Improving the inventory management of parts using an order-based performance approach at Fluiten

Master Thesis Assignment

Emma A. Vlaswinkel

University of Twente
Master Industrial Engineering and Management
March, 2024

Master thesis Industrial Engineering and Management

Improving the inventory management of parts using an order-based performance approach at Fluiten

Author:
E.A. Vlaswinkel (Emma)

University of Twente
Drienerlolaan 5
7522 NB Enschede
Netherlands

Fluiten Italia S.p.A.
Via Leonardo Da Vinci, 14
20016 Pero MI
Italy

Supervisors University of Twente
Dr. M. C. van der Heijden (Matthieu)
Dr. D. R. J. Prak (Dennis)

Supervisor Fluiten
Ing. R. Bollina (Raffaella)

Preface

Dear reader,

This master's thesis, 'Improving the inventory management of parts using an order-based performance approach at Fluiten', marks the end of my studies Industrial Engineering and Management at the University of Twente. I am very grateful for the fantastic years I have enjoyed so much in both Enschede and Turin.

First of all, I would like to thank my first supervisor, Matthieu van der Heijden. He agreed to guide me remotely which allowed me to perform my research abroad. He always helped me on short notice with extensive feedback and suggestions. He also stressed the importance of the order-based approach, which has improved the thesis. I would also like to thank my second supervisor, Dennis Prak, for taking the time to read the report and providing a fresh perspective.

I also want to give special thanks to Fulvio Colombo and Raffella Bollina for their trust, which allowed me to conduct my thesis with Fluiten. They have always supported me with close guidance and were open to any question. Additionally, I want to say *grazie mille* to all my colleagues at Fluiten for the Italian language and culture classes and for giving me a warm welcome in a foreign country.

Lastly, I want to thank my family and friends for supporting me during the research and my entire study period. Thank you to Hot for a close home, Fatale for all the Tuesday nights, Lieke for our tennis matches, and the girls for teaching me Dalmuti. I especially thank Rozan for making sure I never felt alone while writing this thesis and Wouter for being my home in Milan.

I hope you enjoy reading this master's thesis!

Emma Vlaswinkel

Milan, March 2024

Management summary

This research has been performed at Fluiten SpA in Pero, Italy. Fluiten designs and manufactures mechanical seals, rotary joins, and auxiliary systems for various purposes. Their products come in many different designs and sizes and consist of many parts. They currently experience a wrong inventory configuration of parts, where they have too much inventory of parts they do not need and too little of the ones they do need. A significant cause is that it is unknown when new parts will be added to the inventory. Therefore, the core problem that is solved is that 'currently no inventory model is applied to provide knowledge on when and how many parts to buy or produce'. This research aimed to propose a management model to provide this knowledge. Therefore, the main research question addressed in this thesis is formulated as follows:

'How do we design Fluiten's inventory management to maximise order fill rate and minimise the inventory costs?'

For the standard orders in the scope of the research, the production-inventory model is identified as Assemble-To-Order. The 19,610 Stock Keeping Units (SKUs) in this scope are analysed. The Distribution By Value showed that only 7% of the SKUs already represent 80% of the annual usage of 2022. Many parts have too high inventory levels (50% of the SKUs have an inventory coverage of over a year) due to a lack of insight into current stock and a risk-averse policy. The current order policies are highly inefficient and rely heavily on experts' opinions. Demand is currently managed entirely reactively. The current performance was quantified through the analysis of backorders, which showed especially bad performance for assembly orders (an order fill rate of only 55%).

Continuous review policies from literature are considered as the purchasing and production offices analyse inventory levels every other day, a neglectable review period. A promising stepwise approach was presented to classify the SKUs. Two methods for setting individual target fill rates using an order-based approach are considered. The method by Teunter et al. (2017) includes the criticality of a SKU, in this research expressed as the number of orders a SKU is part of, while the method by van der Heijden (2024) uses the Bill Of Materials to translate orders into demand for individual parts. The calculations of the policy parameters are described using the relevant Normal, Gamma and Negative Binomial distributions.

Using the theory, a tailored solution using classification is designed. With different characteristics of the SKUs, demand patterns, expected lead time demand and coefficient of variation in lead time demand, a stepwise approach has been created, enabling a straightforward step from historical demand to demand distributions. Although the demand for individual parts that make up one end-product is not independent, this has been assumed in this research for simplicity.

Using historical data up to and including 2022 in the proposed inventory model, it could be tested on actual demand in 2023 with a simulation. The simulation, created in Python, works with a time unit of a day. Available parts are reserved for incoming orders prioritised based on the upcoming due date. Among others, the output Key Performance Indicators (KPIs) are the average fill rate of the individual parts, the average on-hand inventory value and the order fill rate. The two methods to set individual target fill rates found in the literature have been tested in the simulation on a small scale. The method by Teunter et al. (2017) resulted in an OFR of 85.4%, 1.3% higher than the outcome of the method by van der Heijden (2024), with considerably less effort. This first method has therefore been applied to all parts. The parts currently not kept in inventory have been re-evaluated, which showed that 63 should have a safety inventory to avoid unnecessary costs. The proposed solution showed an increase in the OFR of 17.4% to 78.7% while reducing the average on-hand inventory value by €808,243 to €2,294,966.

A sensitivity analysis tested the proposed solution's robustness. Leaving out the undershoot reduces both the average on-hand inventory value by €319,933 considerably as well as the number

of orders by 10,064, at the costs of a diminished OFR by 3.5%. Applying an (s,S)-policy for less important items results in a slightly increased average on-hand inventory value of €16,535 while increasing the OFR by 1.6%. Four different lower bounds for the individual TFRs show the influence of this parameter on the final performance of the models and serve as input for the management decision on the chosen value. Any required extra Supplier Lead Time significantly increases the required inventory level and highlights the importance of reducing lead times when possible. The findings were combined and discussed with management to formulate a final proposal which balances costs an performance, where undershoot is applied only to the class of A-items with Normally distributed demand, an (s,S)-policy is used for less important items and a lower bound of 85% is applied. An implementation plan was established to ensure that the changes in inventory management proposed in this research are properly embedded in the company.

Using the results from the simulations and insights from the research, we list the following main conclusions for Fluiten:

1. Using the created inventory management tool prototype, we can construct an inventory policy for each SKU. Testing in the simulation showed that the KPIs have improved significantly. Costs are reduced, and the performance towards customers has improved due to a reduction of backorders. The delay of the backorders, of which the number has been reduced by 1,753, has increased by 5.1 days, mostly due to unfair comparison as there is additional flexibility in reality.

KPI (parts)	**Change**	**Obtained value**
Average Fill Rate over SKUs	6.8%	93.0%
Volume Fill Rate	5.0%	95.6%
Target Fill Rate met	17.4%	80.3%
Average On-Hand Inventory value	€-1,015,789	€2,087,420
Number of orders placed	*no comparison possible*	23,893

KPI (orders)	**Change**	**Obtained value**
Order Fill Rate	17.6%	79.0%
Number of parts late	-0.9	1.6
Delay of backorders (days)	5.1	18.6

KPI (labour costs)	**Change**	**Obtained value**
Labour costs of backorders	€-30,484	€19,668

2. 63 selected SKUs should be brought into inventory as holding costs are less than costs of backordering.

Based on the research performed and stated conclusions, recommendations are presented to Fluiten. The main recommendations are the following:

1. Implement recommended inventory control policies using the proposed implementation approach.

   The provided inventory control parameters for SKUs in Classes 2 to 6 should be implemented in the ERP system and methods of ordering. The inventory of parts placed in Class 1 should be removed, and these SKUs should be managed on MRP basis.

2. Improve data quality.

The cost value and the supplier lead times of the SKUs in the ERP system should be made accurate for all parts. The SKU-codes that have been merged or transferred should be connected properly to avoid data loss, after which the old codes should be removed. Some SKUs are used internally and are taken from inventory in large batches, resembling intermittent demand. These internal movements should be separated from customer demand.

3. Implement the created performance KPIs.

   When possible, the chosen KPIs have been determined on historical data to determine the current performance. They provide valuable insight into the performance, which was not available before. Implementing the KPIs enables constant insight into all (future) available data.

Contents

Readers guide

In eight chapters, the research performed at Fluiten is described. We shortly introduce the chapters.

Chapter 1: Introduction
Chapter one introduces the research and describes the company. Using a problem cluster, the observed action problem has been traced back to the core problem. The research design is explained by setting the scope and listing the research questions and deliverables.

Chapter 2: Current Situation
The second chapter analyses the current situation. The production-inventory model is explained, as well as the Stock Keeping Units (SKUs) in inventory with their characteristics. The current order policies are described, and demand and supply are analysed. Calculations regarding back-orders quantify the performance of the current situation.

Chapter 3: Literature study
For chapter three, literature is consulted to find an inventory management model that could be applied to Fluiten. The chapter starts with a classification of inventories at Fluiten according to literature. Then, different methods of SKU classification and inventory policies are described. The different parameters and their formulas using an order-based performance method are given. Last, different demand distributions with which lead time demand can be estimated are described.

Chapter 4: Solution design
The fourth chapter introduces the design of the solution. A tailored approach to classify SKUs is presented and applied. The classes are matched to inventory policies and their corresponding parameters. Constraints to the model are listed.

Chapter 5: Results analysis
In chapter five, the proposed solution's performance is tested using a simulation. Two methods for selecting individual target fill rates using an order-based approach are implemented on a small scale and compared. Parts currently not kept in inventory are evaluated to decide whether they should be. The final results of the proposed solution are presented.

Chapter 6: Sensitivity analysis
Different modifications to the proposed solution are tested to evaluate the performance of the model under different scenarios. Results provide more insight into the robustness of the tool and the best configuration.

Chapter 7: Implementation
This chapter presents a stepwise approach based on literature to guide the implementation of the proposed changes into the companies' practices.

Chapter 8: Conclusions, recommendations, and future research
The final chapter lists the conclusions and recommendations from the research. The practical and scientific contributions are discussed. Lastly, the limitations of the research are described, as well as possible future research.

Acronyms

AHP Analytical Hierarchy Process.

BOM Bill Of Materials.

CODP Customer Order Decoupling Point.

DBV Distribution By Value.

EOQ Economic Order Quantity.

ERP Enterprise Resource Planning.

GNB Generalised Negative Binomial.

IC Inventory Coverage.

IP Inventory Position.

KPI Key Performance Indicator.

KPIs Key Performance Indicators.

MOQ Minimum Order Quantity.

MRP Material Requirements Planning.

NB Negative Binomial.

OFR Order Fill Rate.

SKU Stock Keeping Unit.

SKUs Stock Keeping Units.

SLT Supplier Lead Time.

TBS Time Between Stockout occasions.

TFR Target Fill Rate.

TFRs Target Fill Rates.

VFR Volume Fill Rate.

1 Introduction

While completing my master's degree in Industrial Engineering and Management, I performed research for my thesis at Fluiten SpA. This research aimed to design a solution that improves their inventory management. This chapter presents the company in Section 1.1 and subsequently identifies the problem we study in Section 1.2. This research design is described in Section 1.3.

1.1 Company introduction

Fluiten is an Italian company that designs and manufactures mechanical seals, rotary joints and auxiliary systems for various purposes. It was founded in 1962 by Alberto Delfo Colombo, after which it started manufacturing the first pumps and mixers. In 1981, it moved its production headquarters to its current location in Pero, Milan. In the following ten years, the plant kept expanding to the current size of over 10,000 square meters. Fluiten's products are manufactured here by over a hundred employees (Fluiten, 2023).

A mechanical seal, the company's main product, contains fluid in a vessel where a rotating shaft passes through a stationary housing. An example would be if you had a can with tomato sauce (see Figure 1) which you want to mix. You would have to create a hole in the wall of the can to put the rotating arm through; in this example, it might be a spoon. This spoon should be able to turn without causing any tomato sauce to leak. In this simple example, a simple rubber ring would suffice. However, a better solution is required when the spoon starts spinning quickly and at a higher temperature. The mechanical seals enable this. The company sells its mechanical seals to clients worldwide, both original equipment manufacturers and end users. They work for the petrochemical, chemical, pharmaceutical, food, energy, water and shipbuilding industries.

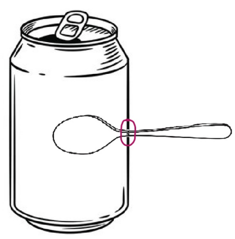

Figure 1: Tomato can and spoon

Figure 2: Type SA mechanical seal

The seals come in many different designs and sizes (see Figure 2 for an example). Each seal consists of 20 parts on average. In addition to the standard existing seals, they also design and manufacture customised seals. Fluiten has a commercial partnership with clients in over 40 countries, offering a highly qualified after-sales service. This means that regardless of the product type or the date of purchase, customers can request a replacement, adjustment, or repair of their products.

Fluiten is also part of the joint-venture Johnson-Fluiten. Together, they design, manufacture, and market rotary joints for water, oil, air, and other fluid applications. The heart of the rotary joint is a mechanical seal from Fluiten. Kadant Johnson has several locations worldwide and contributes by providing a worldwide sales network (Johnson-Fluiten, 2023).

1.2 Problem identification

In this section, the experienced action problem is described (Section 1.2.1). Using the visualised problem cluster, the core problem could be identified in Section 1.2.2.

1.2.1 Action problem

The sales team receives an order arriving from a customer. This might be an order for a standard existing product, in which case they simply put it in the Enterprise Resource Planning (ERP), and it is processed further. The order might also be for a fully tailored product, which first requires new drawings. Once these drawings are made, the tailored parts are produced and sent to the warehouse for temporary storage. The picking team is responsible for collecting all the parts from the warehouse for each order, standard or tailored. These will then go to the assembly team, which assembles and tests the seals. If everything works correctly, the product(s) of the order go to shipping, where the order is packed and prepared.

Fluiten manufactures most of the parts itself. Raw materials arrive and are manufactured into parts by the workshop. However, some of the raw materials go to sub-suppliers for manufacturing when Fluiten does not have the capacity, or outsourcing is the less expensive option. Other parts are never produced internally and are purchased from parts-suppliers. These parts, either produced or purchased, all go into the warehouse.

However, the inventory level configuration is currently not what it should be. Many parts have been in the warehouse for years with no future purpose in sight. Fluiten does have to pay government taxes on them, even though they are of no value to the company. They also take up space in the warehouse. On the other hand, Fluiten is currently short on relevant parts. Suppose an order arrives that requires materials which are not in stock. Purchasing and production offices will aim to minimise the delay, and the sales department has to communicate bad news to the customer.

So, in short, both having too much and too little inventory results in costs. There is too much inventory of the parts they do not need and too little of the ones they do need. The action problem is, therefore, a *wrong inventory configuration of parts*. This research will thus focus on designing an inventory management model with which the inventory levels can be adjusted to match Fluiten's needs.

1.2.2 Problem cluster and core problem

The observed action problem is thus that Fluiten has the wrong inventory configuration for parts. With the current way of managing the inventory, too many stockout occasions occur, while at the same time, high taxes are paid for dead stock. In this section, the process from action problem to core problem is described.

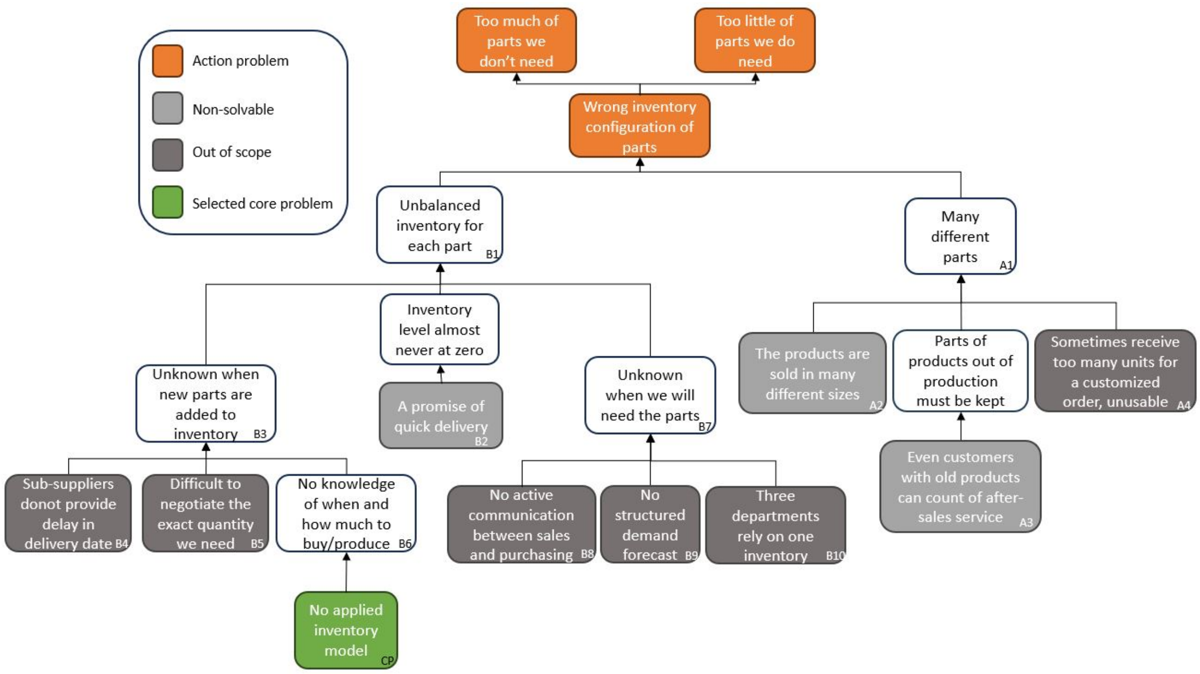

Figure 3: Problem cluster

The problem cluster, constructed after consulting the purchasing and production managers, is visualised in Figure 3. The described causes and reasons are referred to with a letter and a number corresponding to the figures' rectangles. The two underlying main reasons for the wrong inventory level are the high number of different parts in circulation and, therefore, in inventory (A1) and the unbalanced inventory for all these parts (B1).

Many different parts
There are three reasons why there are so many parts in inventory.

1. The first reason (A2) is that they sell about 50 different types of products, composed of about 20 parts on average, in many different sizes. This means that their parts are also produced in many different sizes, resulting in a wide variety of parts (about 30,000), of which the inventory level has to be managed. This is not a problem to solve now, as it is a part of Fluiten's competitive advantage. It is a constraint to the model.

2. Secondly (A3), not only the parts of the products currently in production must be kept, but also parts of older models, as previous customers can always count on their after-sales service. This is also a part of their competitive advantage and an organizational decision.

3. However, the third reason (A4) for the high number of parts is a problem. Customers can also reach out to Fluiten for a fully tailored and customised seal, resulting in designs of parts that are solely useful for them. If Fluiten lets a sub-supplier produce a part for them, they usually also send their raw materials. The sub-supplier then manufactures the maximum quantity of this part they can get from the raw material, even if Fluiten requested only a smaller number. The extra customised parts go into inventory, where they might stay for years, as they are not part of a 'standard' product. This increases the number of parts in inventory, as they usually would not have inventory for these customised pieces. Fluiten recognises this problem, and they are currently working on solving it. It is not a problem in the scope of this research, as tailored parts will be left out of scope (1.3.1).

Unbalanced inventory for each part
Going back to the starting two reasons for the wrong inventory level and analysing the problem of having an unbalanced inventory for every part (B1), we find three underlying reasons.

1. The first reason (B2) for the unbalanced inventory for all parts is that Fluiten wants to promise quick and after-sales service. For standard orders, they want to deliver within five days; for special customers, even within 1 or 2 days. Their after-care entails that customers can return their broken mechanical seal and receive an offer on the options. This might be a repair by replacing a broken part or a recommendation for ordering a new seal. They aim to provide any of these services in 5-7 days and 1-2 days for special customers. As the lead times of their parts are longer than this, the only way to produce and provide service so quickly is to have the parts already in inventory. This is why there is now at least some inventory for almost all parts. However, the exact quantity to have in stock to handle this quick demand has never been determined properly, resulting in a somewhat random inventory balance. The inventory level is now increased just so they do not have zero units.

2. The second reason (B3) is that it is not clear enough when new parts are added to the inventory. This has three causes.

- The first one (B4) is that the sub-supplier's delivery date is planned with Fluiten but not respected nor communicated properly. The company has recognised this problem, and one of the employees is now working on improving it. Managing the relationship is not the purpose of the research, but again, the effect of it, namely the lead times, does serve as input. The actual observed lead times will be used rather than the communicated ones.

- The second cause (B5) is that although the parts-suppliers are more reliable, respecting the planned delivery date about 80 per cent of the time, one of them has minimum ordering quantities, which are often not ideal. If possible, they might then opt for another, more expensive supplier, who does not have this Minimum Order Quantity (MOQ) if the quantity is very unattractive. Solving this relationship regarding the minimum ordering quantity is not part of this research. In the case that the MOQ is reasonable, it will be used; if not, the optimal batch is calculated, which can be used to go into negotiations.
- The third cause (B6) can be summarised as a lack of knowledge of when and how many parts to purchase from the parts-suppliers and sub-suppliers, or to produce, in the workshop. The new orders are checked by the Material Requirements Planning (MRP) every other day. If the inventory level of one of the parts in those orders surpasses the safety stock level, it proposes to order or produce, depending on the type of part. Whether they actually order or produce and which quantity is decided on at that moment. The dedicated employee comes up with a theoretically optimal batch defined from the trend of the volume of the last seven years, intuition and expertise. The lead time of the parts is not taken into real consideration. With this approach, no stability can exist regarding restocking dates.

3. The third reason (B7) for the unbalanced inventory is that it is unclear when they will need the parts. This reason has three underlying causes as well.

