal to zero. Two months later, demand is predicted and comes very close to the actual demand. But, in the final months of 2020 we expect demand since there is demand on product group level where the component is previously used frequently, but apparently the component is used in less finished goods. Figure 5.5d displays a component which is used in almost all finished good demand for a specific product group. In the first months of 2020, there is peak in demand which is not recognized by our method, which resulted in an under-prediction of more than half of the monthly demand. Remarkable is that the actual demand declines afterwards, while our method predicts a relatively stable demand.

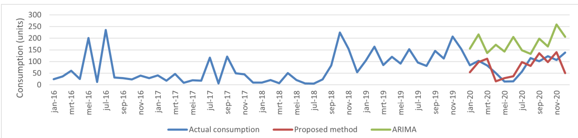

(a) 115104

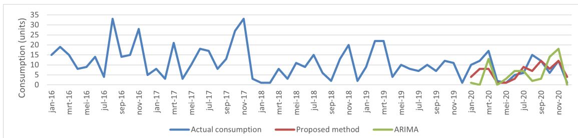

(b) XK50K035

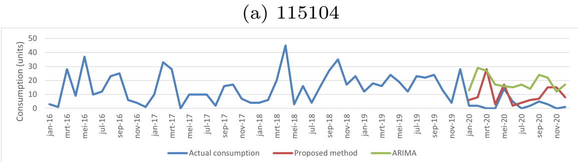

(c) 115109

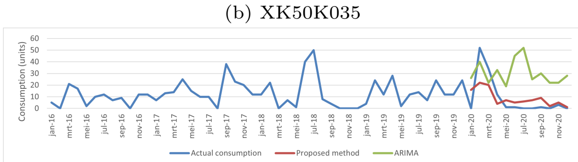

(d) 116431

Figure 5.5: Erratic performance

Intermittent

We also see that both methods have problems in forecasting the intermittent demand pattern as intermittent demand has a lot of periods with zero demands. ARIMA is not able to accurately predict when demand occurs. Our proposed method achieves a better performance. On Figure 5.6a and 5.6b two customer-specific components are displayed, both predicted accurately. These components have very unstable demand, but since the forecast on product group is accurate and our calculations are able to translate this to component requirements, we are able to exactly estimate the number of units required to fulfill product level demand.

This is not the case for both customer-specific parts on Figure 5.6c and 5.6d, where component demand is over-forecasted. The first component is still used in product level demand, but less than 2019 where our component probabilities are based on. The second component is only used in the first months of 2020 and was replaced by another component afterwards.

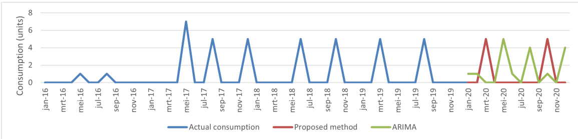

(a) XH1LH034

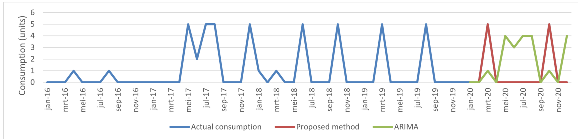

(b) XK50K054

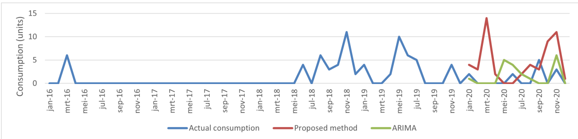

(c) XH1CH449

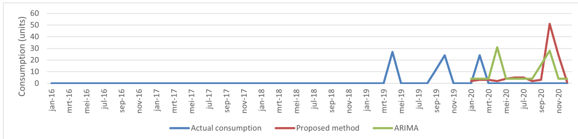

(d) XK20K035

Figure 5.6: Intermittent performance

Lumpy

As described, lumpy demand is characterized by a sporadic and irregular demand pattern. This means that both the timing and the size of demand are uncertain. From the WAPE performance measure in Table 5.4 we can conclude that our method has most troubles in forecasting this demand, but compared to ARIMA the decrease in performance compared to other groups is relatively low. During 2019, the component on Figure 5.7a is introduced for products for one large customer. Because of this limited history is available for our proposed method and ARIMA. However, we see that the peak demand is predicted accurately. The component on Figure 5.7b has more history, but has clear differences in the size of the demand. As our proposed model is able to predict the component demand exactly, we can conclude that we benefit from the sales forecast on group level. However, as Figure 5.7c and 5.7d show, this is not always the case. Both components entered the EOL-phase, which is impossible to capture with quantitative information only.

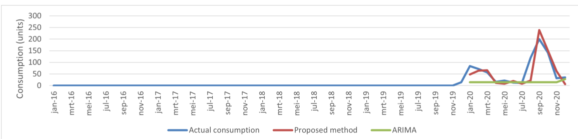

(a) 105383

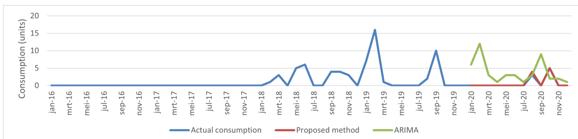

(b) 118431

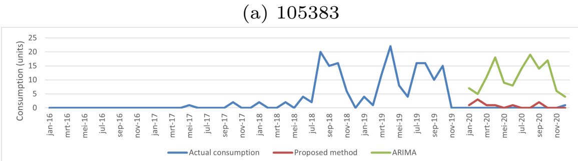

(c) 115113

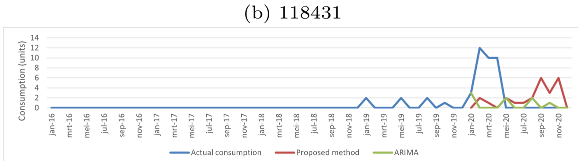

(d) XK08K003

Figure 5.7: Lumpy performance

5.1.5 Validation with component consumption of forecasted group level demand

In the previous sections, we found that our method tends to under-predict the component consumption, which we think is the result of not all product groups being forecasted. In this section, we use the component consumption of the forecasted product groups only to verify this claim. We extract the consumption coming from these groups from the ERP-system. As the date of the actual consumption of the component is not available, we use the date the finished good is requested by the customer as consumption date. Table 5.5 shows the performance of our proposed method. We are unable to calculate the WAPE for intermittent and lumpy demand, since the actual demand for some components in these groups became zero. However, all calculated performance measures improved. From this analysis, we can conclude it is crucial to provide forecasts for all products groups to achieve accurate component forecasts. We still see a negative bias for three demand patterns, but we checked multiple components with a negative bias and found that these under-predictions are mainly caused by unpredictable consumption peaks.

Table 5.5: Forecasting performance consumption from forecasted group demand

	MAD	RMSE	Bias	WAPE
Smooth	38	48	-17	18%
Erratic	51	60	-44	30%
Intermittent	12	20	2	*
Lumpy	13	18	-10	*

*** Undefined, dividing by zero**

Summarizing, we clearly see that our introduced model performs well. ARIMA is outperformed especially in cases where demand is unstable and the sales forecast on product group level gives information on the timing and size of the expected component demand. However, components that entered the EOL-phase or are introduced recently are still unpredictable. Because of this, human judgment is needed to account for these errors. The validation also showed that our method tends to under-forecast. We can attribute this negative bias to the method we used to create the input required for our test. Some products are categorized in the wrong product group, resulting in inaccurate component probabilities. To avoid this, we only used the demand on product group level of the three main categories described in Chapter 2. We showed that the performance of our model improved when we only used the component demand coming from the finished good demand of these three main categories. In the next section, where we evaluate the performance of our proposed inventory control policy, we will only use this component demand to be able to get a fair measure on the performance.