- The first one (B8) is that there is no communication between sales and purchasing. If there is an incoming order and there is no direct problem due to a lack of parts, purchasing will not be aware of the reduced inventory unless they actively scan all orders. Improving this could reduce the number of stockouts. However, this is not part of this research for now, as we first focus on finding the right balance for the inventory. This communication could be implemented in the future.
- The second cause of unknown demand (B9) is the lack of a demand forecast. They work reactively, ordering and producing only when the inventory level is below the safety stock. The quantity of orders and productions is roughly based on historical data combined with intuition and expertise, instead of having a quantified analysis or some structured forecast. This increases the risk of stockout and reduces the control over the inventory. Although this thesis does not aim to optimise the forecast, some form of demand forecast for the parts is required to improve the inventory. A simple forecasting model will be used to forecast the demand for the parts. This way, we can obtain data for the inventory model to work with. If the demand forecast is improved in the future, so will the inventory model.
- The third cause of unknown demand (B10) is that three different departments within the company rely on one inventory, production, assembly and after-sales department. The picking department collects all materials for the three types of workorders and gives the parts to the right department. Currently, there is little insight into these three sources of demand. For now, we consider these demand streams as one, as this is also how it is structured in the MRP. We add the combined incoming demand as a constraint to our model.

Selected core problem
To summarise, reducing the high number of parts in stock will not be taken into consideration, as it is due to either Fluiten's competitive advantage or an out-of-scope problem. The promised service times, resulting in a need for inventory due to long production times, will not be altered.

The uncertainty of demand due to a lack of communication, lack of demand forecast and combined sources will not be solved. A basic demand forecast will be set up to approximate this uncertainty. The uncertainty in the sub-suppliers' delivery date will not be improved, but actual historical lead times will be used as input. If the MOQ for a certain part is reasonable for Fluiten, it will be used; if not, the optimal batch is calculated, which can be used to go into negotiations.

The core problem selected in this research is the missing knowledge of when and how much to purchase or produce. Fluiten starts ordering and producing when the inventory level surpasses the safety stock, but this number has never been appropriately determined. The same goes for the quantity to order or produce, now mainly based on a broad overview of historical data and expert opinion. With knowledge of proper parameters and an approach to the inventory, the stock availability of the parts will improve, as well as the robustness of the purchasing and producing department. The core problem for this thesis is thus that;

'currently no inventory model is applied to provide knowledge on when and how many parts to buy or produce'

1.3 Research design

In this section, the general design of the research is described. This research is structured using the Managerial Problem-Solving Method by Heerkens and van Winden (2017). This method is suitable for this action problem and creates a logical sequence through the research. In Section 1.3.1 the scope is set, making sure that the research is feasible within the time period while still being relevant. The main research question, sub-questions, and corresponding approach are defined in Section 1.3.2. Lastly, the intended deliverables are listed in Section 1.3.3.

1.3.1 Scope

As described in Section 1.2.2, five observed problems are not a part of this research. Improving the relationship with the sub- and parts-suppliers, regarding unknown delivery dates and quantities, will not be taken into consideration. The current three parties who rely on the inventory keep doing this, and no communication system between sales and purchasing will be set up. A basic demand forecast on parts-level is required to serve as input in the inventory model, but creating a high-quality forecast is not the aim.

Within the topic of the core problem itself, an inventory model, there are some boundaries set to ensure that the research can be conducted within the time limit of one academic semester. Regarding purchasing, the research is focused on ordering parts, not the raw materials for the workshop as these are hardly kept in inventory. The input streams taken into account are the workshop, parts suppliers and sub-suppliers. The outgoing streams are the production and after-sales departments.

As the extremely high diversity of parts is a core part of the problem for Fluiten, all spare parts (about 31,000) are initially considered. These are the parts that the existing products are composed of. Both the produced and purchased parts are part of the scope. Some groups of parts/inventory are left out of the scope for various reasons. A more detailed scope analysis will be done in Section 2.2.1. The data used for analysis will be taken from the Enterprise Resource Planning (ERP) system, which has been used since June 1st 2016.

1.3.2 Research questions

From the chosen core problem, the main research question is defined:

'How do we design Fluiten's inventory management to maximise order fill rate and minimise the inventory costs?'

Answering this question will solve the core problem, stating that no such management model exists. The model will help inform when and how much to order or produce. The question sets the shape and direction of the thesis. Sub-questions were created to guide the research process following the chosen approach. They also establish the outline of the thesis. For each question, the motivation and main steps are given.

1. What does the current inventory situation look like?
As our action problem is that our inventory currently has the wrong configuration of parts, the goal of the first phase is to obtain detailed insight into the current situation. The general process of an order and its relation with production and inventory are analysed. This is done by talking with the stakeholders and analysing their handover documents. Data is consulted for a quantitative analysis after this qualitative analysis to understand the existing processes. These questions are answered in Chapter 2 Current Situation.

1.1. What are the process flow and lead time of a customer order, and which production-inventory model is used?
1.2. Which parts are currently stored in inventory?
1.3. What is the current inventory control policy used?
1.4. Which types of demand occur, and how can the demand of parts be forecasted?
1.5. What are the current supply- and production lead times of purchased and produced parts, and how accurate are they?
1.6. What is the performance of the current inventory management regarding backorders, and what are the causes?

2. What inventory management methods are proposed in literature that could be applied to Fluiten, with which the right inventory levels can be chosen?
The second research question has the purpose of solution generation. We did so by consulting literature. With literature research, we found various inventory management models with different characteristics. A way to distinguish a division between the large number Stock Keeping Units (SKUs) and various corresponding policies and parameters are described. This can be found in Chapter 3 Inventory management: A literature study.

2.1 What inventory management found in literature can be applied to the inventory management at Fluiten?
2.2 How can the SKUs be classified according to literature?
2.3 What inventory control policies are available in literature, and how should the corresponding parameters be defined using an order-based performance approach?
2.4 What simple demand forecasting method can be used to estimate the demand for parts?

3. What inventory management methods are most applicable for the SKUs and how should the inventory management tool be designed?
We have now chosen a solution. Based on the information from the literature and discussions with stakeholders, the SKUs could be classified and matched to an inventory policy. Each policy had its parameters to be calculated. Finally, based on the chosen policies and parameters, together with stakeholders a decision had to be made on the tool's design. This is described in Chapter 4 Solution design.

3.1 How should the SKUs be classified?
3.2 What inventory policy is suitable for each type of classification?

3.3 What should the parameters of these policies be?

4. What is the performance of the proposed inventory management tool?
The chosen inventory policies, parameters and prototype tool had to be tested on their performance. First, Key Performance Indicators (KPIs), collected from stakeholders and literature, were defined, and a test simulation was created. After ensuring the test was valid and verifiable, the results of the tool from the simulation could be compared to the current management. These results can be found in Chapter 5 Results analysis. The tool's robustness is tested using sensitivity analysis to provide insight into how it reacts to discrepancies and relaxations of constraints, as described in Chapter 6 Sensitivity analysis.

4.1 How should the performance of the proposed inventory management tool be tested?
4.2 Are the results of the test valid and verifiable?
4.3 Does the proposed inventory management tool result in an improvement compared with the current management?
4.4 How robust is the tool to discrepancies in input settings and relaxations of constraints?

5. How can the proposed inventory management tool be implemented in practice?
After creating the prototype tool, an implementation plan tailored to Fluiten's situation is written in Chapter 7 Implementation to guide the implementation into the current systems and way of working. After all, if the final implementation is not guided and communicated well, the tool has little value.

6. What conclusions and recommendations can be made from conducting this thesis at Fluiten?
We finished the research by concluding, formulating recommendations and reflecting on the research in Chapter 8 Conclusions, recommendations and future research.

6.1 What are the main conclusions from conducting the thesis?
6.2 What are the main recommendations from conducting the thesis?
6.3 What are the theoretical and practical contributions of the research?
6.4 Which limitations and areas for further investigation can be related to the research?

1.3.3 Deliverables

The main research question is answered upon completion of this research. The following deliverables were be presented:

- A prototype for a reusable inventory management tool that enhances decision-making by designing an inventory policy with founded parameters.
- Advice on the appropriate inventory policy for parts with recommended settings for the parts classification, review period, safety stock and order-up-to-level. This is based on the theory, analysis of the current situation and solution design. The performance is determined using a simulation, and the change effect is evaluated via a sensitivity analysis.
- A simulation model in Python with which the performance of an inventory policy can be tested using actual historical orders and demand.
- Recommendations for implementation of the tool and policies and further research.

2 Current Situation

In this chapter, the current situation is described to answer research question *'What does the current inventory situation look like?'*. Section 2.1 presents the main processes, safety stock placement and the customer order lead time. In Section 2.2, the exact SKUs in scope are identified and analysed with a Distribution By Value and inventory coverage analysis. The current ordering policies as well as additional specifications are described in Section 2.3. The demand share of different departments, the demand planning and the demand patterns are presented in Section 2.4. In Section 2.5, the supply with which this demand is filled, from parts-suppliers, sub-suppliers and the workshop, is described. Having gathered all information on the current management, Section 2.6 analyses the current performance and provides a formula with enables future comparison. Section 2.7 concludes the chapter.

2.1 Production-inventory model

In this section, the whole process from an order to a finished product is described (Section 2.1.1), as well as the material flow (Section 2.1.2). These are important to this research, as they indicate how a customer order is transformed into a delivery and incoming goods into a finished product. The safety stock placement is explained in Section 2.1.3. The customer order lead time of end-products is broken down in Section 2.1.4.

2.1.1 Order process

The vast majority of the need for parts comes from arriving customer orders. A flowchart has been created to understand the flow from an order to a finished product. This, with a detailed description, can be found in Appendix A. Improving inventory management could improve the flow of orders, reducing the unwanted buffer between the final departments.

2.1.2 Material flow

In this section, using Figure 4, the way in which the materials move through the company is described. In the following Sections 2.4 and 2.5, each stream is discussed in more detail.

The first incoming stream consists of the raw materials from the raw materials suppliers. The workshop processes these raw materials. Once the parts are produced, they go into inventory. When they arrive at the warehouse department, an employee will scan the necessary barcodes, notifying the MRP which parts can be added to the inventory. If there are orders of which the available parts are already picked and are waiting for completion, the MRP gives a message of this. The remaining parts are stored in the warehouse, and the exact shelf location is entered into the system.

Some parts could be produced by the workshop but are outsourced to sub-suppliers. This can have multiple reasons. For example, a shortage of capacity at Fluiten or an order of such a large quantity that it is worth paying for the fixed costs of externally setting up for a specific part. Again, when they arrive, first possible waiting orders are checked, after which the parts are stored in inventory.

The third and final source of the material consists of the parts purchased from the parts-suppliers. Fluiten never produces these parts, and they are always purchased. Via the waiting orders check, they go to the warehouse.

The biggest 'customer' of the warehouse is the assembly department. The picking team collects the parts for each order. All parts for one order are collected in a box. The boxes with corresponding paper with the order information are then brought to the assembly team. They will assemble the product(s) and perform the required tests.

The second user of the warehouse is the workshop. When producing parts for the final products, they need both raw materials as well as sub-parts from the inventory. The picking team collects the required parts on the production order and brings them to the workshop.

The third user of the warehouse is the after-sales service department. When broken seals from the customer are received, their engineers might be able to fix them using spare parts. These spare parts are collected from the warehouse by the same picking team. All flows are handled on a first-come-first-serve basis.

After assembly and testing, either a new product or a repaired one, the product goes to the shipping department. Here, the product is packed and labelled. As Fluiten is not responsible for the transportation, the due date of orders and thus their delivery performance is based on this moment.

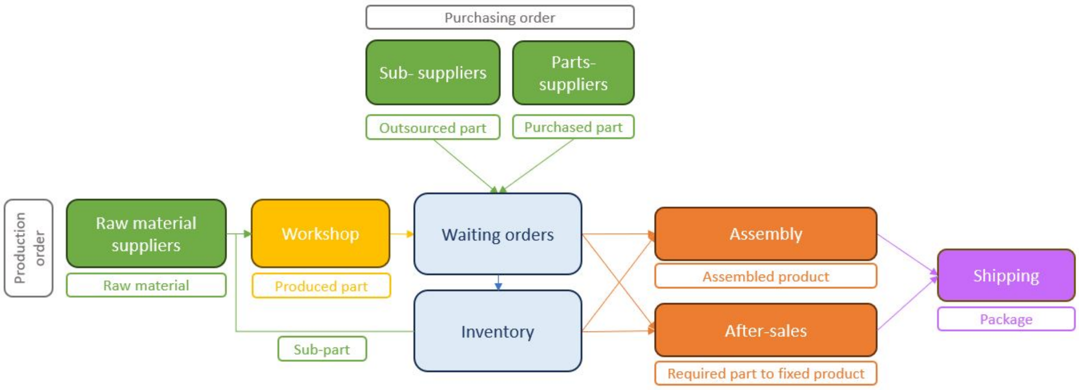

Figure 4: Material flow

2.1.3 Safety stock placement

Before discussing the most suitable safety stock placement, it is important to understand the Customer Order Decoupling Point (CODP) in processes. The CODP is the point in the material flow where the product is tied to a specific customer order. It divides the material flow that is forecast-driven (upstream of the CODP) from the flow that is customer order-driven (downstream of the CODP). The basic manufacturing situations are make-to-stock, assemble-to-order, make-to-order, and engineer-to-order (Olhager, 2010). These four strategies are visualised in Figure 5. Firms with high-volume standardised products are assumed to utilise a level planning strategy, make-to-stock. In contrast, firms with many low-volume customised products are expected to choose a chase planning strategy, make-to-order (Olhager, 2010).

For all standard offered products in the scope of this research, the manufacturing situation can be described as assemble-to-order. Many of the parts can be used for multiple products. In the past, Fluiten has tried to keep some standard final products in stock, thus working in a make-to-stock manner. However, this resulted in products staying in inventory for long, as their demand is hard to forecast. They might keep stock for a particular product in five different sizes, but the ones in inventory are useless if a customer wants the sixth size. Additionally, if there is a slight modification to the design of a product, the ones in stock are immediately obsolete. The assemble-to-order approach is more suitable as the time required for assembly is short compared with the previous steps in the manufacturing and purchasing process. As this has been chosen as the applied approach by the company and it works well for their situation, this CODP is taken as fixed for this research.

2.1.4 Customer Order Lead Time of end-products

The customer order lead time can be defined as the time between when a customer places an order and when the customer receives the product (Kenton, 2023). The steps from an order to a finished product are summarised in Figure 6.

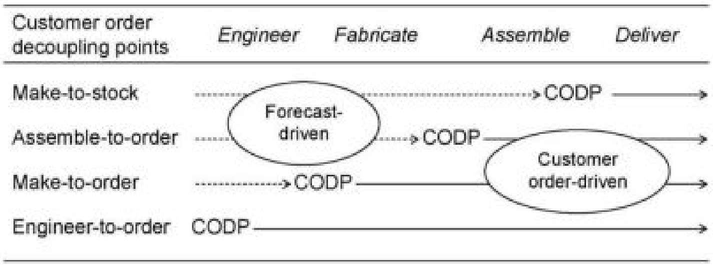

Figure 5: Customer Order Decoupling Point strategies (Olhager, 2010)

The lead time of standard orders depends heavily on the on-hand availability of the parts. The sales office usually checks and adds the order to the ERP system on the same day. If all parts are in stock, only the orange section has to be performed, which could be done in 5 to 10 working days. However, if parts are missing, a replenishment order has to be placed, as shown in the blue section. The picking department will start collecting the parts which are present, storing the incomplete order in the waiting station while waiting for the last ones ('+ delay' in the figure). Incoming replenishment orders are added to the waiting orders with priority based on the closest due date. For initially incomplete orders, the average lead time is 4/5 weeks (7/8 for complex products). Therefore, the parts should be in stock at the arrival of an order. In the customer order lead time aimed for, there is no time for extra sourcing.

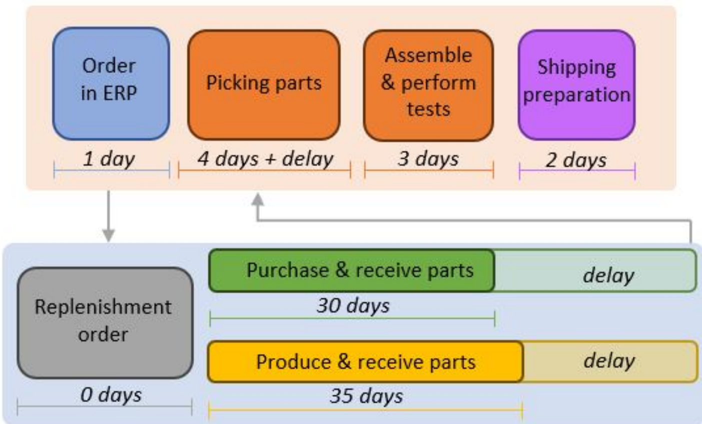

Figure 6: Customer Order Lead Time

2.2 SKUs in inventory

In this section, the SKUs currently in inventory are analysed. The exact selection of parts which are part of the scope (previously defined in Section 1.3.1) is defined in Section 2.2.1. A Distribution By Value is presented in Section 2.2.2, the inventory coverage analysis in Section 2.2.3. Some data cleaning was required, found in Appendix B.

2.2.1 SKUs in scope

To determine the size of the scope, the ERP is analysed. There are 49,160 units which have been either marked as produced, purchased, used or sold in the last seven years in which the ERP has been active. This includes the raw materials and end-products which are out of scope. The inventory that is focused on in this project is the (spare) parts for all products except the auxiliary systems, 31,190 (64%) SKUs. These parts can serve as sub-parts for the workshop, initial parts when assembling a product, and spare parts to sell or use when repairing an existing one. The selection has been visualised in Figure 7.

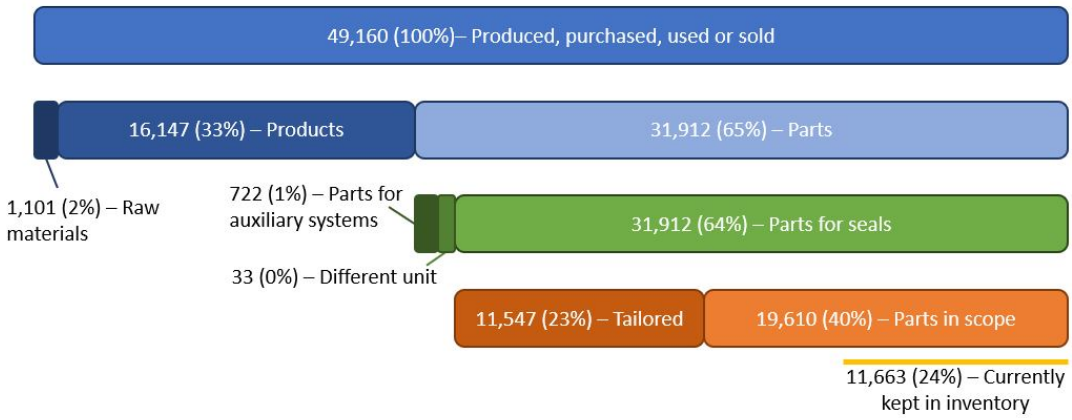

Figure 7: Selection of scope

Of these 31.190 SKUs, 33 are expressed in a unit such as mm/kg/cc/lt/mt rather than the number of parts. As these quantities obviously cannot be compared, they limit the opportunities for analysis. Seeing that of these 33, only five currently have inventory; it is chosen to leave them out of the scope of the main part of the research. This leaves 31,157 SKUs.

As described, Fluiten offers both tailored and standard products and parts. From the 31,157 SKUs saved in the ERP system, 11.818 are tailored parts. Except for a specific selection of parts which are sold so regularly they are considered as standard (271 SKUs), the tailored parts are managed on MRP basis. These tailored parts, which contributed 15,8% of the value in 2022, thus do not require an additional inventory control method and will not be considered. This leaves 19.610 SKUs in scope.

Not all of these SKUs are currently always kept in inventory. 7,947 SKUs have been selected in the past to be managed on MRP basis without keeping permanent inventory. The lead time of an order with any of these parts would undoubtedly be longer. This selection has not been updated; it has not been re-evaluated whether it would be beneficial to keep these SKUs in inventory. They are, therefore, part of the scope, as this research provides a founded recommendation on whether holding inventory would be beneficial for these parts. All other parts have to be kept in stock, as Fluiten wants to be able to provide short delivery times for end-products with these parts.

As mentioned, both the fabricated parts and the parts that are purchased, either from the sub-suppliers or the parts suppliers, are in the scope of this research. Currently, the purchasing office handles 10,038 parts (51%) and 9,572 (49%) parts by the production office. For every production order, consisting of parts which have to be produced, the workshop optimiser checks which will be produced in-house and which will be outsourced. On average, about 50% of the parts will be produced in-house, and the other half will be outsourced.

2.2.2 Distribution By Value

From now on, only the relevant SKUs, as described in Section 2.2.1, are considered. Typically, somewhere in the order of 20% of the SKUs account for 80% of the total annual dollar usage. This suggests that all SKUs in a firm's inventory should not be controlled to the same extent. Creating a Distribution By Value (DBV) is one of the most valuable tools for handling the diversity of disaggregate inventories because it helps to identify the SKUs that are the most important. A

DBV is developed by ranking each unit's annual demand multiplied by the unit's value. The corresponding values of the cumulative percentage of total usage and cumulative percentage of the total number of SKUs are ranked in descending order (Silver et al., 2017, p.28).

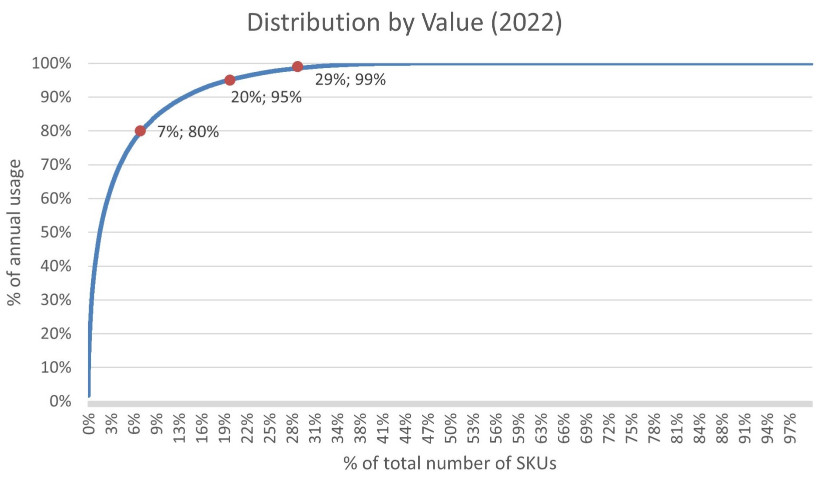

Figure 8: Distribution By Value 2022

Analysing the DBV, visualised in Figure 8, we find that indeed a small share of the SKUs have represented a major share in the usage value of 2022. A full analysis with different years and the optimal part-value calculation can be found in Appendix C. 80% of the usage is reached with only 7% of the SKUs, 20% of the SKUs already represent 95% of the annual usage. Another 28% contributes with the remaining 5%. 49% of the parts did not have any demand this year. This DBV confirms the observed problem of having many parts with many different characteristics and demand patterns. We can use this analysis to make a division between the SKUs and suitably approach them.