5.2 Performance of proposed inventory control

The inventory control policy proposed in Chapter 4 will be tested in this chapter. To test the proposed policy, a simulation model in Python is created which we explain in this section. In this simulation, we use the found base-stock policies together with the actual consumption of the year 2020, which is also the time-period we used to validate our forecast. In the simulation, we determine the average inventory level with our proposed base-stock policies. We compare this with the current inventory performance of Pan Oston. One of the measures is the inventory value per item, which can be calculated by the cost of the item multiplied by the average inventory level. We first run our model to determine our component demand distributions functions. Based on this output, we run our base-stock model using a fixed target fill rate of 90% for each finished good group. For the evaluation of the performance, we first calculate the average inventory levels given the base stock levels we found in our greedy heuristic. Given our base-stock policy, we can calculate this by averaging the inventory level at the start of the period and the inventory level at the end of the period. The average inventory level is multiplied with the unit price to get the average inventory value. We compare the average inventory value with the actual inventory value of 2020, which is recorded in AX. Unfortunately, the service levels on product group level are not measured, which means we are unable to provide a fair comparison. However, we are sure that our inventory policy meets all target fill rates if the demand is predicted accurately.

In our forecast validation, we were interested how our forecast performs for all components. To test the performance of our proposed inventory control, we only measure the performance by comparing the performance of the components that currently have a base-stock policy. These base-stock policies are determined by the judgments of the purchasing department, who are in close contact with the sales team. The other components are ordered when there is actual demand.

5.2.1 Input and output parameters

To be able to perform the simulation, we need different inputs which we retrieve from the ERP system. The first input we use are the actual component inventory levels per month in 2020. Besides, we use the actual consumption during this same time period. However, we perform one change to test our inventory control performance because of our forecasting performance. We showed that our method under-forecasts demand for some components, since not all demand on product group level is forecasted. Because of this, we only extract the consumption of 2020 coming from the forecasted demand on product group level. To determine when to restore the inventory position, we use the proposed base-stock levels, which we determined in our greedy heuristic. For our greedy heuristic, we need the target fill rate for each group. Currently, no data is available about service differentiation at Pan Oston. Because of this, we chose to set the target fill rate for all groups to 90%. Finally, for our simulation we also use the lead time of the components as input to determine when the ordered units arrive or what is still in the pipeline.

The output parameters are the parameters from which we can measure the performance of our proposed base-stock policies. The first output is the average inventory level per item, which is the summation over the inventory level per month divided by 12. Then, we multiply this by the price of the component to determine the average inventory value. Besides, we also calculate the true fill rate of the component, which is the sum of the fulfilled monthly demand from stock divided by the sum of the total monthly demand. Other outputs are the number of backorders and the amount of units ordered.

5.2.2 Simulation results

The results of the comparison can be found in Table 5.6. We see clear differences in the inventory values given our base stock levels compared to the original inventory value. However, we also see that some components are still hard to forecast, resulting in an under-forecast of the demand. Because of this, our average item fill rate is lower than the expected fill rate. Based on this, we can conclude that our base-stock policy performs well if the component demand forecast accuracy is high. As we already showed, our model is able to capture sudden changes in demand. However, we still see components where the actual fill rate is lower than expected because of an inaccurate component forecast.

Table 5.6: Total change in inventory value

Classification	Old inventory value	New inventory value	Improvement	Average item fill rate
Smooth	€365.740	€186.649	48,9%	97,9%
Erratic	€108.057	€65.722	39,1%	94,0%
Intermittent	€255.621	€168.092	34,2%	92,8%
Lumpy	€233.510	€143.594	42,2%	88,9%

Table 5.7 shows the biggest contributors to the improved inventory value and this also confirms that the inventory value is decreased, but the simulated fill rate is not as expected for two of the five components. Figure 5.8 displays the difference between the expected and simulated fill rate compared to the in- or decrease in average monthly inventory value per component. A negative percentage on the x-axis means that the fill rate

is lower than expected. The y-axis shows the in- or decrease in average monthly inventory value. A negative value on the x-axis means that the average monthly inventory value with our proposed base-stock levels is lower than the current policy. From this figure, we can conclude that our base-stock levels can not guarantee the expected service level for all components. Based on earlier configurations and the forecast of demand on product level, our model expects that demand for some components is low, which means a low base-stock level can guarantee a high fill rate. When demand for a component is higher than expected the actual fill rate will be way lower, which is what we for some components. The difference between the expected and simulated fill rate is more than 20% and no big savings are achieved by using the proposed base-stock levels. However, we can also conclude that our proposed base-stock policy performs well when demand is predicted accurately. The biggest contribution is achieved by expensive components, where our model is able to predict demand accurately based on the aggregated forecasts on product level and past configuration data.

Table 5.7: Biggest contributors to improved inventory value

Component	Monthly inventory value decrease	Expected fill rate	Simulated fill rate
115436	€32.633	99,2%	100,0%
105544	€31.850	93,5%	74,0%
105567	€19.879	98,0%	58,0%
115970	€17.111	99,1%	100%
118435	€9.149	99,9%	100%

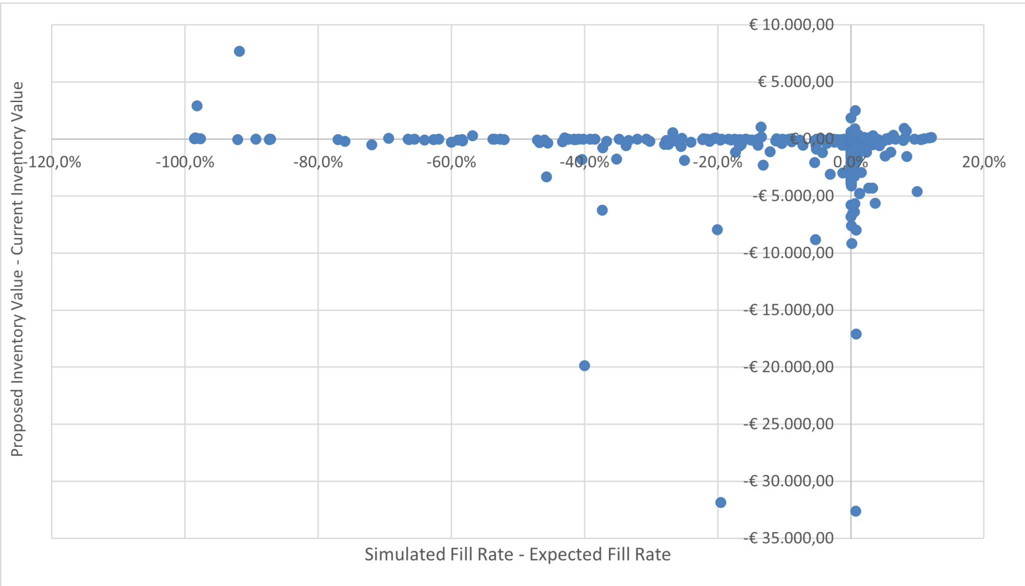

Figure 5.8: Fill Rate difference versus inventory decrease

5.3 Sensitivity analysis

In this section, we will examine how the outcomes of our model changes when input variables are changed. As mentioned, our model uses the following input variables:

- Component unit costs
- Forecast on product group level
- Component lead time
- Target fill rates

The goal of our sensitivity analysis is to provide insights on our model outcomes given different inputs. Therefore, we are only interested in adjusting input variables that can be adjusted by Pan Oston in real life. So, we only focus on adjusting the expected demand on product group level from the sales plan and the target fill rates of the customer groups.