2.2.3 Inventory Coverage

As one of the observations leading to our problem cluster was the problem that there are too many parts in storage that are not needed, it is interesting to analyse the Inventory Coverage (IC). This IC represents the expected time in which the current stock level will be depleted (Silver et al., 2017, p.366). It can be computed by Formula 1.

$$IC = \frac{12I}{D} \tag{1}$$

where

IC = inventory coverage, in months

I = on-hand inventory, in units

D = expected usage rate, in units/year

In Appendix D, the complete analysis can be found. The calculations are based only on the SKUs with inventory, which are 6,789 of the total 19,610. The others have no inventory and thus no inventory coverage. Although an average inventory level would be more representative of the inventory throughout the year, at the time of these calculations, only the on-hand inventory of Tuesday the 26th of September 2023 was available. For this general analysis, the expected usage rate has been calculated by taking the average in the years 2017-2023.

As presumed, many parts have a high IC; the median is 10.2 months. 50% of the SKUs are covered for over a year, while 23% of the parts in stock are of SKUs with inventory enough for at least five years. This can be classified as dead stock. Exactly half of the value, however, belongs to SKUs with a more reasonable IC of 3-12 months. There are 274 SKUs with an IC of less than three weeks, which might be too low and increase the stockout risk.

There are two main reasons for these high inventory levels. The first is that there is no, for example yearly, check of the inventory to make sure everything is still relevant. Looking in the ERP, some parts have been entered during the set-up in 2016, and have hardly/not moved during the last seven years. This inventory coverage analysis helped to identify these parts of which the stock level should be reduced. Secondly, they usually keep a high inventory to minimise the chance of missing parts. There is no properly calculated quantified expectation for future orders in the company. In combination with the fact that Fluiten wants to be able to provide quick delivery, but the supply lead time of parts is very long, the only solution was to keep a lot of stock for everything. Some form of demand information and proper parameters reduce these super cautious reserves.

2.3 Current order policies

This section describes the current order policies (Section 2.3.1) as well as additional specifications of ordering characteristics (Section 2.3.2).

2.3.1 Current policies

Fluiten works in an assemble-to-order manner. Their purchasing and production teams base their activities on inventory levels rather than individual orders.

Purchasing
Every other day, the buyer in the purchasing team receives a list of all parts with a current inventory position below the safety stock set in the MRP with their recommended ordering quantities. However, this recommendation is not trustworthy. The parameters were set years ago for most SKUs and do not apply to reality anymore due to changes in demand data, suppliers and lead times. For example, the MOQs, regularly updated by the supplier, are not implemented.

The buyer enters all parts individually in the ERP and looks at the relevant information of each one, such as previous order size and the historical yearly demands. Only about 11% of the parts highlighted by the MRP are actually purchased (on average 50 of the 446). The exact quantity to order is based on the historical purchases and the buyer's experience.

The process is very inefficient and has to be performed for many parts because the safety stock is not accurate, and, therefore, the list of parts triggered by this value created by the MRP is inaccurate. All in all, the process takes her about 2-4 hours every other day. By setting proper parameters, both the number of parts that have to be checked and the number of steps to perform are reduced, as the buyer can trust the calculated order quantity and will not have to calculate it every time. The process will also be less dependent on the employee's expertise.

Production
The process of checking the overview of parts going below the safety stock, printed by the MRP, is the same for the production office. The critical difference is, of course, that the creation of the parts has to be managed by the production team themselves. The person responsible for optimising the workshop first determines whether a part is produced in-house or outsourced to sub-suppliers. For this research, the workshop is also considered a warehouse supplier and thus not analysed extensively. If a part is outsourced, the sub-suppliers manager takes care of ordering.

2.3.2 Order specifications

Due to the long delivery times, none of the parts for standard products can be produced on order. Already now, the requested due date of a customer order often has to be postponed before confirmation if not all parts are already present.

One of the parts-suppliers has strict policies regarding a Minimum Order Quantity (MOQ). This quantity varies heavily based on the type of part and has to be considered individually. In the case that this quantity is considerably too high for the purchasing team, they decide to order from another supplier. However, for 213 of the SKUs the MOQ should be taken into account as they have to order from this specific supplier.
About 2% of the parts-suppliers offer quantity discounts of about 3-5%. The purchasing team uses this discount if convenient, but it does not outweigh the disadvantages of ordering in unwanted high quantities. By recommendation of the head of purchasing, these economies of scale will not be considered.

2.4 Demand

In this section, the demand for the parts is analysed. The demand share of the different departments using the warehouse is calculated and visualised in Section 2.4.1. The demand planning is described in Section 2.4.2, as well as the demand patterns found in Section 2.4.3.

2.4.1 Multiple sources

The need for parts comes from arriving customer orders for new products or services. From the ERP, the orders of the current year 2023 are analysed as they represent the most actual demand.

On average, 675 customer orders are received monthly, of which 82 are for the service department. These orders can be transferred into demand for parts. As described, the three teams all rely on the same inventory (the workshop uses sub-parts to produce parts). When looking at the number of parts picked (visualised in Figure 9), we find that considerably more parts are picked for work orders for the assembly department. The workshop asks for about 17% of the total parts picked. While 12% of the customer orders are service orders, the service department hardly creates demand when it comes to picking parts. On average, 8.26 pieces of one SKU are required for an order, meaning clear non-unit sized demand.

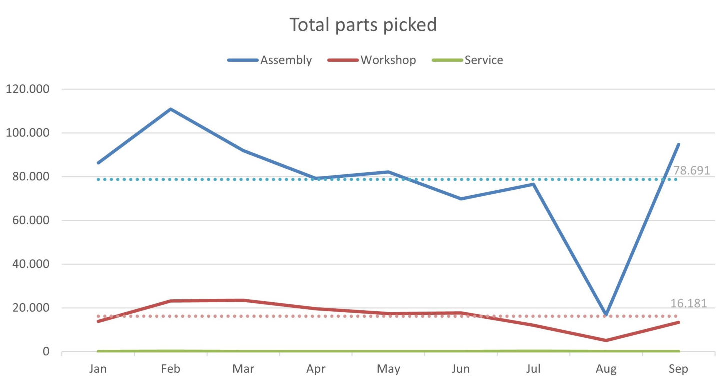

Figure 9: Total parts picked per month in 2023

2.4.2 Demand planning

Currently, demand is managed in an entirely reactive way. The due date promised to the customer is almost fully dependent on the availability of materials, as this is the most time-consuming phase. The production and purchasing offices, determining the incoming flow of parts and thus availability, react only to orders confirmed and entered in the ERP. The products in these orders is split up into the demand of parts using the Bill Of Materials (BOM). When setting the quantity of the new purchase/production, historical demand and purchasing/production data are briefly analysed. There is no active forecasting.

There is no communication with the sales department. This means that demand is only known at the time it occurs and is required. Although improving this communication might be beneficial, as it improves the quality of the forecast of future demand, Fluiten is not yet able to properly implement this. Without additional demand information, the stationary demand for parts can be modelled using set parameters.

2.4.3 Demand patterns

In Appendix G, the full analysis regarding intermittent demand can be found. It has been performed based on monthly demand. The average monthly demand over all years of the SKUs with an average time between demands of at least three months is 0.80. Filtering on an average time period of at least four months and an average demand size, in the months when there were orders, of at least 20 units, we find 215 SKUs. The policy regarding these parts is individually discussed with the production manager.

The parts are not experiencing any seasonality. Although a dip in demand can be found in August for all SKUs, this has to do with the summer closing of the company. During this period, no orders are accepted or scheduled, and thus no demand occurs. Therefore, seasonality will not be considered for the SKUs.

2.5 Supply

In this section, the supply side of the parts is described. The lead time of the parts is difficult to determine. The lead times in the ERP are not up-to-date. There are three sources for the SKUs which all have different characteristics and managers. The lead times and performance of parts-suppliers are analysed in Section 2.5.1. The lead times and management of the sub-suppliers are described in Section 2.5.2, those of the workshop in Section 2.5.3.

2.5.1 Parts-suppliers

Historical orders are analysed to estimate the supplier lead time. The head of purchasing at Fluiten recommends using the time difference between the placement and delivery of the order as a solid approximation of the Supplier Lead Time (SLT). The orders placed between 2018 and October 2023 have been analysed to determine a lead time for each part by taking the average of observed arrivals. If a part has not been ordered in this period, an approximation has been made based on similar parts. Afterwards, the extremes were checked together with the buyer to filter for outliers. Figure 10 shows the lead time of all the individual purchased parts in weeks. There are, for example, 1,169 SKUs with a lead time of 3-4 weeks. It can be confirmed that the lead times are quite long, increasing the importance of on-hand inventory.

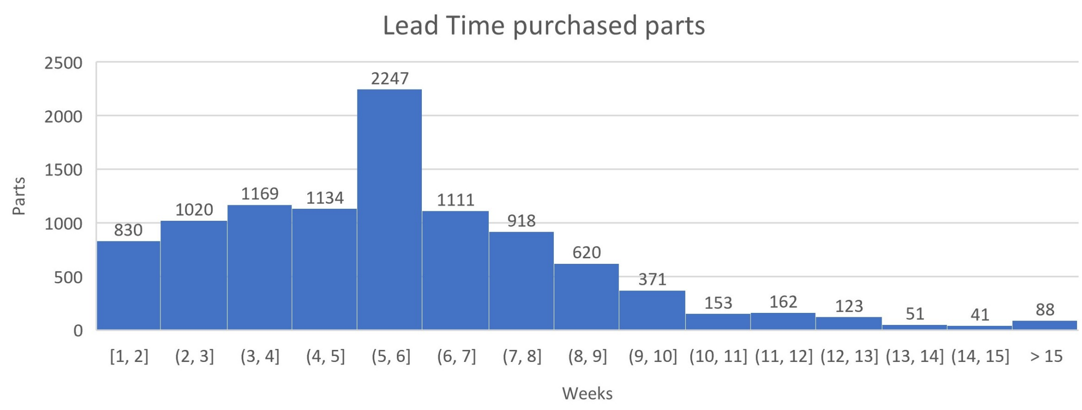

Figure 10: Lead time purchased parts

The suppliers' performance and how well they respect the confirmed delivery date significantly influence the lead time's accuracy. By comparing the date confirmed by the supplier themselves, we do not penalise a long lead time but the credibility of their promises. An analysis of this performance indicates whether additional measures must be taken to avoid backorders. Negative values indicate that the order was delivered early. The full figure can be found in Appendix H.

There are many parts-suppliers. The 25 with at least 40 orders in the selected period have been analysed in detail. The top 10 suppliers delivered 80% of the value of all orders indicating a clear preference. The weighted average delay from the promised delivery date (only taking late deliveries into account), by their share of the total ordered amount, was a delay of only 0.39 days. When ordering the differences between promised and actual delivery date, and taking the weighted average 3rd quartile, we find 3.39 days. This statistical measure is chosen to avoid penalizing early delivery (as, for example, the standard deviation calculation would).
The average value is excellent. The third quartile value is reasonable. The purchasing department has been and is working on improving the reliability of the suppliers. Considering the long supply lead times and flexible choice of supplier for many parts, the random lead times are not initially included in the model on the recommendation of the head of purchasing. An option to include extra lead time will be included in the model.

A general analysis has been done to determine the completeness of orders. For 93% of the orders from the selected 25 suppliers, the delivered quantity was at least the ordered quantity. In the other cases, the rest arrived slightly later. No detailed framework is required.

2.5.2 Sub-suppliers

Regarding the sub-suppliers, who produce the parts that Fluiten could produce themselves as well but at a higher cost, too many factors influence the lead time to use the difference between the two dates as an approximation. The sub-suppliers manager was interviewed to estimate their lead times. When choosing the requested delivery date, he constantly works with five weeks of lead time. This is regardless of the type of parts and the quantity. In this research, we will also take a fixed lead time of five weeks for all sub-suppliers.

2.5.3 Workshop

Like the situation for the sub-suppliers, the lead time for the workshop cannot be determined by comparing the order and delivery date. In cooperation with the sub-suppliers manager, the person responsible for optimising the workshop also works with a lead time of 5 weeks. However, within the workshop, orders can be prioritised based on urgency for Fluiten. In this research, we will also take the five weeks of lead time for all parts of the workshop.

2.6 Performance of inventory management

In this section, the performance of the current inventory management is analysed through the identification of backorders. Section 2.6.1 presents the analysis of backorders both on part- as well as order level. A formula to quantify the labour costs of backorders to enable future comparison between situations is provided in Section 2.6.2.

2.6.1 Occurrence of backorders

We analyse the occurrence of backorders between January and September 2023. We define a backorder as when the picking has started and not all parts are present. The share of parts which were, however, present in inventory is called the fill rate. The date of the picking of each part was compared. If the date of picking of some of the parts was later than others, these were considered backordered. The complete analysis can be found in Appendix E. The backorder analysis is performed on SKU level as well as on workorder level.

Parts - Figure 11 shows the analysis on SKU level, with the SKUs sorted from lowest to highest fill rate and from highest to lowest quantity picked within each percentage. Of the 8,642 SKUs with demand in the analysed period between January 2023 to September 2023, 11% had a fill rate of zero per cent and 61% a fill rate of 100%. The average fill rate is 81%. With the orange line of cumulative picked parts, we see that the parts with a low fill rate only make up a tiny part of the total demand. Additionally, we can conclude that 49% of the pieces picked belongs to SKUs with a fill rate of 100%.

Looking into the SKUs with a fill rate of 0, we find many parts (about 65%) with special backgrounds (such as special contracts with clients or a special combination of materials) as well as 'normal' parts. Currently, these 'special' parts are usually not kept in stock and have an inventory coverage of 0. Inventory management policies for these products will provide a founded proposal for parameters, which can serve as a starting point in the discussion on whether these special parts should be kept in stock.

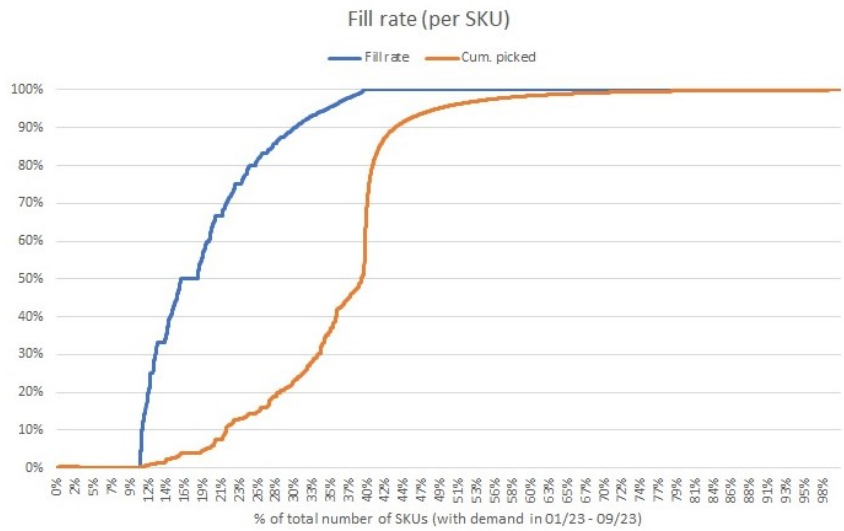

Figure 11: Fill rate on SKU level

Orders - Within workorders, we divide the three types of target departments: the workshop, service, and assembly. The performance of these workorders is visualised in Figure 12. The shares of orders which can be filled directly from inventory are 86%, 83% and 55%, respectively. This shows that especially assembly-work orders are often incomplete. This is more likely to be the case than for the other two types, as the average size of the assembly-workorders is 11.4 SKUs (9 times bigger than the others). The incomplete 45% of assembly-workorders are, on average, 78% complete at the first picking moment. On average, it takes 4.3 days for a workshop order to go from the start of picking until completion. Service orders take 1.3 days, whereas assembly orders, on average, stay 8.0 days waiting for the last SKUs.

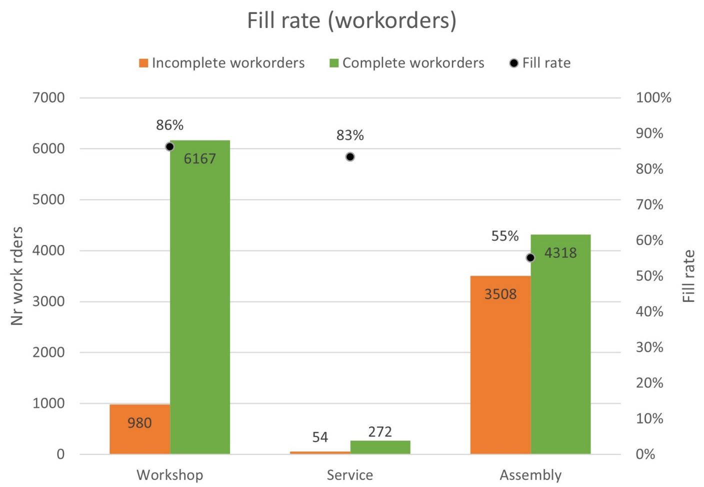

Figure 12: Fill rate on workorder level

To improve the fill rate of the workorders, especially the fill rate of the SKUs required for assembly should be improved.

2.6.2 Backorders quantified

The most significant internal disadvantages of a backorder are the space a half-picked order takes up and the additional manual work that has to be performed. Using an approximation of backordering costs will monetise the improvements in fill rates, which enables a clear comparison between the current situation and future improvements. The space is hard to estimate, but the work can be approximately quantified.

The breakdown of labour costs resulting from a single backorder is described in Appendix F. Within the scope, two employees are actively managing the backorder: the picking and incoming goods employees. Although the exact costs might vary depending on situation-specific circumstances, Formula 2 provides insight into the cost of every backorder, both workshop-, assembly- and service-related.

$$Backorder\ costs = €2.92 + €4.08 * Nr\ backordered\ SKUs \tag{2}$$

The total labour costs consist of fixed costs for every backorder and variable cost depending on the number of SKUs missing in the order.

2.7 Chapter conclusions

In this chapter, the research question *'What does the current inventory situation look like?'* is answered. The sub-questions have guided the research through the various aspects of the current inventory management.

Production-inventory model - The incoming streams of parts into the inventory are the workshop, sub-suppliers and parts-suppliers. The parts from the inventory are used by the workshop, the assembly department, and the service department. The production-inventory model can be described as Assemble-To-Order. The customer order lead time of standard finished products is about 5-10 working days. There is no time for additional sourcing after receiving an order.

SKUs in inventory - There are 19.610 SKUs in the scope of this research, of which 11,663 are currently kept in inventory. From the DBV, we find that only 7% of the SKUs already represent 80% of the annual usage of 2022. These are the parts which should be monitored closely. On the other hand, 80% of the SKUs only represented 5% of the usage. These are the parts of which inventory should be low. As was confirmed by the inventory coverage analysis, many parts currently have an inventory level that is too high due to a lack of insight into the current stock and a risk-averse policy. This level should be reduced.

Current order policies - The current order policies are highly inefficient. Many parts are checked repeatedly, whereas action is taken for only 11% of them. The decisions are heavily based on experience. The ERP system is unaware of the MOQs. Proper inventory management policies reduce the frequency of checking the parts, reduce the dependency on experience and could include minimum order quantities.

Demand - Orders arrive at the sales and the service department. The actual demand for parts from inventory is 83% for the assembly department and 17% for the workshop. The demand for parts for service is minimal. Demand is currently managed in an entirely reactive way. Inventory is managed purely on existing incoming and historical orders. 215 parts have intermittent demand. Non-unit-sized demand has been observed, but no seasonality has been experienced.

Supply - The supply lead times for the parts purchased at the parts-suppliers have been approximated through historical orders. Their delivery date has an average delay of 0.39 days and will not have to be added to the approximated lead time. A lead time of 5 weeks will be taken as a lead time for the sub-suppliers and the workshop.

Performance of inventory management - The current model's performance was quantified through the analysis of backorders. 86%, 83% and 55% of the workshop, service and assembly picking workorders are filled directly from inventory. Most of the workorders are for assembly, so this percentage should be improved. From the analyses, the expectation is that this disruption is caused by the large share of parts (about 25%) with a relatively low fill rate (<80%). The current ordering approach and a lack of proper inventory control policies result in this fill rate.

To conclude, it has been found that many parts have high inventory levels, while this is not required. These should be lowered to avoid unnecessary holding costs. On the other hand, there are many workorders which are missing parts, causing a low order fill rate. The availability of these SKUs (item fill rate), especially those required for assembly workorders, should be improved. Regardless of whether the inventory level should be higher or lower, proper inventory control policies can significantly help determine the correct parameters. These, in turn, reduce the effort required by the production and purchasing teams to manage new supplies.

Research should be conducted to find these suitable inventory control policies. Important characteristics are the non-unit-sized demand and the MOQ for the parts-supplier. To form these policies, some basic form of demand forecast should be found to implement on parts level. As the frequency and size of demand vary heavily over the parts, proper distributions should be chosen.

3 Inventory management: A literature study

In this chapter, a literature study is performed to answer the research question *‘What inventory management methods are proposed in literature that could be applied to Fluiten, with which the right inventory levels can be chosen?’*. Fluiten’s situation is classified according to literature in Section 3.1. Section 3.2 describes different methods to classify the types of SKUs and their characteristics. Different classifications can then be matched to suitable inventory control policies, as described in Section 3.3. The formulas to calculate their corresponding parameters using an order-based performance approach are described in Section 3.4. As these formulas require input regarding demand, a simple demand forecasting method is included in Section 3.5.