5.3.1 Impact of changes in forecasted demand on product group level

As discussed, we assume that the forecast on product group level is fixed. In reality, the aggregate forecast can have errors which are outside the scope of this research. However, it is interesting to see how our method behaves under product group level demand with small deviations. We will change one input variable at the time, keeping other variables equal to find the impact of the single input variable on the output variable.

We decide to change the forecasted demand of the six most important customers with 10%. In our method, a different expected finished good demand results in differences in the component requirements where our base stock levels are based on. We test the base-stock policy using the same simulation with actual component consumption as proposed in Section 5.2. Table 5.8 gives the results of our sensitivity analysis. In the left column, the demand is under-forecasted with 10%. The right column gives the increase in inventory value as a result of a 10% over-forecast. As the table shows, an under-forecast of 10% does only result in a small inventory value decrease. When demand of group 5 is under-forecasted, the inventory value decreases more. This is as expected, since components used in this group are more expensive than component used in other groups. The table also shows that an over-forecast has a bigger impact on the inventory value, since we see that the inventory value is increased by more than 3% for 4 of the 6 product groups when demand is over-forecasted.

Table 5.8: Impact of forecast error on product group level

Productgroup	Inventory value difference under-forecast	Inventory value difference over-forecast
Group 1	-1.01%	1.03%
Group 2	-0.33%	3.23%
Group 3	-0.30%	3.17%
Group 4	-0.37%	0.38%
Group 5	-3.69%	3.83%
Group 6	-0.25%	3.14%

5.3.2 Impact of changes in target fill rates

In our performance test, we used a target fill rate of 90% for each product group. However, the fill rates of the current policy are unknown and therefore also no target is set. Because of this, we are interested in the impact of changes in the target fill rates. We chose to use the following experimental settings:

Table 5.9: Experimental settings target fill rates

Target Fill Rate	10%	20%	30%	40%	50%	60%	70%	80%	90%	95%	97.5%	99%

The results per target fill rate as proposed in Table 5.9 are illustrated in Figure 5.9. From this figure, we can conclude that the marginal effort to arrive at a high fill rate (in terms of inventory value) is higher than when the fill rate is low. For example, the increase in average inventory value from a 10% fill rate to a 50% fill rate is €1.494 per month, which is lower than the increase needed to go from a 90% to 99% fill rate (which is €1.582 per month).

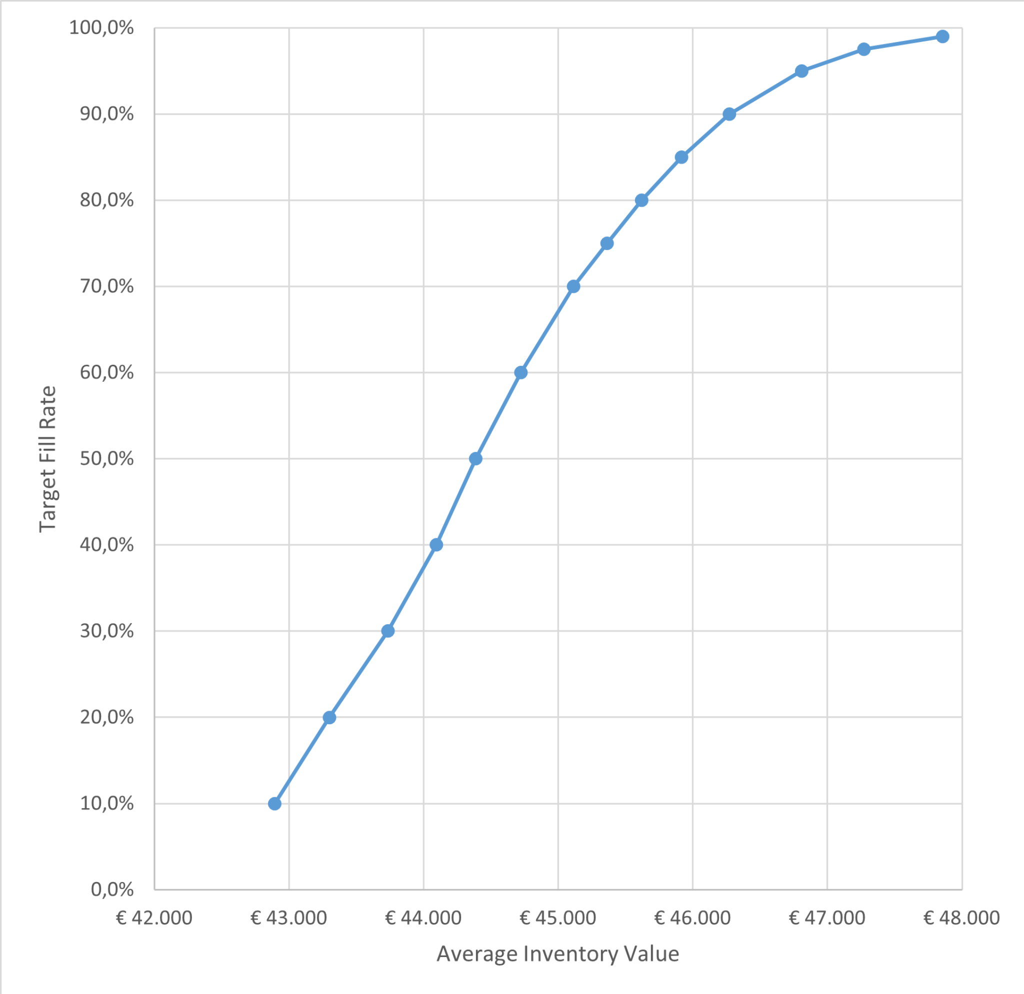

Figure 5.9: Target Fill Rate Sensitivity

5.4 Conclusion

In this chapter, we elaborated on the results following from the method proposed in Chapter 4 and answered the research question *What is the effect and improvement of the proposed method?* First, we compared our proposed forecast approach with a traditional time-series analysis and showed that our approach outperforms this method given the demand characteristics of Pan Oston. However, during the solution test, we found out that our method has the general tendency to under-forecast component demand, which is the result of lacking demand forecasts on product group level. Our method is able to handle the order configuration uncertainty. But, since we calculate the component probabilities based on last years order configuration history, the probabilities are lower than actual in case of new introductions. Because of this, human judgment is still needed to reflect on the outcomes of our model.

To test the performance of the proposed greedy heuristic, we developed a simulation study where we only used actual component consumption from the product groups forecasted. Based on this simulation study, we can conclude that our proposed greedy heuristic performs well. The inventory value can be reduced for a lot of components. However, we also noticed that the actual fill rate of some components differ from our expected fill rates. It is also hard to compare the proposed policy with the current, since service levels are currently not used and not measured. We also included the start value of the inventory in our simulation, which can give an unrealistic performance of our proposed method.