3.1 Classification of inventories at Fluiten according to literature

This research focuses on improving the inventory management at Fluiten. Inventory management can be defined as ‘the continuing process of planning, organizing and controlling inventory that aims at minimising the investment in inventory while balancing supply and demand’. It encompasses decisions regarding purchasing, distribution, and logistics, and specifically addresses when and how much to order (Silver et al., 2017, p.16). All parts within the scope of this research are managed in an assembly-to-order manner and could thus be kept in inventory. Currently, safety stock is not kept for all parts which are analysed in this research. As Silver et al. (2017, p.371) mention, decisions of whether items should be kept in stock go beyond the area of production planning and inventory management. Many considerations, including customer relations, are relevant. Some other factors, such as unit variable costs, costs of a temporary backorder, and the carrying charge are proposed. They conclude that the total relevant costs per year are quite insensitive to the precise setting of the control variables, as more substantial savings can be achieved by answering the question of whether the item should be stocked.

Fluitens combination of the assembly department and the workshop results in the classification of a production-inventory system. Here, the order quantities of parts generated by the inventory model determine the production lot sizes and, as production time increases with the number of parts to produce, the production lead times. These lead times, in turn, affect the parameters of the inventory policy (Noblesse et al., 2014). This production-inventory system has endogenous lead times (Boute et al., 2006). The lead times for the parts produced in the workshop are very dependent on multiple aspects such as occupancy. A lead time of 5 weeks has been chosen for all fabricated parts after consulting the workshop manager, but the effects of internal changes should be monitored closely.

At Fluiten, the ERP system checks the inventory position of all SKUs and highlights the ones requiring attention. Every other day, the purchasing and production offices analyse them. Compared to the lead times and general time scale within the company, this review period of 0-2 days is negligible, and the way of working can be considered continuous. Periodic review policies are thus not considered in this chapter.

3.2 SKU Classification

Managerial decisions regarding inventories must ultimately be made at the level of an individual item or product (SKU) (Silver et al., 2017, p.28). Several monitoring systems and processes can be employed to check inventory imbalances and minimise supply and demand dynamics. To simplify this, items are classified into different groups (Dhoka and Lokeswara Choudary, 2013). This enables companies to decide on production strategy, production and inventory management and customer service for entire SKU classes rather than for each product separately. The main aim of any SKU classification is to use the similarity of products regarding different properties to classify products systematically. SKU classification is also frequently used in forecasting and production strategy (van Kampen et al., 2012).

The previously mentioned aims and the context influence which characteristics to base the classification on. These could include volume and variability, unit cost, criticality, and lead time. The technique can be either judgemental or statistical. Judgemental techniques have been proposed by Partovi and Burton (1993) using the Analytical Hierarchy Process (AHP) and Cohen and Ernst (1988) who used cluster analysis to group similar items.

There is no classification of the SKUs in Fluiten. We, therefore, focus on the straightforward and renowned approaches as the concept is new.

3.2.1 ABC classification

The most well-known approach for classification is the ABC-analysis. It is based on the Pareto analysis, which implies that a small portion of items in inventory contributes to high sales (Dhoka and Lokeswara Choudary, 2013). The aim is that if one focuses on the relatively small number of products that represent a major part of the sales volume, rather significant reductions in inventory costs can be obtained (van Kampen et al., 2012).

Although the analysis usually uses three classes, two classes might be considered enough, selecting only A- and C-classes (Hautaniemi and Pirttila, 1999, Yan et al., 2013, T.H. Willis and J.D. Shields, 1990, D.E. Nicol, 1989). Another option is to identify all three classes, find suitable ordering policies for items in classes A and C, and extract fuzzy rules from these policies for items in class B (Mohamadghasemi and Hadi-Vencheh, 2011).

Although the ABC-analysis is attractive because of its simplicity, it is criticised for considering only one aspect. Studies use additional criteria to improve the traditional ABC-analysis (Hautaniemi and Pirttila, 1999). For example, some inexpensive SKUs may be classified as 'A' simply because they are crucial to the operation of the firm (Silver et al., 2017, p.31). In most articles, there are two criteria used in a matrix form. Additionally, the analysis can be modified manually by moving items from one group to another according to specific criteria, such as problems in procurement (Hautaniemi and Pirttila, 1999).

3.2.2 XYZ-analysis

The XYZ-analysis is the dynamic extension of the static ABC-analysis (Pandya and Thakkar, 2016). It distinguishes between items according to their fluctuations in consumption. It can also be seen as the fast, normal, slow-moving technique, as it groups products based on their consumption rate (Scholz-Reiter et al., 2012). Class X-items have a relatively constant demand, Y-items have stronger fluctuations, and class Z-items have high variation. The division is made based on the coefficient of variation, which the following Formulas 3 and 4 can calculate.

$$\sigma = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(x_i - \hat{x})^2} \tag{3}$$

$$CV = \sigma/\hat{x} \tag{4}$$

When considering not a single time period but rather a sequence, for example, in the case of lead time demand analysis, the following Formulas 5 and 6 can be used on the rare occasion of independent and identically distributed demand (Tibben-Lembke, 2006).

$$\hat{x}_L = LT_{weeks} * \hat{x}_{week} \tag{5}$$

$$\sigma_L = \sqrt{LT_{weeks}} * \sigma_{week} \quad (6)$$

In the probable case that demand is not independent and identically distributed, another approach should be used. With enough historical demand, the standard deviation can be directly estimated based on forecast errors over the lead time. Solely non-overlapping periods should be used to avoid a positive correlation between observations (Van der Heijden, 2021a).

There are several challenges when doing the XYZ analysis. The average might be hard to determine based on historical data taking trends into account, and the period must be chosen carefully. The drawbacks of the analysis are, among others, the categorization of new products without established demand patterns and the insensitivity to seasonal patterns, where the coefficient of variation may be high, but the predictability as well (Kumar Dhoka and Lokeswara Choudary, 2013).

3.2.3 Stepwise approach

Hautaniemi and Pirttila (1999) uses a stepwise approach to classify all SKUs, included in Figure 13. They first separate items using the ABC-classification, although they use only the A- and C-classes. Within the A-class, items with a supply lead time shorter than the final assembly schedule are separated. Those with a longer supply lead time are separated based on the demand distributions. The three classifications are singular demand, lumpy demand and continuous demand. This procedure results in five groups of items in a general case in an Assemble To Order company. Not all groups are necessarily filled. For example, their case-company does not have lumpy demand for any of its items due to independent customer orders for usually one unit at a time.

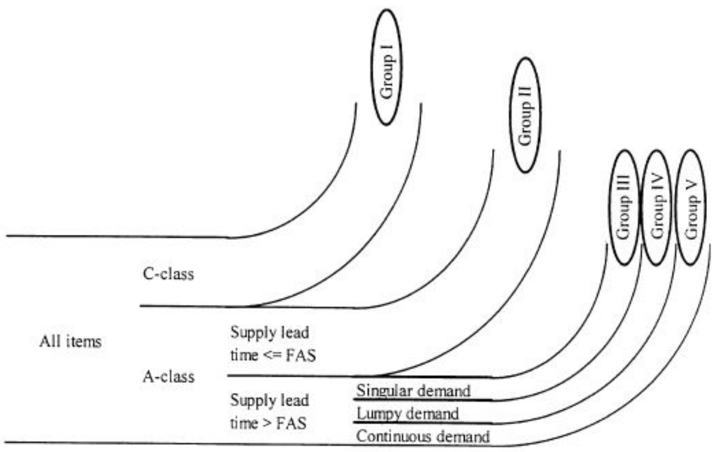

Figure 13: Stepwise approach (Hautaniemi & Pirttila, 1999)

3.3 Inventory control policies

The objective of inventory control is often to balance conflicting goals. One goal is to keep stock levels down to make cash available for other purposes. On the other hand, large batches can give volume discounts, long production runs avoid time-consuming setups, and high raw material inventory reduces the stops in production due to missing materials. Marketing would like to have a high stock of finished goods to be able to provide customers with a high service level (Axsäter, 2015, p.1). All in all, this can be summarised as the role of stock control to 'meet the required demand at a minimum cost' (Wild, 2002, p.7). The fundamental purpose of a replenishment control system is to resolve the following three issues or problems: *How often should the inventory status be determined, when should a replenishment order be placed, how* *large should the replenishment order be.* Managers must first establish how critical the item under consideration is to the firm. The importance of the item helps direct the response to the three questions (Silver et al., 2017, p.240).

In the continuous review case, the stock status is always known. It allows for less safety stock but might miss the reduced setup and shipping costs of batch orders (Silver et al., 2017, p.241). An ordering decision can not be based only on the stock on hand. We must also include the outstanding orders that have not yet arrived, backorders, and possible committed stock (Axsäter, 2015, p.46). This is defined as the inventory position, calculated by Formula 7.

$$Inventory\,position = (On\,hand) + (On\,order) - (Backorders) - (Committed) \tag{7}$$

Once it is determined in which category the item falls and the choice for continuous review is settled, the form of the inventory policy should be chosen (Silver et al., 2017, p.241). The following control systems in Table 1 are the most common (Van der Heijden, 2021c).

	Continuous review
Fixed lot size	(s, Q) or (s, nQ)
Variable lot size	(s, S)

Table 1: Common control systems

A policy can be based on a certain reorder point s, which triggers a new replenishment order. The quantity of this order might be fixed Q or an integer multiple of a fixed quantity nQ. This quantity might be determined by the Economic Order Quantity (EOQ) or ordering requirements. Another approach to the ordering quantity is always ordering up to a certain level S. The policies visualised in Table 1 are described individually.

(s, Q)- or (s, nQ)-policy
These policies are triggered when the inventory position surpasses the reorder point s. A fixed order quantity Q is ordered, or an integer multiple of a fixed quantity nQ. This fixed quantity might be necessary when a supplier limits the freedom in order quantities, for example, due to a fixed number on a pallet. If a single pallet does not suffice, multiple could be ordered (Van der Heijden, 2021c). This policy is also referred to as the two-bin system, as replenishment is triggered when the śecond' bin representing the amount of the reorder point is opened (Silver et al., 2017, p.242). A significant advantage of this system is that it is simple and easy to understand. However, the primary disadvantage is that it is not able to cope effectively with large individual transactions, as the replenishment size Q might not be large enough and the more time-intensive nQ quantity has to be determined (Silver et al., 2017, p.242).

(s, S)-policy
In this system, a variable replenishment quantity is used, which can be calculated with the order-up-to-level S. The quantity to order is always the difference between the level S and the current inventory position. The advantage of this method is most appreciable for A items, as a possible slight improvement in availability has a high influence on the final performance of the inventory management. A disadvantage is the variable order quantity, which suppliers might dislike, and the difficulty of precisely analysing the best parameters (Silver et al., 2017, p.243, Van der Heijden, 2021c).

(S-1,S)-policy
This policy is a special case of the (s, S)-policy mentioned. By setting a reorder point of the order-up-to-level minus 1, an order at the supplier is placed each time there is demand (Feeney

and Sherbrooke, 1966). This policy is appropriate when demand is low but the item is expensive so that the cost of ordering is negligible compared with the cost of holding and shortages (Moinzadeh, 1989). Feeney and Sherbrooke (1966) shows how compound Poisson demand can be used in combination with this inventory policy.

Policy selection
As described, the item's importance influences the suitable inventory control policy. There is no specific model for a specific type of product, but Silver et al. (2017) gives the following rules of thumb in Table 2 when selecting the form of the inventory policy. These have also been used by Mely Permatasari et al. (2017).

	Continuous review
A-items	(s, S)
C-items	(s, Q)

Table 2: Rules of thumb for policy selection

Supplier restrictions might influence the ordering quantity from variable to fixed, even for A-items. For C-items, firms can use a more straightforward approach, such as a simple (s, Q) system. Less effort is devoted to their inventory management because the savings available are pretty small (Silver et al., 2017).

3.4 Calculation of parameters

This section describes the parameters required for the previously described inventory control policies. First, the two relevant criteria for calculating the inventory performance are determined and explained. Then, the definitions and formulas are given with which the policies can be implemented.

3.4.1 Criteria for establishing the safety stock

Fill rate
The cycle service level, the fraction of cycles in which the on-hand stock does not reach zero, is too strict with its definition of stockout, as reaching zero is not necessarily a problem. A more suitable criterion focuses on the performance towards the customers. Therefore, the fill rate is the main criterion for establishing the required safety stock. This is the fraction of customer demand that is met routinely, without backorders or lost sales (Silver et al., 2017, p.249). An equivalent criterion may be the fraction in which a specific item is available off the shelf (Kiran and Loewenthal, 1985). The fill rate is also referred to as P2.

Time Between Stockout occasions
When working with C-items, however, Silver et al. (2017) advocates the use of the Time Between Stockout occasions (TBS). It is a method with which managers are comfortable expressing their risk aversion. It is more straightforward than dealing with probabilities or fractions. As many C-items could be involved in a single customer order, a very high level of service must be used for each item. If only a small expense is added for carrying high stock, values for the TBS such as 5-100 years are not unreasonable (Silver et al., 2017, p.360). In this research, however, the time span of one year will be used. The TBS can thus not be properly measured as output. For this reason, the fill rate will be used for C-items as well.

Order Fill Rate
Most standard inventory models do not take into account connections between items; they assume that demands for each item are independent of the others. This is an item-based approach. However, this might result in good performance on an item basis but poor order-based performance.

The Order Fill Rate (OFR), the probability of filling an entire customer order immediately from shelf is an important service measure in industry, crucial even to assemble-to-order practices (Song, 1998). Larsen et al. (2008) distinguish the OFR from the standard unit-based fill rate by terming the latter the Volume Fill Rate (VFR). Also Boylan and Johnston (1994) propose several performance measures based on partially or completely filled orders or order lines. The basis for using an OFR is the focus on individual customer orders and the ability to fill each of those orders in full from inventory. Equal weight is attached to each order irrespective of its size, corresponding to a cost of shortages related to the occurrence of a shortage (Larsen et al., 2008).

Larsen et al. (2008) propose a model with compound renewal demand methods using an OFR as service level requirement. Larsen and Thorstenson (2008) provides a comparison between using the OFR or VFR with these compound renewal demand processes. Song (1998) uses convolutions, compound Poisson and batch-size distributions to compute the order fill rate.

Teunter et al. (2017)
The OFR can also be used as a service level requirement by using it as input in setting the Target Fill Rates (TFRs) of the individual SKUs. Teunter et al. (2010) have focused on improving the ABC classification by proposing a criterion which includes a penalty cost per backordered item. By also indicating the criticality of an item, the penalty cost can be higher for more critical items. A similar approach is proposed by Teunter et al. (2017), where the targeted system fill rate FR_T leads to individual fill rates using the following Formulas 8, 9 and 10. This FR_T should not be confused with the OFR, as it represents merely the overall assortment fill rate.

$$PriceCriticalityRatio_i \; (PCR_i) = \frac{p_i}{c_i} \tag{8}$$

$$AveragePriceCriticalityRatio \; (APCR) = \frac{\sum_{i=1}^{N} \frac{p_i}{c_i} D_i}{\sum_{i=1}^{N} D_i} \tag{9}$$

$$1 - FR_i \approx (1 - FR_T) * \frac{PCR_i}{APCR} \tag{10}$$

The variable c_i represents the (relative) criticality of a backlog for SKU i per time unit, where the SKU's influence on the OFR can be included. For this research, this criticality is expressed as the number of orders a certain SKU was part of. It should be noted that the fill rate value FR_i can be negative for very expensive SKUs. Some positive lower bound on the fill rate for each individual SKU can be imposed, although this does imply that the achieved system fill rate may exceed its target. Even when setting minimum fill rates of 50%, 70% or 90%, Teunter et al. (2017) still found reductions of system stock value.

van der Heijden (2024)
Another approach has been designed by the supervisor of this thesis van der Heijden (2024). In this approach, an (s,Q)-policy with Normal lead time demand and a First Come First Service inventory control is assumed. Additionally, the distribution of the on-hand inventory at the moment of the arrival of an order is approximated to be equal to those in steady state, which is known to be as Equation 11(Axsäter, 2015, p.92).

$$F_{OHI,i}(x) = \frac{\sigma_{L,i}}{Q_i} \left[G\left(k_i - \frac{x}{\sigma_{L,i}}\right) - G\left(k_i + \frac{Q_i}{\sigma_{L,i}} - \frac{x}{\sigma_{L,i}}\right) \right] \tag{11}$$

If for every order type j, exactly $a_{ij} \geq 1$ pieces of SKU i are required from the set Ω_j, the OFR is 1 minus the probability that the on-hand inventory is at most $a_{ij} - 1$. Focusing solely on the first-order effects and leaving out the second term, we find the following expression (Equation 12).

$$OFR_j \approx 1 - \sum_{i \epsilon \Omega_j} \frac{\sigma_{L,i}}{Q_i} \left[G\left(k_i - \frac{a_{ij}-1}{\sigma_{L,i}} \right) \right] \tag{12}$$

Setting one target weighted average OFR β for all orders with demand m_j for order type j, where $M = \sum_{j \epsilon J} m_j$ and the Ψ_i the set of orders which require SKU i, provides the following model for minimizing the inventory holding costs (Equations 13 and 14).

$$\min_{k_i, i \epsilon I} \quad \sum_{i \epsilon I} k_i \sigma_{L,i} v_i \tag{13}$$

$$s.t. \quad \sum_{i \epsilon I} \frac{\sigma_{L,i}}{Q_i} \sum_{j \epsilon \Psi_i} m_j G\left(k_i - \frac{a_{ij}-1}{\sigma_{L,i}} \right) \leq M(1-\beta) \tag{14}$$

Using the Lagrange relaxation, we find Equation 15.

$$\min_{k_i, i \epsilon I; \lambda} \quad \sum_{i \epsilon I} k_i \sigma_{L,i} v_i - \lambda \left\{ M(1-\beta) - \sum_{i \epsilon I} \frac{\sigma_{L,i}}{Q_i} \sum_{j \epsilon \Psi_i} m_j G\left(k_i - \frac{a_{ij}-1}{\sigma_{L,i}} \right) \right\} \tag{15}$$

Subsequently, setting the partial derivation on k_i to zero gives Equation 16.

$$\sigma_{L,i} v_i + \frac{\lambda \sigma_{L,i}}{Q_i} \sum_{j \epsilon \Psi_i} m_j \left[\Phi\left(k_i - \frac{a_{ij}-1}{\sigma_{L,i}} \right) - 1 \right] = 0 \tag{16}$$

Finally, dividing by $\sigma_{L,i}$ results in Equation 17.

$$v_i + \frac{\lambda}{Q_i} \sum_{j \epsilon \Psi_i} m_j \left[\Phi\left(k_i - \frac{a_{ij}-1}{\sigma_{L,i}} \right) - 1 \right] = 0 \tag{17}$$

Using this formula, for each value of λ, the value of k and corresponding fill rate can be found for every SKU. The value of λ depends on the desired OFR. Although a Normal distribution was assumed for the lead time demand of every SKU, the found near-optimal fill rates can be applied in models using other distributions.

Chosen method

Due to the importance of the OFR for Fluiten, an order-based approach should be used when determining the safety stock. Both the method by Teunter et al. (2017) and the one by van der Heijden (2024) are suitable without compound renewal methods. They approach the order-based performance very differently. van der Heijden (2024) sets the overall OFR as an initial target value, and uses the Bill Of Materials for to translate demand of products back to parts.The quality of the data concerning the BOM is then very important. Teunter et al. (2017) focusses on a weighted average TFR for the SKUs where the influence on orders is included via the criticality, focussing more on the number of orders than on the requirements for one order. As this way of determining criticality has not been done before, it is unknown whether it provides good results. As both methods could be applied to Fluiten, a small-scale experiment should be performed to determine which is most suitable for the full scale.

3.4.2 Undershoot

A replenishment action is not necessarily taken when the stock level is precisely at the reorder point (Silver et al., 2017, p.251). The undershoot can be defined as the difference between the reorder point and the inventory when an order is placed (Gutierrez and Rivera, 2021). This can occur through non-unit-sized demand or high demand during a review period. The first cause will be described in detail as relevant to this research.

Non-unit-sized demand might result in undershoot when a certain order drops the inventory position below the reorder point. This is visualised in Figure 14. The inventory position when placing the order is thus not the reorder point but already a smaller value.

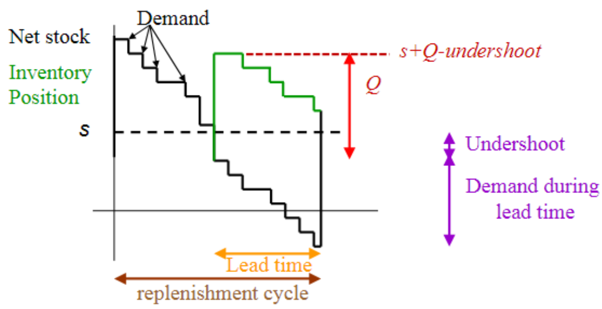

Figure 14: Undershoot due to non-unit-sized demand in a (s,Q)-system (Van der Heijden, 2021b)

A stockout occurs if the sum of the undershoot plus the total lead time demand exceeds the reorder point s (Silver et al., 2017, p.328). The order quantity and the undershoot are related since the order quantity should even out the negative influence of the undershoot for the next cycle (Gutierrez and Rivera, 2021). Considering a sequence of customer orders D_i, independent and identically distributed, the following Formulas 18 and 19 can be used to estimate the mean and variance of the undershoot Z for discrete orders (Silver et al., 2017, p.329).

$$E[Z] = \frac{E[D^2]}{2E[D]} - \frac{1}{2} \tag{18}$$

$$Var[Z] = \frac{E[D^3]}{3E[D]} - \frac{1}{4}\left(\frac{E[D^2]}{E[D]}\right)^2 - \frac{1}{12} \tag{19}$$

The expected value, and the second and third moments are calculated differently for the Normal and Gamma distribution. For parts with Normally distributed demand we use Formulas 20, 21 and 22 (Van der Heijden, 2021b);

$$E[D] = \mu \tag{20}$$

$$E[D^2] = \mu^2 + \sigma^2 \tag{21}$$

$$E[D^3] = \mu^3 + 3\mu\sigma^2 \tag{22}$$

For parts with Gamma distributed demand, the following Formulas 23, 24, 25, 26 and 27 can be used (Al-Ahmadi and Yanikomeroglu, 2010);

$$\alpha = \frac{\mu^2}{\sigma^2} \tag{23}$$

$$\beta = \frac{\mu}{\sigma^2} \tag{24}$$

$$E[D] = \frac{\alpha}{\beta} \tag{25}$$

$$E[D^2] = \frac{\alpha(\alpha+1)}{\beta^2} \tag{26}$$

$$E[D^3] = \frac{\alpha(\alpha+1)(\alpha+2)}{\beta^3} \tag{27}$$

3.4.3 Reorder point

Stock is held either because it is convenient to buy in bulk or because the item is required in a shorter time period than the supply can provide it. In the latter case, there is some uncertainty about the quantity required, so some safety stock is needed. It is found that the major uncertainty is caused by customers and their unpredictable requirements (Wild, 2002, p.96). The reorder point to aim for thus covers the quantity required and the uncertain variability. In case of the continuous review and undershoot, as relevant to this research, the reorder point can be calculated in the following manner with formula 28(Silver et al., 2017, p.259,Van der Heijden, 2021c).

$$Reorder\ point\ s = \hat{x}_L + E[Z] + k\sqrt{var[x_L] + var[Z]} \tag{28}$$

It is important to note that the user should always have the option of adjusting the reorder point to reflect factors that are not included in the model (Silver et al., 2017, p.259). The parameter k, also known as the safety factor, determines the coverage of the uncertainty of demand. It has to be made sure that its value is at least as large as the lowest allowable value (e.g. zero) (Silver et al., 2017, p.269). Also Presutti and Trepp (1970) and Silver and Rahnama (1987) restrict k to non-negative values. It can be calculated based on many different particular shortage costs or service measures used. In this research, the fraction of demand satisfied from the shelf (P2) is relevant.

Fill rate, Normal - When working with the fill rate, we find the reorder point s depending on the demand distribution of the part. If demand is normally distributed, the safety factor k can be found using the following Formula 29 and the Normal loss function. If undershoot is included, this standard deviation should be added to the denominator.

$$G_u(k) = \frac{Q}{\sigma_L}(1 - P_2) \tag{29}$$

Fill rate, Gamma - If the variability of lead time demand is higher than reasonable for a normal distribution and a Gamma distribution is more fit, the following Formulas 30 and 31 can be used (Silver et al., 2017, p.745, Tyworth et al., 1996).

$$P_2 = 1 - \frac{ESPRC}{Q} \tag{30}$$

$$ESPRC = \alpha * (1/\beta) * [1 - F(s; \alpha + 1, (1/\beta))] - s[1 - F(s; \alpha, (1/\beta))] \tag{31}$$

Fill rate, Negative Binomial - With a high variance-to-mean ratio for the expected lead time demand of slow-movers, the discrete Negative Binomial function might be fit (Agrawal and Smith, 1996). The following Formula 32can be used to determine the expected units short per replenishment cycle for this distribution (Silver et al., 2012).

$$ESPRC = \frac{rp}{1-p} * [1 - F(s - 1; r + 1, p)] - s[1 - F(s; r, p)] \tag{32}$$

For items with intermittent demand, it is recommended to separate two components of the demand process - namely, the time between consecutive transactions and the magnitude of individual transactions (Silver et al., 2017, p.122). Croston (1970) provides an updating procedure for this but warns that infrequent updating introduces a marked lag in responding to the underlying parameters. Syntetos and Boylan (2005) have tested several methods with respect to their forecasting accuracy, concluding that different estimators are most suitable for different sets of demand data and different accuracy measures can lead to different conclusions.

3.4.4 Order quantity

We minimise the relevant costs when deciding on the appropriate order quantity. Also for the order quantity Q, manual override should be possible to incorporate factors which are not included in the model or assumptions which do not hold (Silver et al., 2017, p.147). One of the earliest and most well-known results of inventory theory gives the following Formula 33 to define the Economic Order Quantity, using the costs of ordering (A), the item's value (v) and the holding costs percentage (h).

$$EOQ = \sqrt{\frac{2A\mu}{vh}} \tag{33}$$

In practice, the order quantity often has to be an integer. The best value is one of the two surrounding integers (Axsäter, 2015, p.54). The exact value can be chosen mathematically, but rounding the EOQ to the nearest integer usually works well enough for simplicity. When the choice has to be made between two small numbers, and the choice may significantly impact results, Axsäter (2015) provides the following approach; the lower value $Q = n$ should be chosen if $Q^*/2 \leq (n+1)/Q^*$. Lastly, limits on order sizes such as a Minimum Order Quantity or a Fixed Order Quantity should be considered. In the case that the EOQ is less than the MOQ, the best allowable order quantity is the supplier or production minimum (Silver et al., 2017, p.165).

3.4.5 Order-up-to-level

The order-up-to-level S determines the size of a flexible replenishment order. The order should increase the inventory position to S. The reorder point s is known for the continuous policies relevant to this research. As long as we know that the inventory position is exactly this reorder point s when the order is placed, the previously determined order size can be used following Formula 34 (Axsäter, 2015, p.107).

$$S = s + Q \tag{34}$$

However, in the case of undershoot, this quantity has to be added to the order using Formula 35. This is visualised in Figure 15 by Van der Heijden (2021c).

$$S = s + Q - E[Z] \tag{35}$$

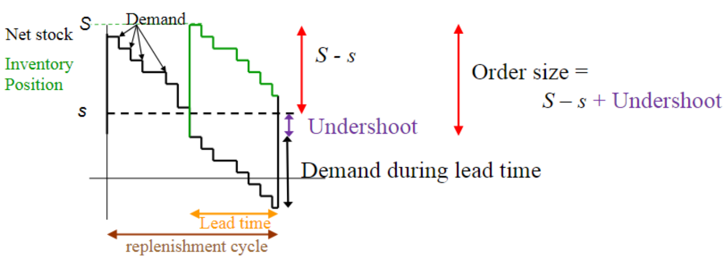

Figure 15: A (s,S)-system with undershoot (Van der Heijden, 2021c)

3.5 Demand forecasting

In many formulas described in the previous section, the demand during a specific period was one of the input variables. The forecasted quantity is of high importance to the exact control policies. While the relative impact of demand uncertainty on the inventory varies across industries, it will almost always be a significant factor for inventory management (Silver et al., 2017, p.73). Typical for forecasts is that they concern a relatively short time horizon; it is seldom necessary to look more than one year ahead (Axsäter, 2015, p.7). Forecasts can be based on a combination of an extrapolation of what has been observed in the past and informed judgment about future events (Silver et al., 2017, p.73). Informed judgment might concern a planned sales campaign or a weather forecast. When using historical data, time-series forecasting models can be applied, which can relatively easily be implemented in computerised inventory control systems (Axsäter, 2015, p.7).

3.5.1 Distributions

Demand can be forecasted using demand distributions. The normal distribution is convenient to use, widely tabulated and built into spreadsheets. The impact of using other distributions is *usually* relatively small, particularly when recognizing the other inaccuracies present (estimates of the parameters) (Silver et al., 2017, p.275). A key property is that it can be described using only two parameters, the average and the standard deviation (Wild, 2002, p.98).

However, if the ratio $\sigma_L/\hat{x}_L$ is greater than 0,5, the normal distribution provides issues with the right tail of the distribution when choosing the reorder point. In this case, it should be considered to use the Gamma distribution Silver et al., 2017, p.275.

In the case of slow-moving items with an average lead time demand below ten units, the demand should be approximated by a discrete distribution rather than a continuous one. The normal distribution is squashed against zero usage and is replaced by the skewed Poisson distribution (Silver et al., 2017, p.352; Wild, 2002, p.98). However, this is only fit when σ_L is within 10% of $\sqrt{\hat{x}_L}$ (Silver et al., 2017, p.352). As for this distribution, the standard deviation is equal to the square root of the average demand; this relationship simplifies the review level formula for slow-moving items (Wild, 2002, p.108). For slow-moving items with a larger variance, the Gamma distribution is not suitable as it is a continuous distribution. Then, the Negative Binomial (NB) could be a good option (Agrawal and Smith, 1996). For NB distributions with a real-valued parameter r, the Generalised Negative Binomial (GNB) distribution can be used (Jain, 1971). Here, the following recursive Formulas 36 and 37 can be used as probability density functions, using success probability p and the number of failures r;

$$f(0; r, p) = (1 - p)^r \tag{36}$$

$$f(k; r, p) = f(k - 1; r, p) * \frac{p(k + r - 1)}{k} \tag{37}$$

3.6 Chapter conclusions

In this chapter, a literature study was performed to answer the research question *'What inventory management methods are proposed in literature that could be applied to Fluiten, with which the right inventory levels can be chosen?'*. The situation at Fluiten is compared with existing literature to highlight situation-specific remarks, highlighting the decision of whether a part should be kept in inventory, endogenous lead times and continuous policies.

There have been good proposals to classify the SKUs in literature. The proposed methods, ABC-analysis, the coefficient of variation in historical demand, and the stepwise approach from Hautaniemi and Pirttila (1999), will be used to develop a tailored model in the next chapter.

The common continuous review policies and the formulas for their parameters are described. Using the created classes, these policies will be assigned to the SKUs. As input in these parameters, not only item-based KPIs but also the Order Fill Rate is an important performance measure. Using either the technique proposed by Teunter et al. (2017) or van der Heijden (2024), an order-based approach can be implemented into the individual TFRs for each SKU. Both techniques will be tested to find the most suitable one. Parts might experience undershooting due to non-unit-sized demand. The expected magnitude of customer orders for intermittent demand should be estimated separately from the frequency of the arrival of orders. For the purchased parts with a MOQ, it is essential to consider this when determining the order size.

To have some demand forecast, needed to estimate policy parameters such as demand during lead time, the different suitable distributions are described. Based on the historical demand of the SKUs, either the Normal, Gamma, Poisson or Generalised Negative Binomial distribution can be used.

4 Solution design

In this chapter, the question *'What inventory management methods are most applicable for the SKUs and how should the inventory management tool be designed?'* is answered. In Section 4.1, the classification steps, demand distributions and required additional user input are described. The following Section 4.2 matches the created classes to inventory control policies. The formulas to calculate the necessary parameters are listed in Section 4.3 as well as constraints to the model. The conclusion of this chapter can be found in Section 4.4. The data connections used for this solution are described in Appendix I.

4.1 SKU characteristics

In this section, the individual characteristics of the SKUs are determined. Section 4.1.1 describes the created tailored classification method. In Section 4.1.2, the SKUs are matched to a distribution with which to estimate lead time demand. Section 4.1.3 describes the additional required user input.

4.1.1 SKU classification

Figure 16 shows the method used to classify the SKUs. Using this classification, the SKUs which should not be kept in stock or have insufficient data to work with are excluded first. Parts with intermittent demand are separated due to their special characteristic. Then, a Pareto analysis determines the importance of a SKU and, consequently, the level of detail and attention the inventory policy receives. Within the A-class, the distribution with which demand can be approximated is determined based on the coefficient of variation in lead time demand. For the C-items, slow-movers with a lead time demand of less than 10 are separated.

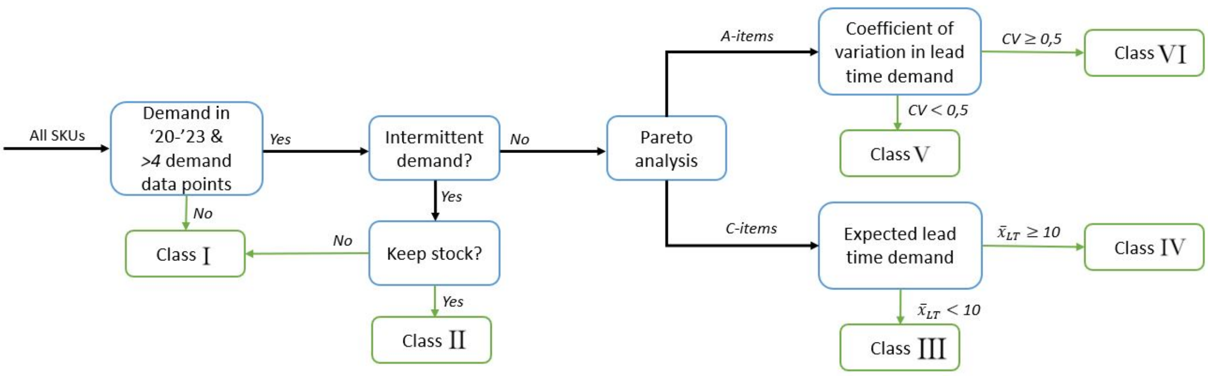

Figure 16: SKU classification method

The method is roughly based on the stepwise approach by Hautaniemi and Pirttila (1999). Three alterations were made. Firstly, SKUs with no demand in the period September 2020 - August 2023 are separated. Due to being outdated or replaced, these have so little demand that they should be managed on order. Parts that have been introduced so recently that they do not have sufficient training data, at least 4 demand data points, are placed in the same class. Secondly, the scope has been chosen in such a way that the research tests what the inventory policies should be if none of the parts could be managed based on MRP. The Supplier Lead Time/Final Assembly Schedule step has thus been taken out. Thirdly, the items are processed with an analysis based on lead time demand mean and variance rather than singular/lumpy/continuous demand classifications to simplify the step from historical demand to demand distributions. This concerns distributions suitable for high variation or slow-movers.

As was often opted for in literature and Hautaniemi and Pirttila (1999) choose to do as well, only two classes were used. From the DBV in Section 2.2.2, it became clear that a small share of the parts already makes up a very high share of the annual value. The advantage of creating a class B is thus likely very small.

Historical demand
To classify the SKUs following this method, historical demand is used. The weekly demand for each SKU was taken from the ERP system. This way, the demand from all three demand sources is combined. The demand for each part is thus not connected to the demand of the final product(s) it belongs to. Weeks without demand before the creation of a SKU should not be considered as '0' when taking the average demand, as the part did not exist yet. Using the first movement dates, all weeks before the first existing date are considered 'empty'.

The total available data is from January 2018 until the end of 2023. The actual demand in the year 2023 will be used to test the created policies. Different data selections will be used based on the purpose of each analysis. The period from September 2020 - August 2023 is used in the Pareto analysis. The years 2019-2022 will be used as training data to determine the parameters that will be tested with the data from 2023.

Based on the historical data, the first class can be filled with the 4,671 SKUs without demand between September 2020 and August 2023. From the remaining 14,939 SKUs, 596 were only introduced in the last four weeks of 2022 or in 2023. As their training set is thus minimal to non-existent, they are considered Class 1 in the main part of this research. Afterwards, a separate recommendation can be calculated. 552 individual SKUs were selected by the production manager to not be kept in stock (anymore). These parts are either in transition to another combined code or are repair kits. 81 of the 215 intermittent demand parts, as identified in Section 2.4.3, are not kept in stock either. This sums as 5,900 SKUs in Class 1.

The remaining 134 intermittent demand parts are placed in Class 2.

Pareto analysis
Using the Pareto analysis visualised in Figure 17, the division between the C- and A-items can be made for the remaining 13.576 SKUs. The analysis uses the annual usage value of the SKUs, which is determined by multiplying the annual demand by the item's value. The three years, September 2020 - August 2021, September 2021 - August 2022 and September 2022 - August 2023, are used to include historical data in this analysis. This division over classes does not influence the calculation of parameters and thus does not influence the future policy test. The most recent year is therefore included in this analysis to include recent changes in sales. As the most recent value should carry the most importance, a weighted average is used. The weights have been chosen together with the purchasing office, implementing the weights they currently (unconsciously) use when analysing past usage.

$$Demand\,value = 0.1 * (20/21) + 0.3 * (21/22) + 0.6 * (22/23) \tag{38}$$

After this analysis, 4,656 SKUs (23,7%) are placed in the A-class, contributing to the annual usage value with 95%. The other 8,920 (45,5%) items are classified as C and complete the last 5%.

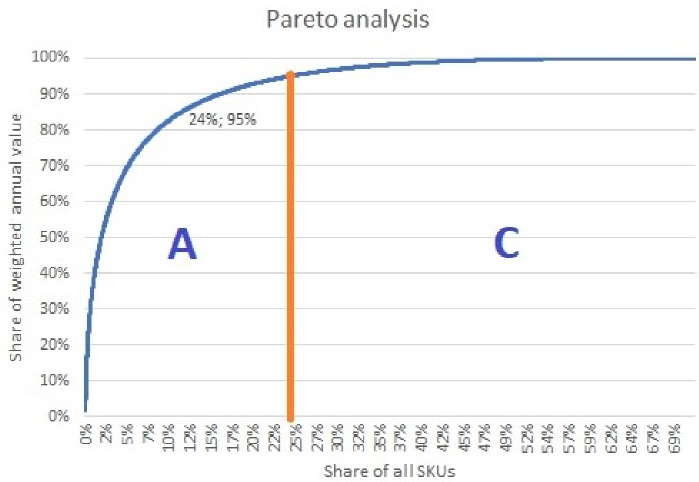

Figure 17: Pareto analysis

Lead time demand parameters
Significant historical data, used to determine lead time demand parameters, is available for almost all parts. As the actual demand in 2023 will be used to determine the performance of the policies, this will not be used to determine the lead time demand characteristics. The analysis of this data did not result in any significant trend or seasonal effects. The absence of seasonal effects was confirmed by the purchasing team. They had not looked into a trend previously.
The average lead time demand of parts is calculated by initially taking the average weekly demand during 2019-2021. Simple exponential smoothing is then applied using the weekly demand in 2022 following Formula 39. All C- and A-items are categorised into 25 groups depending on part type in cooperation with the production manager. Based on minimizing the Mean Squared Error, a value for alpha is chosen for each group. These values vary from 0.01 to 0.1, which are reasonable for the update period of a week. For items which were introduced in 2022, the average weekly demand is calculated over the available weeks. Following Formula 5, the average weekly demand is multiplied by the supplier lead time in weeks to find the lead time demand.

$$\hat{a}_t = \alpha x_t + (1 - \alpha)\hat{a}_{t-1} \tag{39}$$

Using the separation point of a lead time demand of 10, C-items are split into two classes. Class 3, with less than 10 items during lead time, is filled with 8,539 SKUs. The other 381 have a higher average lead time demand and go into Class 4.

The standard deviation has been calculated directly from existing historical lead time demand on the condition that there were at least three data points. These data points were obtained by summing weekly demand for the length of the supply lead time, ensuring non-overlapping periods to avoid a positive correlation. In the case that insufficient data was available, Formula 6 is used to approximate the standard deviation of lead time demand using the standard deviation of weekly demand.

Within the A-items groups, two classes will be used, which, based on a proper separation point, enables a step to a lead time demand distribution. Based on statistical recommendations from the literature framework, the division point is a CV for a lead time demand of 0.5 (Silver et al., 2017, p.275).

The coefficient of variation can then be determined by dividing the standard deviation by the average lead time demand. We find 293 A-items with a CV of less than 0.5 and 4,363 with a

CV higher than 0,5. Looking into the weekly demand for the parts with a high CV, we find that there are some weeks without any demand and other weeks with a high number of orders, resulting in a high standard deviation.

Final classification
Following the previous analysis and the chosen method of classification, the SKUs can be divided over the classes as shown in Table 3.

Class 1	Class 2	Class 3	Class 4	Class 5	Class 6
5,900	134	8,539	381	293	4,363

Table 3: Number of SKUs per class

4.1.2 SKU demand distribution

The analysis of the coefficient of variation of lead time demand also provides insight into the distribution of the lead time demand. The focus is usually on cases with normally distributed demand. It is convenient from an analytical standpoint, it is built into spreadsheets, and the impact of using other distributions is usually relatively small. However, there are some cases when another approach might be more fit.

Having identified the parts with intermittent demand and a need for inventory, they are placed in Class 2. For each of these parts, it has been confirmed that the demand has always come from a single customer. Therefore, a compound distribution will not be used. Fluiten wants to keep an inventory level that should always be able to cover a single customer's order with a certain probability. After analysing the non-zero demand data and its statistical characteristics (CV<0.5), a Normal distribution is considered fit for the expected size of an order. The mean and standard deviation of the order size have been determined using all non-zero demand data since 2018.

Slow movers, with a lead time demand of fewer than ten units, might benefit from a tailored distribution. Silver et al. (2017) comment on the Gamma, Laplace and Poisson distribution while also referring to other references for other options. The Poisson distribution was not fit, as no parts passed the restriction of a σ_L is within 10% of $\sqrt{\hat{x}_L}$. Due to the high variance, the Generalised Negative Binomial distribution was tested using the generalised method to include the possibility of non-integer parameters. This distribution provided a good fit and is therefore used for these SKUs.

As the ratio $\sigma_L/\hat{x}_L$ is greater than 0.5 for many A-items as well as the fast-moving C-items, the Gamma distribution was tested, which resulted in a good fit, as it acknowledged high standard deviations. For this reason, both fast-moving C-items and A-items with a high CV value, Classes 4 and 6, will be analysed with the Gamma distribution.

All parts in Class 5 are assumed to be normally distributed as they have a CV value of less than 0.5.

4.1.3 User input

Having classified the SKUs in the six classes and determined the corresponding demand distribution, the tools user input is required for the calculations for classes two to six.

The first two variables are the ordering costs (A) and the holding costs (h). These influence the parameters of the policies as they are part of the EOQ calculation.

The third variable is an *extra safety lead time*. As described in Section 2.5, the supplier lead times in this research have been based on interviews with experts and analysis of historical orders.

However, we hereby assume that all deliveries are complete and following the determined lead time. Occasionally, a supplier communicates that he is busy and lead times are longer. The user can include extra safety measures in these situations by setting an extra safety lead time for specific SKUs. The expected demand and its variation during lead time are influenced by this number.

The last (set of) variable(s) enables setting the fill rate of either the individual SKUs. Section 3.4 provided two suitable methods to include the order-based approach in the fill rates of the individual items. The method of Teunter et al. (2017) requires a targeted weighted fill rate over all SKUs as well as some expression of criticality for each SKU. The method by van der Heijden (2024) requires a targeted OFR. As both methods to set individual TFRs are tested, both sets of input are required for the testing. The final individual fill rates directly influence the final parameters for the policy, as it is a part of the formulas calculating the safety factor. The higher the fill rate, the higher the reorder point and, thus, the higher the inventories.