Our sensitivity analysis showed us that the average inventory value increased more when demand is over-forecasted compared to the decrease of the inventory value when demand is under-forecasted. Based on our sensitivity analysis regarding the target fill rate setting, we can conclude that the target fill rate is increasing when more inventory is held. However, the growth in target fill rate is decreasing. In other words, when the fill rate is low the marginal increase with the same amount of extra inventory is higher than when the fill rate is higher.

6 Solution Implementation

In the previous chapter, we analysed the performance of our proposed method. In this chapter, we explain how Pan Oston should implement our method. We first present the created tool in Section 6.1. In Section 6.2, we explain which steps are required to implement the model by using this tool. Section 6.3 explains how the tool should be maintained after implementation. Section 6.4 describes the technical risks but also risks related to change management during the implementation and use phase. The chapter is finalized with conclusions.

6.1 Python tool

The solution model is developed using Python and has an interface that is developed for Pan Oston. The interface is created to make sure employees can use the tool without having in-depth knowledge of Python. In the interface, input files can be loaded, the model can be executed and output data is displayed. Figure 6.1 shows an example of the main interface of the tool with fictive input and output. After selecting all input files, the component probability mass functions can be calculated. After execution, the probability matrix can be displayed and the algorithm proposed in Chapter 4 can be run. Finally, the base stock policies can be shown. To gain insights for the purchasing department, we also give the difference between the current and proposed base stock levels to be able to select the components that need the most attention.

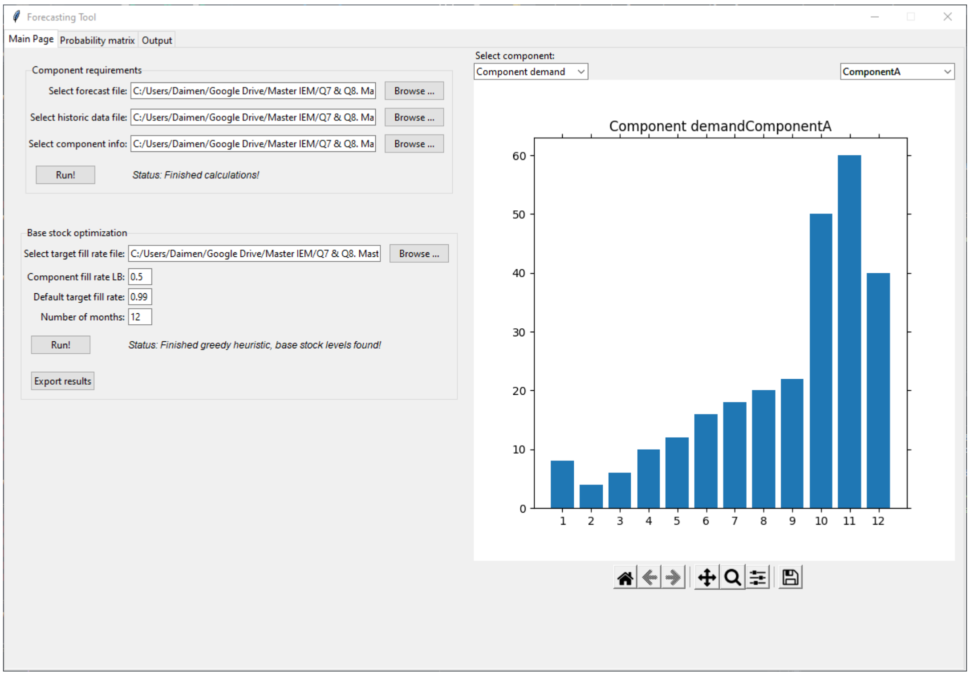

Figure 6.1: User Interface

6.2 Roadmap for implementation

To implement the model in order to compute efficient base-stock policies, a few steps are required. First of all, a team with employees from all involved departments must be composed. This team should include employees from sales, purchasing and the AX-team. The first thing that the team should do is to agree on the product groups and make guidelines to make sure that newly introduced products are assigned to the correct product group in the future. The created product groups should be as detailed as possible, but sales employees should also be able to forecast the number of units of each product group within reasonable time. Then, these forecasts should be extractable so they can be used as input for our model.

The next thing which is important to make the tool work is to agree on the target fill rates for each product group. As the fill rates are used to determine the monthly base stock levels for each component, it is crucial to agree on the fill rates to make sure the right base stock levels are found. These fill rates should also be stored in a file that can be used as input for our tool.

Regarding the technical implementation, it is important to provide the historical order data in the same way as it is currently extracted from the ERP system. The model is developed in Python to make it work with the current ERP-system as the input for the tool is directly extracted from AX. However, Pan Oston already announced the introduction of a new ERP-system. If it is not possible to extract the historical order data in the same way, the data should be pre-processed or the Python code should be adjusted.

6.3 Maintaining the tool

It is important to also maintain the tool after a successful implementation. As described, Pan Oston is a company where products have short life-cycles, which is also the case for the components. This means that EOL-components and introductions should be managed accurately. Besides, prices and lead times are also subject to changes. It is important that the component lead times and prices are updated by the purchasing department frequently. Furthermore, the sales team should provide the expected quantities on product group level each month. As a result, new base stock levels can be computed by the purchasing department.

The component probabilities computed from the BOM information are bound to change. Since we use the historical orders of the past year, each day changes to the probabilities of component usage in the product groups are possible. During the introduction-phase, a component should be purchased the same way as the current way of working. When a product is produced and sold with a new component, the probability of this component in this product will automatically increase until the component is present during the entire planning horizon. So, a product needs to be introduced one year ago to compute an accurate probability distribution function. This is the same for EOL components. When the decline phase of a component starts, the probability of this component in finished good will decrease, as it is used in less finished goods. After a full year without a single occurrence, the model will return a demand of 0 with a probability of 1 for this component. So, the purchasing department has to keep this in mind during the interpretation of the model outcomes.

However, the procedure above is not needed when a completely new product group is introduced. After introduction, there is no historical data on which components are required to produce finished goods in this product group. If this is the case, our model is not able to determine component demand coming from this product group. However, the required data comes available directly when a sales order is delivered, which means we can calculate a probability distribution function directly. It remains important to interpret the outcomes of the base-stock policies for these components, since the data sample used to compute the distribution function is small.

6.4 Risks

In the previous sections, we discussed how the tool should be implemented and maintained. During both phases, several risks can show up, which we can categorize as technical risks and risks related to change management.

6.4.1 Technical risks

The technical risks of this project are:

- Historical data is not available in the new ERP system.
- Data conversion from old ERP not done correctly.
- Adjustments to tool needed, but insufficient knowledge available.

The technical risks are related to the required input data. The component probabilities are calculated using the sales order data of the past year. When historical data is not available due to the implementation of the new ERP-system, our model is unable to compute the component probabilities. Finally, it could occur that adjustments to the tool are needed, but that insufficient knowledge within the company is available to make these changes. This risk can make our tool useless in the future.

6.4.2 Change management risks

Also risks related to change management can show up. The most important risks are:

- Risk of accepting the output of the model.
- Changes in component demand not captured.
- Tool not used anymore.