4.2 Choice of inventory control policies

With the different SKUs classes come different suitable policies. The chosen policies depend on the characteristics of the classes and the policies as found in Chapter 3. Table 4 shown the combinations.

Class 1	**Class 2**	**Class 3**	**Class 4**	**Class 5**	**Class 6**
MRP	(S-1,S)	(s,Q) - GNB	(s,Q) - Gamma	(s,S) - Normal	(s,S) - Gamma

Table 4: From class to policy

Class 1 contains the parts which should be managed based on MRP data. These parts either did not have demand in the last three years, do not have sufficient data for the large-scale analysis or have been selected by the production manager. Parts with intermittent demand which are not kept in stock are also in this class.

Class 2 contains intermittent demand items for which there should be inventory. An (S-1,S)-policy is selected for this group to ensure that if an order arrives, it can be filled from stock, but no unnecessary additional safety stock is kept. The size of the order is estimated using the Normal distribution as the coefficient of variation of the order size is less than 0.5.

SKUs which were classified as C-items from the Pareto analysis are placed in Classes 3 and 4. These parts will be managed by a (s,Q)-policy as it is simple and easy to understand. The C-items are relatively unimportant and do not require special attention, but an appropriate simple control policy should be used because of their large number. A fixed lot size ensures appropriate inventory levels with minimal required effort in a single reorder moment, as currently the quantity for orders is calculated and entered manually. For Class 3, the GNB distribution is used, whereas Class 4 lead time demand is estimated with the Gamma distribution.

Class 5, with A-items and a coefficient of variation of lead time demand of less than 0.5, is managed by a (s,S)-policy using the Normal distribution for demand forecast. The (s,S)-policy is more detailed with a variable lot size. Determining the size of the reorder requires more attention and effort, but it is worth it due to the importance of A-items. From the analysis of the coefficient of variation in lead time demand, a Normal distribution has proven to be fit. Undershoot will be included for SKUs in this class.

Finally, Class 6 is matched to a (s,S)-policy with the Gamma distribution. This distribution is more suitable due to high variability in lead time demand. Possible undershoot is taken into account.

For the SKUs with an MOQ, the replenishment order size should be at least as large as this restricting value.

4.3 Calculation of parameters

Having collected all characteristics of each SKU and connected them to a policy, the corresponding required parameters should be calculated. Following the formulas found in chapter 3, we find the following approaches as shown in Table 5.

Policy	Reorder point	Quantity	Order-up-to-level
(S-1,S)	$\hat{x}_L + k\sigma_L$	$S - IP$	$s + 1$
(s,Q) - GNB	$Formula\, 32$	EOQ	-
(s,Q) - Gamma	$Formula\, 31$	EOQ	-
(s,S) - Normal	$\hat{x}_L + E[Z] + k\sqrt{var[x_L] + var[Z]}$	$S - IP$	$s + EOQ - E[Z]$
(s,S) - Gamma	$Formula\, 31 + E[Z]$	$S - IP$	$s + EOQ - E[Z]$

Table 5: Parameters calculation

The safety factor k can be calculated depending on the targeted performance rate, using the formulas in chapter 3 for the Normal distribution. As the safety stocks and reorder points cannot be negative, the safety factor k will be set to zero in case of a negative value resulting from the formula. This is the case for 5 SKUs and might result in higher performances than targeted for these parts. The targeted performance rate is expressed as the fill rate. The safety stock for SKUs with the Gamma distribution can be calculated directly with Formulas 30 and 31. Formula 32 provides a direct method for the safety stock calculation of GNB distributed SKUs. The undershoot parameters will initially be applied to Classes 5 and 6 as this detail will have the most impact on these high-value items. The undershoot variance is considered negligible for Class 6 due to its small values and incompatibility with the Gamma distribution. Different settings will be tested in the sensitivity analysis in chapter 6.

The targeted fill rate used as input in these formulas will be determined by using either the method by Teunter et al. (2017) or van der Heijden (2024). Both will be tested on a small scale, after which one will be chosen to apply to the full problem.

4.3.1 Constraints

In its current design, the model comes with some constraints. Firstly, the arriving replenishment orders are assumed to be complete and delivered within the specified supply time, which might not always be the case. Analysis as described in Section 2.5.1 showed that no extended framework was required to include insufficient delivery performance. However, the user of the tool can apply *extra safety lead time* to include extra safety measures.

When calculating the policy parameters, future demand is forecasted using distributions which are fitted to historical demand. It is hereby assumed that the demand of individual parts is independent. In reality, however, all parts which make up one Bill Of Materials will see demand when an order is placed for the final product. As in Fluiten many parts are used for multiple final products; this was a too-extended analysis. The chosen approach will still be a considerable improvement as there currently is no forecast at all.

As already mentioned by Silver et al. (2017, p.275), several inaccuracies will influence the outcome of the inventory management (the estimates of the parameters of the distribution, estimates of cost factors etc.). More precise estimates will also improve the performance of the inventory policies. Additional information not known to the model should be included by human interference.

4.4 Chapter conclusions

In this chapter, the design of the inventory management tool is created while answering research question *'What inventory management methods are most applicable for the SKUs and how should the inventory management tool be designed?'*. Based on a tailored stepwise approach based on Hautaniemi and Pirttila (1999), historical demand and additional SKU characteristics could be used to place the parts in 6 Classes. First, a class is created with parts to manage on MRP data (5,900 SKUs). Intermittent demand is separated (134), as well as low-value items (8,920). These low-value items are split based on their estimated lead time demand (8,539 <10 and 381 >=10). The division between A-items is made based on the coefficient of variation of lead time demand (293 <0,5 and 4.363 >=0,5).

Depending on the characteristics of the SKUs in each class, different control policies were the most suitable and thus matched to the Classes. The intermittent items in Class 2 will be managed as (S-1,S), where the Normal distribution is used to estimate the size of a customer order. Low-value items in Classes 3 and 4 are managed with the (s,Q)-policy, with the GNB and Gamma approach, respectively. The A-items which are approximated with the normal distribution use the (s,S)-Normal-policy, those with the Gamma distribution are managed with the (s,S)-Gamma-policy.

An overview of the formulas used to calculate the parameters corresponding to the policies has been given. The individual TFRs for the SKUs will be determined by using either the method by Teunter et al. (2017) or van der Heijden (2024). Both will be tested on a small scale, after which one will be chosen to apply to the full problem. The tool is constrained by likely policy input inaccuracies and the dependent demand of individual parts. Additionally, incomplete and considerably late replenishment deliveries are not included. However, the model can take these uncertainties into consideration using extra safety lead time.

5 Results analysis

In this section, the solution designed in the previous chapter is tested to answer the question *'What is the performance of the inventory when applying the proposed inventory management tool?'*. This has been done using a simulation with the actual demand in 2023 and user input, as can be found in Section 5.1. The details of the simulation, such as output and test on accuracy, are described in section 5.2. As Chapter 3 suggested two different methods for determining the TFRs of individual items using an order-based approach, both are tested on a small scale in Section 5.3. Section 5.4 provides a detailed example of a part and an order in the simulation. The re-evaluation of parts for which currently no inventory is kept is performed in Section 5.5. The final numerical results are shown and analysed in Section 5.6, which highlights successes as well as possible alterations which might give improvements. Section 5.7 concludes the chapter. The last sub-question regarding robustness is answered in the following Chapter 6.

5.1 Input

Historical data up to and including the year 2022 has been used to determine the policy parameters. To test the performance of the results, a simulation will be done. Here, the chosen policies and parameters will be tested on the actual demand which occurred in 2023. The model will use 2022 as a warm-up period to avoid high backorders during initialisation. As a starting inventory, the inventory levels at the start of 2022 will be used.

The actual demand for parts of 2023 has been determined by taking all demand from the ERP. For each part, demand for the assembly, workshop and sales departments are thus summed, as this is also the way they currently manage the warehouse. The fill rates of individual parts are determined based on all this historical demand.

As the assembly of a final product cannot start before all parts are present, the order fill rate is an important performance measure as well. To be able to test this, the performance of the historical assembly work orders is analysed. Only these are selected, based on both the priorities of the company; these orders are directly meant for customers, and the analysis in Section 2.6.1, which highlighted the lowest fill rates for these orders. These 9,950 orders for 4,415 different final products from 2023 are split into the demand in parts. Looking back at the performance of each part handling the combined demand, we can determine whether this order would have been filled.

As described in the previous chapter, several user input is required next to the historical data. The values in Table 6 have been chosen in cooperation with the operational and administrative departments. The ordering costs have been approximated by dividing the labour costs of the ordering teams by the number of orders placed, resulting in €37.24. The holding costs are based on the annual taxes, which have to be paid for the level of inventory kept, which is about 20% of the purchasing/production price. As a starting scenario, no extra safety lead time is included. If required, extra lead time in weeks can be added to every single SKU.

Ordering costs	€37,24
Holding costs	20%
Extra SLT	0 weeks

Table 6: Input to the model

5.2 Simulation

The simulation to test the determined policies has been created in Spyder - Python. The exact code and explanation of the lines can be found in Appendix J. The time unit is set to be a day. The actual demand for each part is available for each exact day. For the orders, the day on which picking was started has been chosen as the day on which the order arrived. If, in reality, the delivery of a replenishment order is expected on, for example, Wednesday, the assembly team can wait with an order until it has arrived. To incorporate this flexibility in the simulation, replenishment orders are always delivered on the Monday in the week of arrival.

In the case that not all demand can be filled directly from inventory, the available parts are already reserved for this specific order. The priority of the orders with the same picking day is based on the deadline of the order. The order with the closest deadline is picked first, regardless of whether it will be complete. This choice has been made to closely resemble reality.

To determine whether an order can be filled, every part in that order has to be present. Each part is represented as an order line. It might have been the case that, when picking the order last year, part number 7 was backordered and thus picked later. Would this actual day be selected as the requested date, the model would already be working with extended picking time. The selected requested date for each order line is, therefore, the day on which the first part of the order was picked. For each order line, it is determined whether there was sufficient inventory of that part on that day.

The policies determining whether a replenishment order should be placed are based on the Inventory Position (IP). The IP is determined weekly using the following Formula 40.

$$Inventory\ Position = -Outstanding\ backorders + On\text{-}hand\ inventory + Pipeline \qquad (40)$$

5.2.1 Output

The simulation tests the calculated policies using actual historical orders of 2023. The aspects determining whether the policies are successful have been translated to KPIs. A more detailed explanation of model and the calculations of the KPIs are provided in Appendix J.

SKU-level

- **FR SKUs** - The first KPI is calculated by determining the fill rate for each SKU, which share of demand was filled from inventory, and then taking the average over all the SKUs.
- **TFR met** - As SKUs have different individual TFRs, lower obtained fill rates should not necessarily be penalised, as maybe the target was low as well. This KPI determines the share of SKUs that obtained at least their targeted fill rate. To calculate this KPI regarding the actual situation, the actual historical fill rate is compared with the Target Fill Rate used in this research.
- **VFR** - The Volume Fill Rate is determined by taking the total quantity of pieces of all SKUs (belonging to a certain class) that were filled from inventory and dividing it by the total demand. It provides insight into the performance regarding the volume of demand.
- **Average OHI** - The average number of pieces in inventory times the value of a SKU makes up the average On Hand Inventory value. High average inventory means high costs, which are undesired.
- **Backordered pieces** - The sum of all individual pieces which were not present in inventory at the time of demand are backordered pieces.
- **Orders placed** - Placing a replenishment order requires time and attention from the buying team. The costs of having to place an order have to be outweighed by other advantages. This KPI counts the total number of orders placed.

Order-level

- **FR** - The overall fill rate represents the share of orders which were completely filled from inventory on the day of arrival.

- **Backordered orders** - The total number of orders which were not complete on the day of demand is represented by this KPI.
- **FR each order** - This KPI shows how complete each order was on average. If, as shown in Section 5.4, 13 out of the 14 parts required for an order could be filled from inventory, this order was 93% complete.
- **Parts late** - If an order is incomplete, it means that some of the required parts were not present. This KPI is the average of the number of parts which were late for all backordered orders.
- **Delay** - If an order was incomplete due to a missing part, it had to wait for these parts before going to assembly. This delay represents the number of days until the late part(s) arrived and the order was complete. The KPI has been calculated by taking the average number of days of the delayed orders.

5.2.2 Verification & validation

In order to determine whether the simulation model is an accurate representation of the actual system, it should be tested whether the model is valid. Law (2015, p.246) propose a practical approach on how to do so, composed of both *verification* and *validation*.

Verification
Verification is concerned with determining whether the assumptions have been correctly translated into a computer program (Law, 2015, p.247). To test this, the model has to be thoroughly traced and debugged, making sure that the right logical paths are taken. Several SKUs have been traced to test that the simulation worked well. An example is given in Section 5.4.

Validation
Validation is the process of determining whether a simulation model is an accurate representation of the system for the objects of the study (Law, 2015, p.247). This means that the model can be used to make decisions one would make if the actual system could be experimented with. Important to note, however, is that it will always only be an approximation to the actual system. The best way to test the validation of a system is to confirm that its output data resemble the results in the actual system. In this case, however, there are no clear inventory policies. Therefore, they cannot be implemented into the simulation. Several aspects highlighted by Law (2015, p.255) have been discussed with experts, following a method called *face validity* as proposed by Silver et al. (2017, p.267).

The buyer confirmed that the right MOQs have been implemented in the simulation. The president and production manager have agreed with the decision to design the simulation on a daily basis, as it is close to the actual situation in which they orders are placed every other day. They also agree with the warm-up period of a year and the selection of assembly orders. After the addition of the KPI regarding the delay of an order, the president confirmed that all important numbers are produced as output. A limitation of the simulation, however, is that occasional early demand information, when a customer places an order significantly earlier than their requested due date, cannot be implemented as the simulation only reacts to demand when picking starts. With the examples in Section 5.4, the credibility of the model was enhanced.

5.3 Best method for individual target fill rates

The service level requirement to determine the parameters of the inventory management is expressed as fill rates for the individual SKUs. These fill rates should be chosen in a way that includes the SKU's influence on the OFR. Section 3.4 provided two methods with which this could be done. Both methods have been tested in the simulation on a small selection of the full dataset to determine which method works best for this research. This approach is then applied to the full scope. Section 5.3.1 describes the selection of the small scale and parameter setting, after which Section 5.3.2 presents the results.

5.3.1 Selection

The selection of parts and end products on which the approaches are tested has been made based on the historical orders of 2022. All end products which were produced in this year have been ranked on the number of orders they appeared in. From this list, the end products which were ordered at least 11 times were selected. This resulted in a selection of 112 of the 10,016 products, representing 15% of the total number of individual products which were sold. Using the Bill Of Materials, the SKUs which are required for the production of these products are identified. Filtering on the ones within the scope of this research leaves 531 parts.

The method described by van der Heijden (2024) requires a target order fill rate. This value is set to 93%, which was obtained with a λ of 14. The lowest individual TFR is 52%. The weighted average fill rate by the expected weekly demand of each SKU is 95%. This weighted fill rate is relevant to make a proper comparison with the second method. It has to be mentioned that this method was quite labour-intensive due to the low data quality and required a considerable number of tests to find the best value for λ.

The approach by Teunter et al. (2017) requires a weighted TFR over all SKUs as well as an expression of criticality per SKU. The TFR has initially been set to 93%, which should not be linked to the 93% used in the first method, as the value concerns individual parts for this second method. The criticality of a SKU has been determined as the number of orders this SKU was a part of. As proposed in the paper itself, a minimum is set for the individual TFRs. This minimum is 52%, which corresponds to the lowest value resulting from the first technique. The influence of these lower bounds is tested in the sensitivity analysis in Chapter 6. After setting this minimum, the obtained weighted average fill rate is 95%, matching the value from the method by van der Heijden (2024).

Both sets of TFRs have been ranked in ascending order. The corresponding line chart can be found in Figure 18. The horizontal axis does not refer to individual SKUs, meaning that there is not necessarily a SKU at the fifth percentage which has proposed TFRs of both 52% and 60%. The graph does, however, visualise the distribution of the proposed TFRs over all SKUs, highlighting among others, the higher values for the method by van der Heijden (2024). Additionally, we find that the set minimum value has restricted the value of 9% of the SKUs when using the approach by Teunter et al. (2017). Analysing the individual TFRs of the SKUs shows that the largest differences in TFRs can be found for parts which participated in a small number of orders. As only the method by Teunter et al. (2017) takes both the total demand as well as the number of orders this demand originates from into account, this explains why this method recommends lower TFRs for these SKUs.

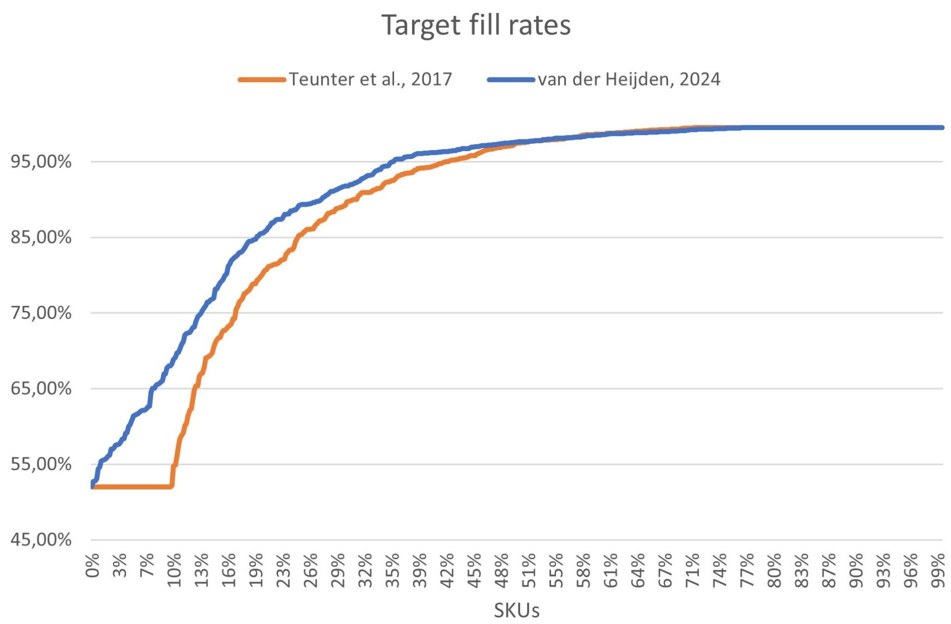

Figure 18: Target fill rates per SKU in ascending order

The two sets of TFRs have both served as input to the control policies of the individual SKUs to determine the final parameters. These parameters and policies have been tested in the simulation. The same selection of end products and SKUs has been used to narrow down the selection of orders of 2023 which are included in the test. The number of orders analysed is 1,362.

5.3.2 Results

The two sets of final inventory control parameters have been tested on the selection of demand and orders in 2023 suitable for the selection of parts and end products. The in Section 5.2.1 described KPIs have been used. Their values are shown in Table 7.

		Teunter et al., 2017	**van der Heijden, 2024**
SKU-level	**FR SKUs**	95.4%	95.4%
	TFR reached	84.9%	84.7%
	VFR	96.0%	95.7%
	OHI	€397,066	€412,548
	#BO	23,947	25,888
Order-level	**FR**	85.4%	84.1%
	#BO	199	216
	FR order	97.0%	96.6%
	#Parts late	1.5	1.5
	Delay	11.0	10.8

Table 7: Fill rate method comparison

For most of the KPIs we find slightly better values for the method by Teunter et al. (2017), but the difference in performance is often small. This method does, however, result in significantly less on-hand inventory costs. This can be linked back to the TFRs as also visualised in Figure 18, which are lower in the method by Teunter et al. (2017) for most of the SKUs. As the fill rates are similar for both methods, the target fill rates and corresponding on-hand inventory are lower for the correctly selected less important SKUs. Combined with the fact that this method was significantly more convenient to compute, the approach proposed by Teunter et al. (2017) is used to determine the fill rates for the full-scale scope. The same method of determining the criticality is used with the same overall weighted TFR of 93%, but the minimum TFR is set at 80%. This results in an obtained weighted TFR of 95%. Analysis shows that the obtained TFR varies little with changes in the lower bound (1% every circa 10% of change), while these bounds might have a large impact on the performance of a single SKU.

5.4 Case example

Detailed examples of the movements of both a part and an order are visualised in this section.

Part
A part ('4993064A3') with backorders is chosen to visualise all state possibilities. This part is classified in Class 3 and is thus managed with an (s,Q)-GNB policy. Its reorder point is 4, and the order quantity is 20. The supplier lead time is 2 weeks, but, as mentioned, replenishment orders will always arrive on the Monday of the expected week of arrival. The targeted fill rate for this part is 96%. Its daily states are visualised in Table 8 and Figure 19.

Each day begins with a starting on-hand inventory (top of the table). Several things can happen during the day, starting with the possible arrival of a replenishment order. In the simulation, this can only take place on Monday. In the case that there were still waiting backordered pieces at the start of the day, these are filled first. If there is demand for a specific part on a day, all individual pieces are either filled from inventory or backordered. Even if not the complete demand can be filled, all pieces which are present are already reserved for the incoming request. At the end of the day, the on-hand inventory is updated, taking the starting value and adding the incoming order minus past backorders and new demand. The pipeline with possible past orders is added, and unfilled demand is subtracted to find the inventory position. Based on this position, a new order is placed at the end of the day.