The first risk is accepting the output of the model without interpretation of the output. We showed that our model performs well, but not always captures sudden changes in demand as the probabilities could be wrong when products are phased out or recently introduced. So, when quantitative information is present, which is not reflected in the qualitative model, the base-stock level should be properly adjusted. Another risk is that the employees change back to their old behaviour and stop using the tool. This can occur when they do not see the added value of the tool or when the tool is hard to maintain.

6.5 Conclusion

In this section we presented our tool and described how it should be implemented. Moreover, we also described how the tool should be maintained and which risks can possibly occur. We recommend to first agree on the product groups to forecast. Then, all products should be categorized in the correct product group to make sure we can use this input in our tool. Besides, the target fill rates for each product group should be agreed on.

When the tool is implemented, component information such as lead times and prices should be updated frequently to achieve the best greedy heuristic performance. Besides, we advise to run our tool each time the sales forecast is updated to make sure deviations in the proposed and actual base-stock levels can be found at an early stage. Finally, we described several risks related to the technical implementation, but also related to change management.

7 Conclusions and Recommendations

This chapter gives an overview of the main findings of the research. Section 7.1 focuses on answering the research question of this study. Thereafter, we provide our recommendations and discuss the limitations and propose some topics for further research. The chapter is finalized by the practical and scientific contribution of our study.

7.1 Conclusion

We used a problem cluster to identify the core problem. We found out that due to customer behaviour products are only forecasted on an aggregate level, with limited insights in the product configuration requested which makes it hard to forecast component demand. Besides, inventory control parameters were based on gut feelings. Based on this, we formulated the following research question:

How can Pan Oston use available historical and future sales data to improve their inventory levels, while maintaining their current service levels?

To answer this question, we first identified the current situation. The current sales forecasting process is similar to approaches from other companies. Since Pan Oston is a project-based company, selling their products to other companies only (B2B), they use a judgemental sales forecasting method based on customer information. We can distinguish three distinct time zones during the forecast horizon: the first consists of customer orders only, the second is a combination of customer orders and forecast and the third consists of forecasts only. Within Pan Oston, both ETO-, CTO- and FTO-strategies are used. These strategies introduce the configuration uncertainty, which makes it hard to translate the sales forecast to component requirements.

In a literature review, we first explored the supply chain management in similar production environments. With ATO and CTO strategies, finished goods are assembled from components in stock. We described the demand planning process and focused on the inventory control and supply planning process in CTO environments. We found several optimization methods for the supply planning in ATO and CTO operations. However, all of them assumed stationary demand on product group level, which does not hold in the context of Pan Oston.

In the first part of our research, we developed a method to compute the component demand probabilities. We implemented this method using Python. The method uses historical configuration data and forecasted sales on product group level as input. We compute the probability mass functions of component demand of the future sales. In the second part of our research, we developed a model that can optimize base-stock levels given the forecasted demand of components is known.

We tested the performance of our component forecast using different performance measures. Remarkable is that our method does perform best for smooth patterns, but the difference compared to other patterns is smaller than expected. From this, it can be concluded that the use of the sales forecast on product group level is beneficial to the forecastability of a lot of components. In our simulation study, we investigated the performance of our proposed base-stock levels on the components where base-stock levels are currently also used. According to this study, the proposed base-stock levels are an improvement when forecasting accuracy is high. The table below depicts per demand pattern how much the average inventory value is reduced by implementing the proposed base-stock levels. However, we also showed that when demand is under-forecasted, the actual fill rate is lower than expected. This means that our base-stock levels will result in a lower component availability. To conclude, our quantitative inventory model provides a good basis for the parameter settings, but human knowledge and judgement are still very important in setting the correct settings.

Table 7.1: Inventory value reduction per demand pattern

Demand pattern	Smooth	Erratic	Intermittent	Lumpy
Inventory value reduction	48,9%	39,1%	34,2%	42,2%

7.2 Recommendations, Limitations and Scientific Contribution

7.2.1 Recommendations

The main recommendation resulting from this study at Pan Oston is to implement the tool proposed in this report. For this, the implementation plan of Chapter 6 can be used. Each month, after the adjustments on the sales forecast, the tool should be ran to check the optimized base stock policies. Despite a worse performance for some components due to unpredictable demand, the proposed method gives the ability to make the trade-off between a fill rate on product group level and inventory costs. Because of this, we advise to use the output of the tool to compare the proposed base stock levels with the current base stock levels and make decisions based on this.

Furthermore, we used product groups available in the current ERP system to cluster products. As Pan Oston has plans to introduce a new ERP system, it should be investigated which groups are needed for which customers and to incorporate these groups in the new ERP. These groups can also be beneficial for other operational or strategic purposes, such as measuring the performance regarding business objectives.

Finally, we recommend to maintain the input of the tool frequently. Our model uses historical data to estimate the order configuration uncertainty and derive component demand from this. To keep the probability density function accurate when products are phased out or introduced, the input can be adjusted or manual judgment can be used to correct the under- or over-forecasts.

7.2.2 Limitations and Further Research

Our proposed solution has a few assumptions and limitations. First of all, one limitation is the accuracy and reliability of data used in this research. We noticed that some finished goods were not classified in the correct product group. As we use these groups as input of our mathematical model, this can have consequences for the outcomes. Because of this unreliable data, we were also unable to validate our approach based on the sales forecast on product group level, instead we used actual sales quantities of the product groups of the past year.

Besides the unreliable data, we used target fill rates for each product group to find near-optimal base-stock levels. Unfortunately, the current fill rates are unknown, so it is unknown if the used target fill rates are close to the current fill rates. Because of this, the performance of our model and the current method are hard to compare. We performed some experiments to test the impact of lower fill rates on the inventory value, but were still unable to make a fair comparison with the current situation.

A third limitation is found in the use of historical data to estimate how finished goods are configured. Our proposed method uses historical data to estimate the component demand. We showed that our base-stock policies perform well, but that some components are still unpredictable resulting in a low fill rate and/or high inventory costs. This is especially the case when products or components are phased out or introduced. Information from sales is still required to account for this.

In addition to the limitations, we can also provide directions for further research. These are outside the scope of this research, but can be interesting for Pan Oston to investigate in the near future. An aspect that is not considered in this study is lead time uncertainty. The effect and how this uncertainty can be incorporated in our model can be investigated in the future.

Another aspect that was not considered in this study is the demand uncertainty of the finished goods. We assumed that these forecasts were accurate and modelled it as fixed. Currently, these forecasts are only based on the judgments of experts as Pan Oston is a B2B company with close customer contact. However, it can be beneficial for Pan Oston to investigate the effect of using machine learning techniques to generate a statistical forecast in the future. Machine learning techniques can be useful for estimating success probabilities of orders.

Next to this, we recommend Pan Oston to explore customer differentiation opportunities. In the first part of our sensitivity analysis, we showed that an over-forecast on product group level for one customer resulted in a very high increase on the average inventory level compared to other product groups. We tested different target fill rates for all product groups, where a higher target fill rate resulted in higher average inventory levels. It could also be interesting to investigate the impact of customer differentiation on the target fill rate.

Our method uses historical data to train estimate the configuration uncertainty. We advise to update the probabilities each time the sales forecast is updated, which is monthly. In our approach, this means that the first month is left out and the most recent month is added to recalculate the probabilities. Further research should determine if there are possibilities to incorporate an updating procedure (such as Bayesian updating) where all past data is used.

Lastly, further research should investigate how Pan Oston can incorporate assembly lead times in the model. As explained, we assumed that the assembly lead times are negligible, so demand for component occurs in the same month as demand for the finished good. This assumption holds for a lot of product groups, but is not the case when sub-assemblies are produced to stock. It could be beneficial to determine the average assembly lead time per product group and find ways to incorporate this in our mathematical calculations.

7.2.3 Practical and Scientific Contribution

We created a decision support tool that Pan Oston can use to find base-stock levels for their components under product group fill rates constraints. In the current practice, a lot of decisions are based on judgments from employees. Our approach is a data-based method that uses past data to find component probability functions of product groups and combines this with the sales forecast on product group level to determine monthly component demand. Then, the user of the model is able to change the desired fill rates and optimize inventory based on the chosen fill rates. This means that the employees are able to make the trade-off between a higher average inventory value and the fill rate by themselves.

The scientific contribution of this master thesis lies mainly in the design of the solution, which integrates different modelling techniques to compute base stock policies. We presented a mathematical model to compute base-stock policies for components for a manufacturer with non-stationary demand on product group level using probability mass functions. We solved instances with more than 800 (customer-specific and common) components and over 80 customer- and product groups. To the best of our knowledge, our approach to calculate non-stationary base-stock policies is new, since other literature only assumed stationary demand for tractability. The solution introduced in this report is suitable for other production companies that experience the same type of uncertainty in their process.

A User manual

This appendix contains the user manual of the tool. It is written in Dutch.

Handleiding toepassing

Deze handleiding is geschreven voor het gebruik van de toepassing om een voorspelling van de componentvraag te verkrijgen. De handleiding geeft eerst uitleg over de benodigde data en het format waarin de data gepresenteerd moet worden. Daarna wordt de interface gepresenteerd, wordt beschreven hoe de toepassing gebruikt moet worden en welke mogelijke fouten kunnen optreden tijdens het gebruik.

Inputdata

De gecreëerde toepassing heeft een drietal bestanden nodig om allereerst de componentbehoefte te kunnen voorspellen en vervolgens base-stock levels aan te bevelen. Een vierde bestand is optioneel. Belangrijk is dat alle inputdata voorzien is van kolomkoppen, zodat de toepassing de juiste input kan selecteren. In de map 'Voorbeeldbestanden' zijn bestanden geplaatst die het juiste format hebben. Hieronder wordt beschreven hoe deze bestanden eruit zien.

Sales forecast

Allereerst moet de sales forecast per productgroep ingeladen worden. Belangrijk is dat de tool 12 maanden vooruit een voorspelling nodig heeft om te kunnen werken. Indien dit voor een bepaalde productgroep onbekend is, moet '0' ingevuld worden.

Historische sales order data

Ten tweede heeft de toepassing data nodig waarop de historische configuratie van de eindproducten wordt bepaald. Dit bestand moet vanuit het ERP door middel van een SQL-query geëxporteerd worden naar het .csv format. De volgende kolommen moeten ingevuld zijn (en kolomkoppen moeten exact overeenkomen):

– Date (datum)

– CustID (klant ID)

– SalesID (salesorder nummer)

– FG (eindproduct ID)

– FGQty (aantal van eindproduct)

– Customer (klant)

– CatFG (productgroep)

– ItemID (component ID)

– ItemQty (aantal componenten in product)

Component gegevens

Voor het uitvoeren van de voorraadoptimalisatie zijn een aantal gegevens van alle componenten benodigd. Deze data kan vanuit het ERP systeem naar een .csv bestand geëxporteerd worden en moet de volgende kolommen bevatten: component, levertijd, prijs en eenheid.

Gewenste fill rate

Optioneel is een fill rate per productgroep. Voor elke productgroep kan een fill rate ingegeven worden. Dit is het percentage van producten die geleverd kunnen worden waarbij alle componenten voorradig zijn. 'Alle componenten' zijn de componenten die in de base-stock bepaling worden meegenomen. Componenten die voor de base-stock policy uitgesloten worden, zullen niet meegenomen worden (dus er wordt aangenomen dat deze componenten geen problemen opleveren voor uitlevering eindproduct).