Day '23	**243** *Thu*	**244** *Fri*	**245** *Sat*	**246** *Sun*	**247** *Mon*	**248** *Tue*	**249** *Wed*	**250** *Thu*	**251** *Fri*	**252** *Sat*	**253** *Sun*	**254** *Mon*	**255** *Tue*
Start OHI	19	19	3	3	3	3	3	0	0	0	0	0	19
Delivery	0	0	0	0	0	0	0	0	0	0	0	20	0
Waiting BOs	0	0	0	0	0	0	0	1	1	1	1	0	0
Demand	0	16	0	0	0	0	4	0	0	0	0	0	0
Filled	0	16	0	0	0	0	3	0	0	0	0	0	0
Backordered	0	0	0	0	0	0	1	0	0	0	0	0	0
Final OHI	19	3	3	3	3	3	0	0	0	0	0	19	19
Pipeline	0	0	20	20	20	20	20	20	20	20	20	0	0
IP	19	3	23	23	23	23	19	19	19	19	19	19	19
Order	0	20	0	0	0	0	0	0	0	0	0	0	0

Table 8: Part Case example

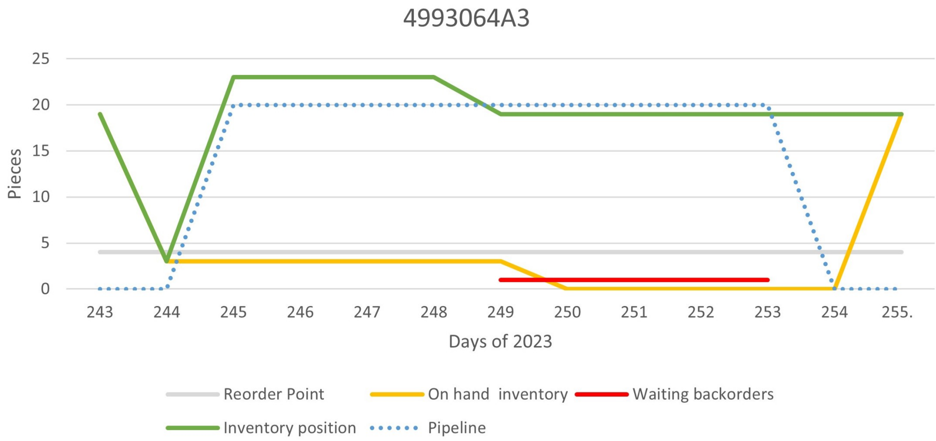

Figure 19: Part case example

In this case example, an order for 4 pieces arrived on day 249 while the starting on-hand inventory level was only three, and no new replenishment order arrived. One piece is backordered; the other three are already reserved for the incoming demand. The last piece is added after the arrival of the replenishment order on day 254. This singular backordered part is the only stockout in 2023, resulting in a fill rate of 98%, meaning that the TFR is met.

Order

An incomplete assembly order has been visualised ('W22/2341-09'). The actual picking of the order started on the 109th day of the year. Fourteen SKUs were required for this end product. The characteristics of the picking are visualised in Table 9, with the final conclusions in Table 10.

Workorder	Part	Q req.	Day	Q fill.	FR	Delay
W22/2341-09	5416019C5	100	2023-109	100	100%	0
W22/2341-09	5605016U41	100	2023-109	100	100%	0
W22/2341-09	5791019WB5	100	2023-109	100	100%	0
W22/2341-09	665585E	300	2023-109	300	100%	0
W22/2341-09	6666815V	200	2023-109	200	100%	0
W22/2341-09	7825019WDY	100	2023-109	100	100%	0
W22/2341-09	7831019EZ1	100	2023-109	100	100%	0
W22/2341-09	7832039C6	100	2023-109	46	46%	5
W22/2341-09	V100001128W44	300	2023-109	300	100%	0
W22/2341-09	V100020080F5	100	2023-109	100	100%	0
W22/2341-09	V100021105F5	100	2023-109	100	100%	0
W22/2341-09	V100031046J	200	2023-109	200	100%	0
W22/2341-09	V100061056Q	100	2023-109	100	100%	0
W22/2341-09	V200AA030E000	200	2023-109	200	100%	0

Table 9: Order Case example 1

For each order, the order lines corresponding to it are analysed and summarised. 13 of the 14 parts had enough inventory to fulfil demand, and one of them did not. As one of the 14 parts was missing, the order fill rate became 93%. However, the final product could not be assembled as the picking was incomplete. The delay of the workorder is the maximum delay of all the individual parts. The replenishment order for the incomplete part arrived after 5 days, meaning that the order suffered a delay of 5 days.

Workorder	**#Parts**	**#Fill.**	**#BO parts**	**FR order**	**Complete**	**Delay**
W22/2341-09	14	13	1	93%	NO	5

Table 10: Order Case example 2

5.5 Re-evaluation of parts without inventory

As described in Section 2.2.1, 7,947 of all SKUs in scope were currently managed on MRP-basis without constant inventory. In contrast to the parts in scope which have inventory, Fluiten has accepted having to communicate long waiting times to the customer if one of these parts is required. However, it has not recently been reevaluated whether the costs of poor performance might weigh up to the costs of keeping inventory for these parts. Now that properly founded inventory policies have been designed in case Fluiten decides to keep the parts in inventory, using the simulation, it can be analysed whether this decision would be beneficial.

Using the available historical data, we can find the actual performance of these SKUs in 2023 while they were managed based on the MRP without any safety stock. For each of these parts, the total demand, as well as the backordered demand, is available. This performance can be compared with the performance of the hypothetical situation where Fluiten keeps these parts in stock. A part (1,640) of the SKUs have already been placed in Class 1 and should thus indeed be managed MRP-based, due to previously described characteristics. The other 6,307 SKUs have been placed in other classes as if it was mandatory to keep them in inventory. Using the simulation, the average on-hand inventory and its value, as well as the backordered demand that would occur with the calculated inventory policies, are determined.

As has been concluded by Silver et al. (2017, p.371), the decision of whether an item should be stocked is influenced by many variables and goes beyond the area of production planning and inventory management. However, in order to decide for each of the many SKUs whether there should be a permanent inventory level, some generalisations had to be made, which included at least the main relevant factors on SKU-level. For this purpose, Formula 41 has been created together with the production manager using the previously calculated labour costs of backordering and the holding costs. It allows for comparison of the costs related to a single SKU.

$$Savings\ (€) = (Actual\ BOs - Simulated\ BOs) * 4.08 - Average\ OHI * v_i * 20\% \qquad (41)$$

Without the SKUs placed in Class 1, an inventory management policy was designed for 6,307 parts. 1,553 of these SKUs had demand in 2023, of which 873 actually experienced backorders and could thus have profited from safety inventory. For 248 SKUs, the savings-value following Formula 41 was positive. As, in general, the company aims to lower the inventory levels, only parts with savings of at least €10 are selected to bring into inventory. This selection consists of 63 SKUs and has been checked with the president of the company. The other 6,244 will thus not be kept in inventory but managed on MRP-basis like they are now. This results in the final division of SKUs over the classes as visualised in Table 11.

	Class 1	**Class 2**	**Class 3**	**Class 4**	**Class 5**	**Class 6**
Change	*+6,244*	*-21*	*-5,090*	*-15*	*-31*	*-1,087*
Final	12,144	113	3,449	366	262	3,276

Table 11: Division over classes in proposed solution

5.6 Simulation results

The results of the proposed inventory management, having moved the SKUs to manage on MRP-basis to Class 1, in the simulation can now be compared with the actual results of 2023. The analysis is based both on part- as well as order-level. The actual results of 2023 have been determined using the same approach as in Chapter 2, now with the full year 2023 in data.

It is important to note that there was more flexibility in obtaining the actual results of 2023 than in the simulation. Any hints for future orders, negotiations, changes in the product range or the internal lead times could not be included in the model. Additionally, when handling rare parts or high quantities, the sales office could negotiate a later delivery date to give the workshop/suppliers additional time to prepare the parts. This has happened especially often with large incidental customer orders, which occurred most frequently for the SKUs placed in Classes 2 and 3. The assembly employees could also wait to initiate picking if they know that the order is incomplete, which reduces the calculated actual delay. The results and general conclusions are presented first, with more in-depth conclusions following.

In Table 12, the results of the simulation for SKUs aggregated per class are visualised. For both the average fill rate over all SKUs and the VFR, we find an overall improvement (from 86.2% to 91.4% and from 90.7% to 95.8% respectively). Within classes, the fill rates of Classes 2 and 3 have decreased. The decrease is especially significant for the VFR of Class 3, but, as can be concluded from the merely small decrease in average fill rate, this is due to a high number of backorders for a small number of SKUs. The TFR is met for 79.3% of the SKUs (coming from 62.9%), showing significant improvements for Classes 4, 5, and 6, and a decrease in Classes 2 and 3. The average on-hand inventory value is considerably lower than last year (from €3,103,209 to €2,294,966, meaning a decrease of €808,243). The number of individual pieces which were backordered is more than halved (reduced by 54% to 49,038). 37,553 orders were placed.

		FR SKUs	TFR met	VFR	OHI (€)	#BO pieces	#Orders
Class 2	*Actual*	92.5%	85.8%	94.4%	36,790	1,458	-
113	*Proposed*	89.2%	69.9%	88.6%	30,284	2,942	280
Class 3	*Actual*	87.6%	75.8%	90.9%	352,977	4,883	-
3,449	*Proposed*	87.0%	71.2%	84.9%	98,913	8,066	6,889
Class 4	*Actual*	91.6%	31.1%	90.8%	38,919	19,019	-
366	*Proposed*	98.0%	87.2%	95.3%	17,665	9,775	1,052
Class 5	*Actual*	91.2%	31.3%	92.4%	363,142	24,718	-
262	*Proposed*	93.9%	76.3%	95.8%	144,200	13,626	4,445
Class 6	*Actual*	83.7%	54.5%	89.4%	2,311,380	58,639	-
3,276	*Proposed*	95.3%	81.6%	97.4%	2,003,903	14,629	24,887
Overall	*Actual*	86.2%	62.9%	90.7%	3,103,209	108,717	-
7,466	*Proposed*	91.4%	79.3%	95.8%	2,294,966	49,038	37,553

Table 12: Simulation Parts Results

On an order basis, as in Table 13, we find an increase in fill rates from 61.3% to 78.7% due to a reduction in the number of backorders of 1,721 (3,840 to 2,119). Each individual order is filled on average 96.2%. If an order is incomplete, it misses, on average, 1.6 parts (a decrease of 0.9), and experiences a delay of 19.2 days (an increase of 5.8 days, but it occurs less frequently).

	FR	#BO orders	FR each order	#parts late	Delay (day)
Actual	61.3%	3,840	90.9%	2.5	13.5
Proposed	78.7%	2,119	96.2%	1.6	19.2

Table 13: Simulation Orders Results

The labour costs of backorders can be approximated as costs using Formula 2. It shows in Table 14 a decrease of €26,910, reaching €20,031, as the number of backordered orders and parts has reduced.

	Unit costs	Actual #BOs	Actual costs	Proposed #BOs	Proposed costs
Orders	€2.92	3,840	€11,213	2,119	€6,187
Parts	€4.08	9,544	€38,940	3,393	€13,843
			€50,152		**€20,031**

Table 14: Backorder costs results

In conclusion, as monetary results of the new inventory management policies, we find a decrease of €807,743 in the on-hand inventory, meaning a reduction of the costs of taxes (20%) of €161,549, and a reduction in labour costs of backordering of €30,121.

5.6.1 Results analysis

Further analysis of the simulation results and background provided the following conclusions. These can be used to run a focused sensitivity analysis aimed at improving the proposed inventory management policies.

Overall, we find good results regarding the fill rates, on-hand inventory and backorders, both on a part- and order-level. All individual classes experience a reduction in on-hand inventory, but this does come at the cost of lower average fill rates for Classes 2 and 3. Further analysis shows that the set minimum TFR often bounds the TFR set for the SKUs in Classes 2 and 3. Different lower bounds can be tested to evaluate the effect on the fill rate as well as the on-hand inventory value.

Class 3 also experiences a considerable increase in backordered pieces. This is partly due to extraordinarily high demand for some SKUs (25% of the backordered pieces of Class 3 are caused by only twelve SKUs). The simulation showed that the fixed-order quantity was sometimes not enough to increase the inventory position sufficiently after a (unexpected) large order. It might be beneficial to test a variable order size on the SKUs which do not have this yet, even if it comes at the cost of more manual work.

On the other hand, there are also parts with considerably higher demand in 2023 than was forecasted based on the preceding years. This is all due to exceptionally large incidental customer orders. 43 SKUs experienced a demand of at least 20 times bigger than estimated. These parts usually had a low demand in previous years (looking into a specific SKU, for example, with a demand of 2 in 2022) and then one or a couple of big orders in 2023 (two orders of almost 100). These parts are all placed in either Class 2 or 3, which results in the lower share of SKUs, which have met their TFR. Their average fill rate is 47%. In reality, Fluiten accepts having to communicate a long waiting time to the customer. They have managed these big orders on MRP-basis and only started assembling the order when all parts were present, which explains the high *actual* fill rates.

Analysing the order performance, no significant conclusions could be made. There is no single type of order or end-product which often backorders. The increase in delay could not be explained using data from the simulation. The difference is, most likely, due to the greater flexibility in reality in setting the promised delivery date and initiating the picking process.

The results have also been discussed with the president of the company. Currently, the priority lies in reducing the on-hand inventory value and thus costs; additional (lower bound) TFRs and policies should be tested to find the optimal balance between costs and performance. Additionally, the influence of the current inclusion of undershoot for Classes 5 and 6 is tested, as well as some extra safety lead times.

5.7 Chapter conclusions

The performance of the proposed inventory management has been tested using a simulation in Python to answer research question *'What is the performance of the inventory when applying the proposed inventory management tool?'*. The simulation simulates demand on a daily basis and reserves parts for each incoming order based on the due date. The chosen policies and parameters, which have been trained using historical data up to and including 2022, are applied to the actual demand in 2023. The placement of an order is based on the inventory position. The performance of the inventory management is tested using several KPIs, among which the average fill rate of the SKUs, the VFR, the average on-hand inventory value and the order fill rate.

Using face validity, several aspects as proposed by Law (2015) have been discussed with experts to verify and validate the simulation model. Two models for setting the individual TFRs using an order-based approach were tested, and both provided promising results. On the full scale, the targeted performance fill rates for each SKU have been set using the approach by Teunter et al. (2017) where the criticality of a SKU is determined by the number of orders it appears in and the lower bound is set to be 80%. Parts which are currently managed on MRP-basis are re-evaluated to determine whether they should be kept in inventory, which resulted in the inclusion of 63 SKUs. Using a case example for both a part and an order, the movements within the model are visualised.

The results of the simulation show a significant improvement in the average fill rate of the SKUs (86.2% to 91.4%) and a decrease in backorders (from 3,840 to 2,119) while simultaneously diminishing the on-hand inventory levels (from €3,103,209 to €2,294,966). The decreased fill rates of certain SKUs could be due to the chosen lower bound when setting TFRs. Especially Classes 2 and 3 suffer from the unfair comparison, where in reality, there was more flexibility in obtaining good results than in the simulation. Possible modifications to the proposed solutions are different lower bound TFRs, variable order sizes, the exclusion of undershoot and the application of extra safety lead time, which will be tested in the next chapter.

6 Sensitivity analysis

In this chapter, several modifications to the initial proposed inventory management policy are tested to answer the sub-research question *'How robust is the tool to discrepancies in input settings and relaxations of constraints?'*. Section 6.1 introduces the different experiments and subsequently analyses the performances of each one individually. In Section 6.2, the different alterations are combined into the proposed best solution, presented with its relative performance. Finally, Section 6.3 concludes the chapter.

6.1 Modifications

Based on Section 5.6.1, some alterations to the initial solution policy were proposed. These alterations include leaving the undershoot out of scope, changing the fixed order quantity to a variable, adjusting the lower bound of the individual TFRs, and applying extra safety lead time. The changes in results when including these modifications are visualised by comparing the initially proposed solution of Chapter 5. In the following sections, the main KPIs and average values for each of these alterations are presented. All KPIs and the part-based performance aggregated per class can be found in Appendix K.

6.1.1 Without undershoot

Based on the literature found in Section 3.4, undershoot should be taken into account with non-unit-sized demand. Its characteristics have been taken into consideration when setting parameters for Classes 5 and 6. Analysis showed that reorder points are increased, but the order-up-to-level remained relatively equal, resulting in a small range between s and S and thus a likely high number of orders.

In this experiment, undershoot has been excluded from the parameters. Classes 2, 3 and 4 are not affected as the undershoot was not taken into consideration in the first place. As the reorder points are decreased, a lower fill rate, both for parts and orders, is expected. The question is how much worse the performance will be, and how large the savings. If the inventory value is significantly lower, it might be that the undershoot is taken into account for too many parts which results in parameters that are set too high.

	FR SKUs	TFR met	VFR	OHI (€)	#Orders
Overall	-1.2%	-2.9%	-0.7%	-319,933	-10,064

Table 15: Changes in parts results: without undershoot

FR	#parts late	Delay (day)
-3.5%	0.1	0.0

Table 16: Changes in orders results: without undershoot

Tables 15 and 16 show the results of the modification. We find a small decrease in the fill rate KPIs (-1.2% for the average fill rate of SKUs, -2.9% for the share of parts which have reached their TFR and -0.7% for the VFR), and very considerable decreases in the average on-hand inventory (€-319,933) and the number of orders (-10,064). The decrease in fill rates is explained by the 450 SKUs which already had some backorders and now perform slightly worse, as well as 295 SKUs which experienced new backorders. A large share of these newly backordering SKUs (59%) have a TFR equal to the set lower bound of 80%. These backorders might have been avoided with a higher lower bound for the TFR.

The average on-hand value has been reduced due to the omission of the undershoot in the reorder point, especially for items in Class 6. As these classes concern A-item, the average value of the item is high, and a small reduction in inventory quantity quickly leads to a considerable reduction in inventory value. Remarkable also is the significant reduction in the number of orders. Due to the inclusion of undershoot, the difference between the reorder point and the order-up-to-level became relatively small, resulting in very frequent small orders for Class 6. This occurs especially frequently for parts which already have a high fill rate (the average obtained fill rate of parts with a reduction of at least 5 orders is 98%), which indicates that the inclusion of undershoot here has an unnecessary negative influence.

With the reduction in fill rate on parts-level also comes a decrease in the order fill rate (-3.5%). This reduced order-based performance is due to the 295 SKUs which experience new backorders. As leaving out the undershoot provides high advantages on the on-hand inventory and the number of orders, it might be an option to exclude it and reduce the negative impact on the order fill rate using other modifications, such as increasing the lower bound on TFRs.

6.1.2 (s,S)-policy for C-items in Classes 3 and 4

For the C-items in Classes 3 and 4, an (s,Q)-policy was chosen as it is straightforward and time-efficient in the current manual way of ordering where the order quantity is calculated and entered manually. In this experiment, their policy is changed to (s,S), using Formula 34. The parameters for Classes 2, 5 and 6 are unchanged. If the results are significantly better than the initial outcomes, the additional manual effort is worth it.

	FR SKUs	TFR met	VFR	OHI (€)	#Orders
Overall	1.9%	3.3%	0.2%	16,535	-3,606

Table 17: Changes in parts results: (s,S)-policy for C-items

FR	#parts late	Delay (day)
1.6%	-0.0	-0.8

Table 18: Changes in orders results: (s,S)-policy for C-items

Tables 17 and 18 show the results of the modification. We find an improvement in the fill rates (1.9% for the FR SKUs, 3.3% for the TFR met and 0.2% for the VFR). The on-hand inventory value has increased a little with €16,535, the number of orders are reduced significantly (-3,606). This can be explained by the fixed order quantity of the (s,Q)-policy, which sometimes took multiple days to recover from big demand. Especially the SKUs in Class 3 experience an improvement in performance at a relatively low cost, as the proposed average on-hand inventory value is still only about a third of the actual value. Over 96% of the increased on-hand inventory value is distributed over about a third of the SKUs.

Improvement can also be found in the order-based performance as the order fill rate has increased by 1.6%. The delay has decreased (-0.8 days). All in all provides this adjusted policy very good results with reasonable costs. Although individual orders might take more time, significantly less have to be placed and manual labour is therefore reasonable.

6.1.3 Different (min) TFR

As was concluded in Section 5.3, the lower bound on the TFR used in the method by Teunter et al. (2017) significantly influences the overall actual weighted average TFR. In their research, they found considerable cost reductions even when setting high minimum values. In the initially proposed solution, a lower bound of 80% was set. Four different lower bounds are tested to analyse the impact, values of 70%, 75%, 85% and 90%.

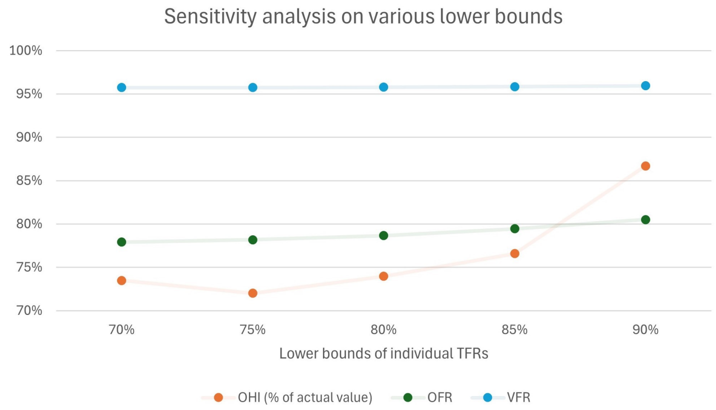

Figure 20: Sensitivity analysis on various lower bounds

Figure 20 shows the changes in the on-hand inventory level (expressed as a percentage of the current actual value), the order fill rate and the part volume fill rate. We find that the inventory costs are lowest with a lower bound of 75%. The bound of 70% does give lower reorder points, but due to the exact timing of the incoming demands and thus reordering, we keep higher inventory on average. We see an increasing order fill rate (from 77.9% to 80.5%) with the increasing lower bounds. The VFR changes marginally from 95.7% to 96.0%. These indicate that using different lower bounds does not greatly influence the order-based performance, meaning that the right SKUs had already received a higher fill rate.

After consulting the president of the company, a lower bound of 85% is chosen to profit from better order-based performance without a significant increase in costs. This increased lower bound, compared to the initially proposed 80%, provides especially improved performance for parts in Class 6 at a relatively low price.

6.1.4 Add extra safety supply lead time

In Section 2.5.1, the supply lead times of all parts are estimated. These have been either based on historical orders in the case of the parts-suppliers, or a fixed lead time of 5 weeks for the sub-suppliers and the workshop.