Gebruikersinterface

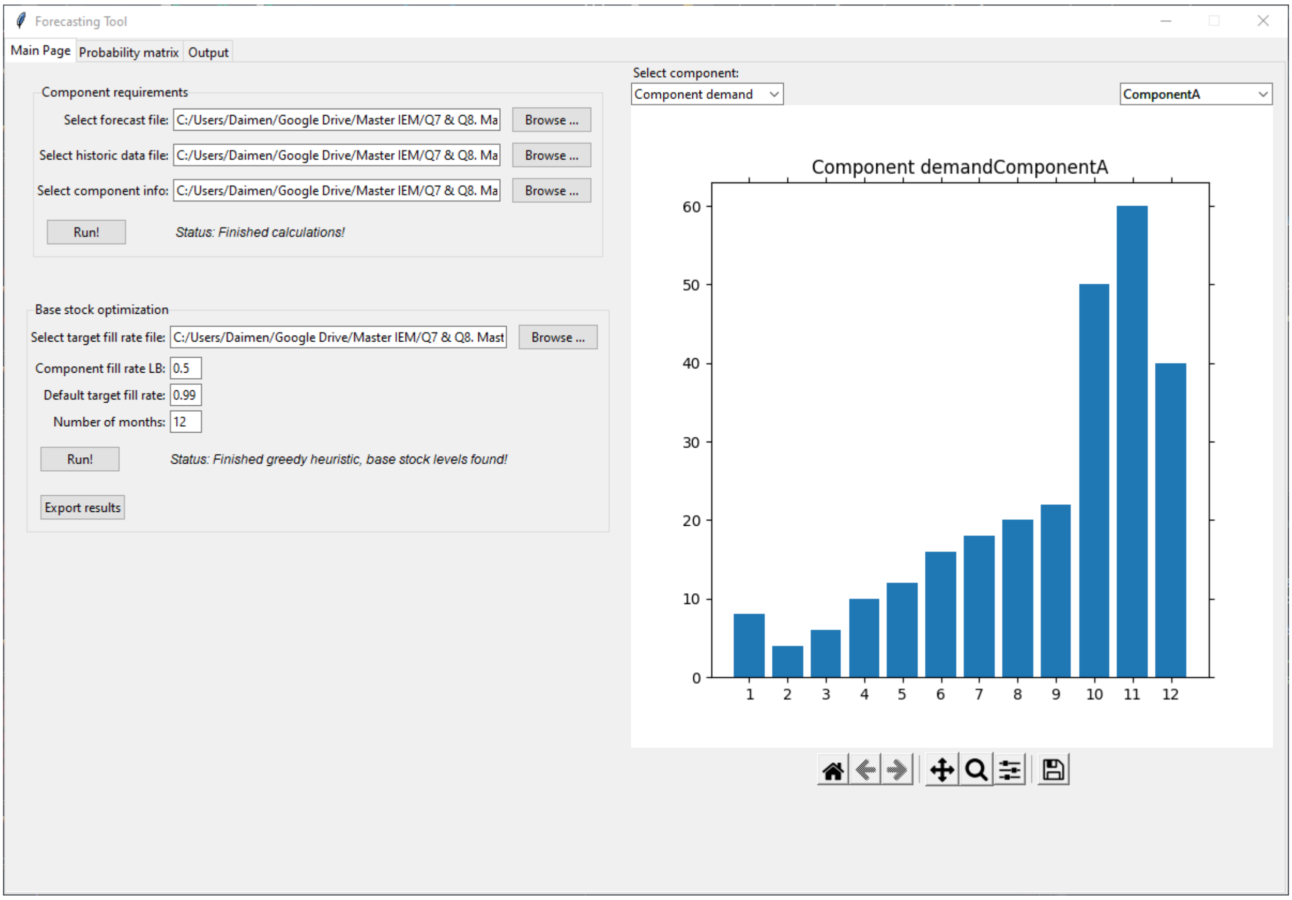

Gebruik van de toepassing

Stap 1: Laad alle data in het programma

Klik op de 'Browse . . . ' knoppen om alle bestanden in te laden. Het scherm hieronder zal weergegeven worden, waarna het bestand geselecteerd kan worden. Om te controleren of de sales forecast juist is ingeladen, kan de grafiek aan de rechterkant gebruikt worden. De voorspelde sales aantallen zullen weergegeven moeten worden. Indien een lege grafiek wordt weergegeven na het laden van de data, is de data niet juist ingeladen.

Als alle data is ingeladen, kan Stap 2 worden uitgevoerd.

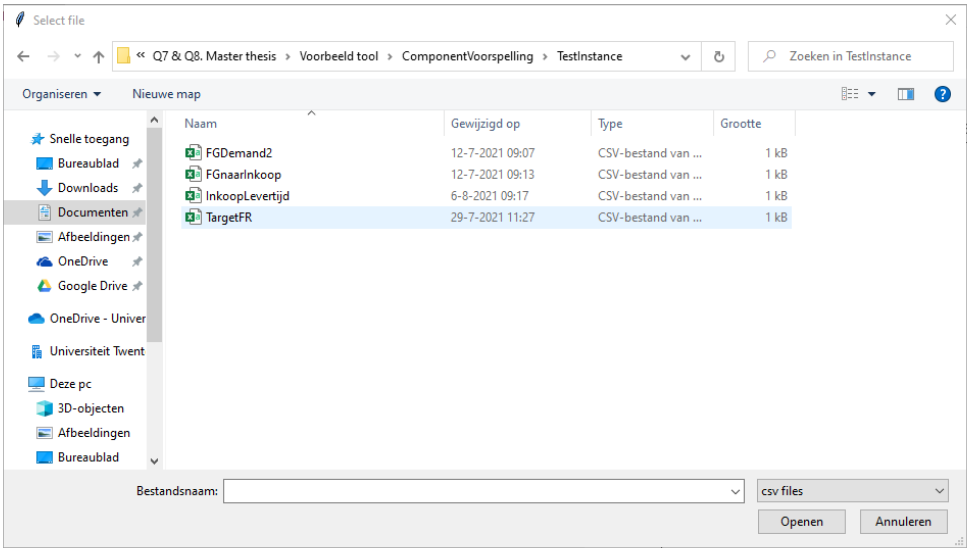

Gebruikersinterface

Stap 2: Bereken componentvoorspelling

De componentvoorspelling kan uitgevoerd worden met de 'Run!' knop. Als de status 'Finished calculations!' wordt weergegeven, zijn de berekeningen afgerond en kan het tabblad 'Probability matrix' geopend worden. In dit tabblad kan de kans dat een component in een productgroep gebruikt wordt gecontroleerd worden. Ook kan de vraag per component bekeken worden. Als beide niet beschikbaar worden, voldoet één van de input bestanden niet aan de beschreven vereisten en zal dit aangepast moet worden voordat de tool opnieuw gedraaid kan worden.

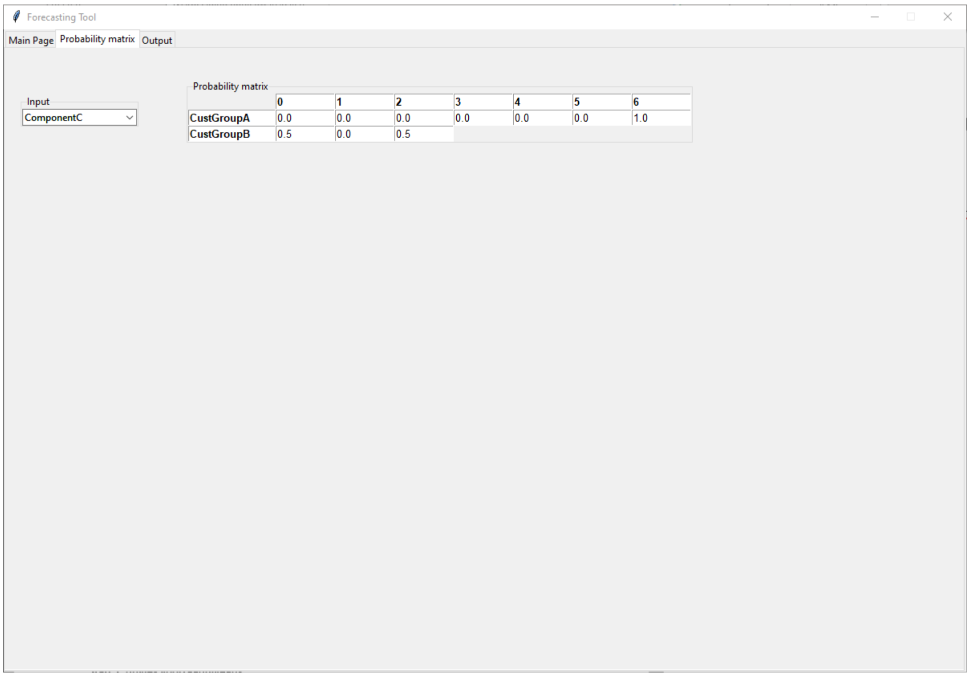

Stap 3: Advies voorraadniveaus

Na het berekenen van de componentvoorspelling, kan de tool een advies voor inrichting van de minimum voorraad geven. Dit advies is gebaseerd op de berekende componentvoorspelling en de ingevoerde gewenste fill rates. Een aandachtspunt is dat het model geen schaalvoordelen (zoals MOQ en bestellen in veelvouden) meeneemt. Het kan dus voorkomen dat een voorgestelde hoeveelheid onder de MOQ van een leverancier valt. Na het runnen zal de status 'Finished greedy heuristic, base stock levels found!' weergegeven worden en wordt het tabblad ‘Output’ beschikbaar. Hier wordt de verwachte fill rate gegeven en worden de componenten met de grootste verschillen tussen de huidige en geadviseerde voorspelling weergegeven.

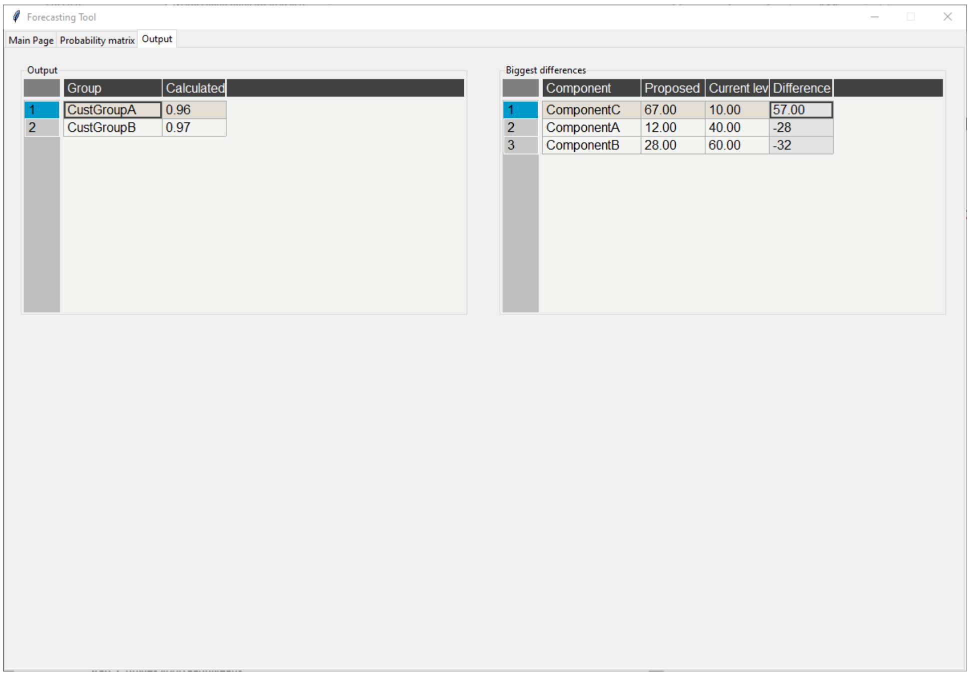

Stap 4: Sla uitkomsten op

Na het uitvoeren van alle stappen kunnen de resultaten naar Excel worden geëxporteerd. Dit Excel bestand bevat de volgende tabbladen:

- Sheet 1: ’*Expected component demand*’ bevat de verwachte componentvraag die is berekend op basis van alle input.
- Sheet 2: ’*Proposed base-stock levels*’ bevat de voorgestelde minimumvoorraden op basis van de voorspelde componentvraag.
- Sheet 3: ’*Expected component fill rates*’ geeft de verwachte component fill rates weer. Deze fill rates zijn bepaald op basis van de voorspelde componentvraag, waardoor de verwerkelijke fill rates kunnen afwijken.
- Sheet 4: ’*Expected FG fill rates*’ geeft de verwachte fill rates op eindproductniveau weer. Deze fill rates zijn bepaald op basis van de historische samenstelling en de component fill rates. Aangezien beide onzeker zijn, zijn deze fill rates ook een verwachte waarde en kunnen deze in werkelijkheid afwijken.

REFERENCES

Agrawal, N., & Cohen, M. A. (2001). Optimal material control in an assembly system with component commonality. *Naval Research Logistics (NRL)*, *48*(5), 409–429.

Axsäter, S. (2006). *Multi-echelon systems: Structures and ordering policies.* Boston, MA: Springer US. Retrieved from https://doi.org/10.1007/0-387-33331-2_8 doi: 10.1007/0-387-33331-2_8

Bogner, E., Löwen, U., & Franke, J. (2017). Systematic consideration of value chains with respect to the timing of individualization. *Procedia CIRP*, *60*, 368-373. Retrieved from https://www.sciencedirect.com/science/article/pii/S2212827117300161 (Complex Systems Engineering and Development Proceedings of the 27th CIRP Design Conference Cranfield University, UK 10th – 12th May 2017) doi: https://doi.org/10.1016/j.procir.2017.01.015

Cheng, F., Ettl, M., Lin, G., & Yao, D. D. (2002). Inventory-service optimization in configure-to-order systems. *Manufacturing & Service Operations Management*, *4*(2), 114–132.

Chen-Ritzo, C.-H., Ervolina, T., Harrison, T. P., & Gupta, B. (2010). Sales and operations planning in systems with order configuration uncertainty. *European journal of operational research*, *205*(3), 604–614.

Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Operational Research Quarterly (1970-1977)*, *23*(3), 289–303.

Dolgui, A., & Prodhon, C. (2007, 12). Supply planning under uncertainties in mrp environments: A state of the art. *Annual Reviews in Control*, *31*, 269-279. doi: 10.1016/j.arcontrol.2007.02.007

Fildes, R., & Goodwin, P. (2007). Against your better judgment? how organizations can improve their use of management judgment in forecasting. *Interfaces*, *37*(6), 570–576. Retrieved from http://www.jstor.org/stable/20141547 doi: 10.1287/inte.1070.0309

Fildes, R., Goodwin, P., Lawrence, M., & Nikolopoulos, K. (2009). Effective forecasting and judgmental adjustments: an empirical evaluation and strategies for improvement in supply-chain planning. *International Journal of Forecasting*, *25*(1), 3-23. Retrieved from https://www.sciencedirect.com/science/article/pii/S0169207008001362 doi: https://doi.org/10.1016/j.ijforecast.2008.11.010

Hautaniemi, P., & Pirttilä, T. (1999). The choice of replenishment policies in an mrp environment. *International Journal of Production Economics*, *59*(1), 85-92. Retrieved from https://www.sciencedirect.com/science/article/pii/S0925527398000267 doi: https://doi.org/10.1016/S0925-5273(98)00026-7

Heerkens, J., & van Winden, A. (2012). *Geen probleem, een aanpak voor alle bedrijfskundige vragen en mysteries.* Business School Nederland.

Ho, C.-J., & Ireland, T. (1998). Correlating mrp system nervousness with forecast errors. *International Journal of Production Research*, *36*(8), 2285-2299. Retrieved from https://doi.org/10.1080/002075498192904 doi: 10.1080/002075498192904

Holt, C. C. (1957). *Forecasting seasonals and trends by exponentially weighted moving averages.* Defense Technical Information Center.

Jacobs, F. R., Berry, W. L., Whybark, D. C., & Vollmann, T. E. (2011). *Manufacturing planning and control for supply chain management: Apics/cpim certification edition.* McGraw-Hill Education.

Pan Oston B.V. (2020). *Over ons: Het verhaal van pan oston.* Retrieved from https://www.panoston.nl/about-us

Silver, E., Pyke, D., & Thomas, D. (2016). *Inventory and production management in supply chains.* Taylor & Francis.

Song, J.-S. (1998). On the order fill rate in a multi-item, base-stock inventory system. *Operations Research*, *46* (6), 831–845.

Srinivasan, R., Jayaraman, R., Rappold, J., Roundy, R., & Tayur, S. (1998). *Procurement of common components in a stochastic environment* (Tech. Rep.). Cornell University Operations Research and Industrial Engineering.

Steele, D. C. (1975). The nervous mrp system: how to do battle. *Production and Inventory Management*, *16* (4), 83–89.

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. Retrieved from https://doi.org/10.1057/palgrave.jors.2601841 doi: 10.1057/palgrave.jors.2601841

van den Bogaert, J. (2017). Buyer and vendor managed inventory optimization in configure-to-order systems. *Eindhoven University of Technology*.

Winters, P. R. (1960). Forecasting sales by exponentially weighted moving averages. *Management Science*, *6* (3), 324–342. Retrieved from http://www.jstor.org/stable/2627346

Zhao, X., & Lee, T. (1993). Freezing the master production schedule for material requirements planning systems under demand uncertainty. *Journal of Operations Management*, *11* (2), 185-205. Retrieved from https://www.sciencedirect.com/science/article/pii/027269639390022H doi: https://doi.org/10.1016/0272-6963(93)90022-H