In the tool, it is possible to add extra safety supply lead time to the parts. This, for parts in Classes 3, 4, 5, and 6, influences their policy parameters. This is useful to include when, for example, a certain supplier communicates a temporarily longer lead time. An extra safety supply lead time of 1 and 2 weeks has been tested. This modification has changed both the policy parameters and the arrival date in the simulation. Figure 21 visualizes the changes on the on-hand inventory level (expressed as a percentage of the current actual value), the Order Fill Rate and the Volume Fill Rate.

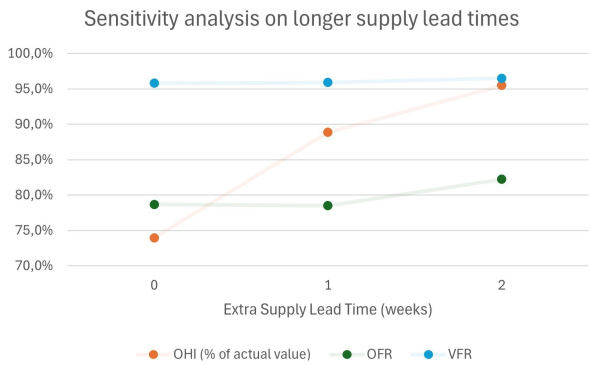

Figure 21: Sensitivity analysis on longer SLTs

We find that the inventory levels are increased significantly when adding additional SLT to all parts (from 74% to 89% and 96%). The OFR increases very slightly (from 79% to 79% and 82%), highlighting that the large added investment in inventory does not lead to a correspondingly large improved performance.

Initially, no extra safety supply lead time will be added, as the lead times used in the initial parameters are the most accurate ones we have. If, in the future, additional time is required, the tool can incorporate this. The analysis showed that additional lead time comes at a high cost regarding increased inventory value, and highlights the importance of reducing lead times when possible.

6.2 Best solution

For the optimal solution, we combine the findings from the individual modifications, as visualised in Figure 22. Section 6.1.1 showed that the inclusion of undershoot had a negative influence on Class 6, but did improve performance at reasonable costs for Class 5. Applying a (s,S)-policy for the C-items in Classes 3 and 4 had a significant positive impact on the fill rate KPIs. Increasing the lower bound of the TFR was a bit costly, especially for Class 6, but provided good improvements for both the classes as well as order performance. Additional safety lead times provided the expected results but are not necessary at this moment.

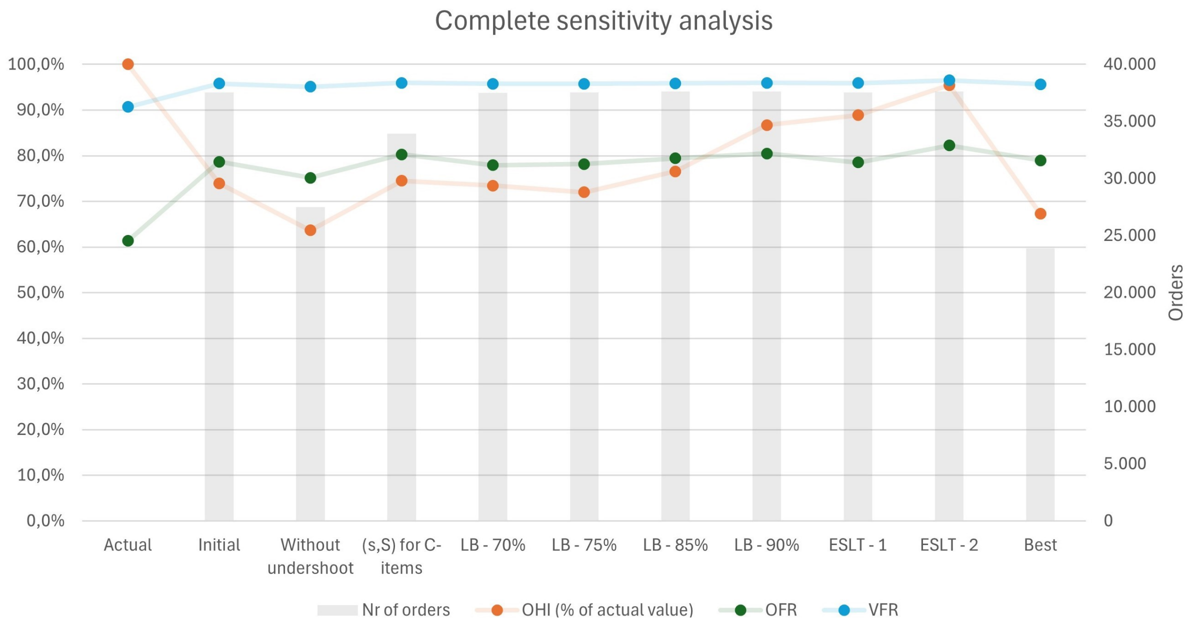

Figure 22: Complete sensitivity analysis

As the optimal solution, undershoot will be applied only to Class 5. C-items in Classes 3 and 4 are managed with an (s,S)-policy (still with the GNB and Gamma distribution). The variable for the lower bound has been discussed with the president of the company. Although the priority is still to reduce the average on-hand inventory value, in combination with the other modifications, the lower bound can be increased to avoid a high negative impact on the order-based performance. The minimum value for individual TFR is thus 85%.

	FR SKUs	TFR met	VFR	OHI (€)	#Orders
Best	93.0%	80.3%	95.6%	2,087,420	23,893
Diff. Actual	6.8%	17.4%	5.0%	-1,015,789	-
Diff. Proposed	1.6%	1.0%	-0.1%	-207,547	-13,660

Table 19: Best solution parts results

	FR	#parts late	Delay (day)
Best	79.0%	1.6	18.6
Diff. Actual	17.6%	-0.9	5.1
Diff. Proposed	0.3%	0.0	-0.6

Table 20: Best solution orders results

	Costs
Best	€19,668
Diff. Actual	€-30,484
Diff. Proposed	€-363

Table 21: Best solution backorder costs results

Tables 19, 21 and 21 show the results of the simulation using the described best policy. The values of the KPIs of this best policy are also compared with the actual situation as well as the initially proposed solution by showing the differences in values.

In line with the priority of the management of Fluiten, the on-hand inventory value has significantly decreased by about one-third of the original value (€-1,015,789). Due to the omission of undershoot for Class 6, the inventory value and the number of orders have also been reduced compared to the initially proposed solution (€-207,547 and -13,660, respectively). Both the average fill rate over the SKUs as well as the overall VFR have improved (by 6.8% and 5.0% compared to the actual situation). Especially the share of SKUs which have met their TFR has increased considerably compared to the actual historical situation (17.4%). This shows that not only the overall fill rate has increased, but this has happened for the right SKUs, which needed a higher TFR after applying the approach from Teunter et al. (2017).

This order-based approach has proven very useful, as the order fill rate is significantly increased (by 17.6% compared to the historical situation). In the case that an order is late, on average, even fewer parts are missing (0.9 fewer). However, the estimated delay is longer than the historically calculated delay by 5.1 days. As described, this is most likely due to higher flexibility in reality, by, for example, promising a later delivery date or postponing the picking of the parts. With the reduced number of backordered parts and orders comes a sizeable reduction in backordering labour costs of €30,484 compared to the actual costs.

6.3 Chapter conclusions

Research question *'How robust is the tool to discrepancies in input settings and relaxations of constraints?'* has been answered. The modifications to the initial solution, proposed in Chapter 5, have been tested to analyse the robustness of the proposed tool and policies. The exclusion of the undershoot provided better results for Class 6 by significantly reducing the average on-hand inventory value (€-303,904) as well as the number of orders (-10,064), and its negative influence on the order fill rate could be reduced using other modifications. The undershoot did perform well for Class 5. Using an (s,S)-policy for the C-items in Classes 3 and 4 improved almost all KPIs (improved the average SKU fill rate by 1.9% and the OFR by 1.6%) at the cost of a slightly higher average on-hand inventory value (€16,535). Although additional manual work is required for a single order, 3,606 has to be placed. Four different lower bounds to use in the TFR approach by Teunter et al. (2017) (70%, 75%, 85% and 90%) have been tested, which showed that the order-based approach has selected the right SKUs and especially the lower bound of 90% comes at high costs without high reward. The decision on which minimum value to choose depends on the priority of the company. If required, extra safety supply lead time can be included for parts in Classes 3, 4, 5, and 6, but this does significantly increase the on-hand inventory.

The best solution presented, chosen together with Fluiten, includes the undershoot only for Class 5, manages Classes 3 and 4 with an (s,S)-policy and uses a lower bound of 85%. The average on-hand inventory value is reduced by one-third of the original value to €2,087,420 while providing significantly better availability. The order-based approach has proven useful, as the order fill rate has significantly increased from 61.3% to 79.0%.

7 Implementation

In this chapter, a step-by-step approach is provided to answer the question *'How can the proposed inventory management tool be implemented in practice?'*. Section 7.1 introduces some theory on the implementation of changes in a business environment. Section 7.2 identifies the initial coalition of stakeholders responsible for the implementation of the new policies. The theory applied to Fluiten provides concrete steps, which are presented in Section 7.3. Section 7.4 concludes the chapter.

7.1 Implementation theory

The implementation of these new inventory policies represents a change in the current approach. In order to help employees accept and embrace changes in their current business environment, change management is necessary. It is a structured approach to shifting/transitioning individuals, teams and organizations from a current state to a desired future state (Tamilarasu, 2012). If not done properly, resistance may occur. For example, people may resist if they do not believe in the added value of the change, they feel like they had no input and experience change being imposed upon them, they are not convinced the change will succeed, or they believe it is not the right time for a change (Tamilarasu, 2012).

Kotter (1995) propose an eight-step method to transform an organization. He starts by establishing a sense of urgency and forming a powerful guiding coalition. A vision should be created and communicated, and others should be empowered to act on this vision. To avoid people giving up on the change, short-term wins should be generated. However, these should not be seen as 'victory' but as a motivation to keep tackling the problems and implementing changes. Finally, change has to be rooted in the company to avoid slowly turning back to the old behaviour.

7.2 Stakeholders

The eight-step method by Kotter (1995) is followed to ensure a structured approach to implementing the changed inventory policies. The sense of urgency is known company-wide and has been the motive for this research. The coalition of stakeholders who develop a shared commitment to the change is initially made up of the stakeholders with whom the new policies were developed. During this process, the vision regarding the requirements and priorities of the new policies was created. The main stakeholder is the owner and president of the company. During semi-weekly meetings, he has been kept up-to-date on the progress of the research as well as the direction, ensuring that the solution is fitted to the company. During the implementation of the change, he can maintain the required urgency and priority of the change, empowering the other stakeholders to act on the shared vision. The two main other stakeholders are the production manager and the purchasing manager, who contributed to the research with their detailed knowledge of current procedures as well as important characteristics of the parts and processes. In the implementation, they will move the change from a strategic to an operational level, staying in close contact with their teams.

Kotter (1995) predicts that the coalition will grow during the phases of implementing change. This means, for example, that not only the head buyer but also his team will be involved, as well as the assembly team and data specialists.

7.3 Actions

As highlighted in Section 1.3.3, advice with recommended settings and a prototype for a reusable inventory management tool were created. Additionally, a code has been written with which an analysis of the backorders can be performed. The recommended settings can be implemented using the described implementation theory.

Implementation of recommended settings
The first step before the implementation of the settings is to update their parameters (using the data connections and instructions as given in Appendix I). As described in Chapter 4, the current parameters are based on historical demand up to and including 2022 to be able to test the policies on the actual demand of 2023. As the actual parameters in the ERP should include all available historical data, the parameters was updated. This updated set of parameters is delivered to the coalition of stakeholders as one of the deliverables of this research.

As a second step, a small-scale pilot should be performed. This is an accessible way to familiarise everyone with the new changes, as well as obtain the short-term wins Kotter (1995) highlighted. For the pilot, the new inventory settings will be loaded into the ERP system for a selection of SKUs, chosen by the initial coalition of stakeholders. An information session should be held to guide the production and purchasing team through the creation, interpretation and significance of the parameters. Afterwards, they will base their ordering activities on these new policies. The assembly team is informed of the changed management for the selected SKUs and encouraged to share any experiences and observations.

After three months, the pilot should be evaluated. The same method of determining the number of backordered SKUs as was created and used in this research can be used for a quantified analysis of the performance of the newly implemented policies. An evaluation session with the coalition team, considering feedback from all involved parties, will result in a consensus on the success of the implemented changes. If the pilot is deemed successful, the policies can be loaded into the ERP system for all SKUs analysed in this research. If there were issues, the first steps should be iterated to find the cause of the problems. Possibly additional information should be included, or priorities have shifted, which requires a different overall target fill rate.

Inventory management tool
Even after the implementation of the settings for all SKUs, the coalition should keep monitoring the inventory management. New demand data and additional information must be loaded into the model regularly. This research and the prototype have highlighted the possible benefit of properly designed inventory management. If, based on the outcomes, Fluiten decides to invest in this further, professional software could be implemented. Highlighted important aspects such as the intermittent demand items, minimum order quantities and the order-based approach should be taken into account when designing this software. In the future, other departments of Fluiten such as the purchasing of raw materials, might benefit from professional inventory management as well.

Backorders analysis
Although the data was available, there was no insight yet regarding backorders. During this research, a code has been written with which an analysis of the backorders is performed, both on part- as well as on order level. This provides valuable knowledge regarding the performance of the inventory management. The production and purchasing managers can very easily perform the analysis on (selections of) any historical data. Its way of working has been explained.

7.4 Chapter conclusions

Research question *'How can the proposed inventory management tool be implemented in practice?'* has been answered. Theory on the proper implementation of change highlighted its importance and presented an eight-step method. The initial coalition of stakeholders to manage the change is identified as the president, the production manager and the purchasing manager. Based on literature, a clear stepwise approach is presented on how the proposed settings from this research can be implemented in Fluiten. The prototype of the inventory management tool has highlighted the possible benefit of implementing professional inventory management software. The KPIs regarding backorders can very easily be calculated with a created code.

8 Conclusions, recommendations and future research

Within this thesis, a prototype of an inventory management tool has been created and tested using a simulation. This is done by analysing the current situation and performing a literature study. Using insights from theory, a proposed solution was formulated. Testing several modifications in the sensitivity analysis improved the solution, resulting in a final proposal. Lastly, an implementation plan was written. In this chapter, the final sub-question is answered *'What conclusions and recommendations can be made from conducting this thesis at Fluiten?'*.

Section 8.1 lists the conclusions. The recommendations can be found in Section 8.2. Section 8.3 explains the practical and scientific contributions. The limitations and suggestions for future research are provided in Section 8.4.

8.1 Conclusions

The solved core problem is that 'currently no inventory model is applied to provide knowledge on when and how many parts to buy or produce'. Analysis of the current situation highlighted the high inventory levels, a small share of SKUs making up the majority of the annual usage value, inefficient ordering policies and incomplete assembly orders. As orders can be placed every other day, continuous review is considered. Out of the literature review, the stepwise approach from Hautaniemi and Pirttila (1999), the order-based target fill rate methods of both Teunter et al. (2017) and van der Heijden (2024), and the Normal-, Gamma- and Negative Binomial distributions are selected to be implemented in the final model. Using historical demand data and simple exponential smoothing, the mean and standard deviation of the SKUs was calculated. These were used in the formulas to calculate the inventory policy parameters. The method by Teunter et al. (2017) provided better results with less effort and has been applied to the full scope. Parts currently without inventory were re-evaluated to determine whether this would be beneficial. A sensitivity analysis tested the tool's robustness.

As expected, the exclusion of the undershoot led to lower fill rates. However, for Class 6 the reduced on-hand inventory value and number of orders were significant, indicating that currently the undershoot increased the parameters too much for a large share of the parts. As mainly parts with a TFR restricted by the lower bound experienced worse performance, it was chosen to exclude the undershoot for this class and alter the lower bound value to reduce inventory and the number of orders, but not the fill rate. Undershoot is still applied for Class 5. Working with a flexible order quantity for C-items in Classes 3 and 4 does require more manual work per order, but as the quantity of orders to place is reduced and the performance is improved, the (s,S)-policy is applied to these parts as well. The analysis on various lower bounds showed that a value of 85% provided good results, while still significantly reducing the inventory costs.

After this research, we conclude that the core problem has been solved. With the formulated implementation plan, the gap between the core problem and the action problem of 'a wrong inventory configuration for parts' can be bridged. We list the following conclusions for Fluiten:

1. Using the created inventory management tool prototype, we can construct an inventory policy for each SKU. Testing in the simulation showed that the KPIs have improved significantly. Costs are reduced, and the performance towards customers has improved due to a reduction of backorders. The delay of the backorders, of which the number has been reduced by 1,753, has increased by 5.1 days, mostly due to unfair comparison as there is additional flexibility in reality.

KPI (parts)	**Change**	**Obtained value**
Average Fill Rate over SKUs	6.8%	93.0%
Volume Fill Rate	5.0%	95.6%
Target Fill Rate met	17.4%	80.3%
Average On-Hand Inventory value	€-1,015,789	€2,087,420
Number of orders placed	*no comparison possible*	23,893

KPI (orders)	**Change**	**Obtained value**
Order Fill Rate	17.6%	79.0%
Number of parts late	-0.9	1.6
Delay of backorders (days)	5.1	18.6

KPI (labour costs)	**Change**	**Obtained value**
Labour costs of backorders	€-30,484	€19,668

2. The 63 in Section 5.5 selected SKUs should be brought into inventory as holding costs are less than costs of backordering.
3. The order-based approach by Teunter et al. (2017) using the number of orders a SKU was part of as an expression of criticality has proven to produce good order-based results.
4. To properly implement the new policies, the stepwise approach from Section 7.3 can be used.

8.2 Recommendations

Based on the execution of this research and the stated conclusions, this section provides a list of recommendations for Fluiten. The initial stakeholders will be the president, the production manager and the purchasing manager.

1. Implement recommended inventory control policies using the proposed implementation approach.

   The provided inventory control parameters for SKUs in Classes 2 to 6 should be implemented in the ERP system and methods of ordering. The inventory of parts placed in Class 1 should be removed, and these SKUs should be managed on MRP basis.

2. Improve data quality.

   The cost value and the supplier lead times of the SKUs in the ERP system should be made accurate for all parts. The SKU-codes which have been merged or transferred should be connected properly to avoid data loss, after which the old codes should be removed. Some SKUs are used internally and are taken from inventory in large batches, resembling intermittent demand. These internal movements should be separated from customer demand.

   These improvements have been, to the best of our efforts, performed manually in this research. Structurally fixing the issues will provide more trustworthy and accurate results.

3. Implement the created performance KPIs.

   When possible, the chosen KPIs have been determined on historical data to determine the current performance. They provide valuable insight into the performance, which was not available before. Implementing the KPIs enables constant insight into all (future) available data.

4. Improve communication between sales and production/purchasing teams.

   Currently the production/purchasing teams have no insight into the sold products until they receive a message from the ERP system that a certain part has surpassed the reorder point. With some form of communication regarding the quantity and type of products/parts (to be) sold, large upcoming orders will be anticipated and sudden significantly reduced inventory will be known.

5. Investigate automating purchasing/production for non-critical SKUs

8.3 Practical and scientific contribution

This section describes the practical and scientific contribution of the research.

8.3.1 Practical contribution

This research is performed at Fluiten. The practical contribution to the company consists of several products. Firstly, a thorough analysis of the current situation has been provided, including the identification of obsolete parts and a detailed performance on both parts and order levels. Using the approach visualised in Figure 16, any (future) parts can be classified easily. A prototype of a tool is created with which the parameters of the inventory management policies based on historical demand can be calculated. Any future changes to the inventory policies can be detailedly tested using the created simulation model.

8.3.2 Scientific contribution

As found in the performed literature research, Hautaniemi and Pirttila (1999) propose a stepwise approach to classify SKUs using Classes A and C from the ABC-analysis. In this research, their approach has been adjusted. The initial separation of A- and C-items remains to focus on the importance of a part, but two different decision points are introduced to simplify the step from a class to its parts' demand distribution. The generalised negative binomial distribution, which enables working with real-valued parameters, has successfully been applied to the slow movers using a target fill rate approach.

A newly proposed order-based performance approach to choosing individual target fill rates by van der Heijden (2024) has been compared with the approach by Teunter et al. (2017) where criticality is expressed in the number of orders a SKU is part of. Compared to the published paper, additional lower bounds for the individual TFRs have been tested.

The thesis can be a case example of how different policies, distributions and input parameters perform on a large-scale production-inventory model.

8.4 Limitations and future research

Due to a bounded complexity and scope, the research has some limitations. Improving these, as well as other research directions, are proposed as future research.

- As there are currently no clear inventory policies or parameters, the current situation could not be modelled in the simulation. The 'actual' KPIs have been calculated based on historical data. However, this comparison is not fully fair, as additional efforts to obtain the observed results were impossible in the simulation. After implementing the new policies, the newly generated historical data should be used to compare the KPIs again.
- As mentioned in the recommendation to improve data quality, quite some data cleaning was required to come to the parameters for the demand distributions. With improved

data quality and a more complex forecasting method, the parameters of the policies will be better.

- The supply lead times for parts-suppliers have been approximated based on the recommendation of the head of purchasing. Due to the high variability in supply time for parts from the workshop and sub-suppliers, a fixed lead time is used. If this variability can be reduced and individual lead times can be assigned, the parameters of the policies will improve for these produced parts. As the sensitivity analysis showed, having to work with longer lead times significantly increases inventory holding costs.
- The possibility of substituting certain parts with others has been left out of scope. Future research including this option might reduce inventory costs without harming order-based performance.
- All types of orders and end-products have received the same priority, while in reality, Fluiten has customers with a special contract. Future research can include the higher importance of the parts of orders for these customers.
- The order-based approach for setting the TFRs of individual SKUs by van der Heijden (2024) has provided good results, but, both due to limitations in data quality as well as more demanding computations, required high efforts. More research into the usability and performance of this method is necessary to determine whether it provides good results.

References

Agrawal, N., & Smith, S. A. (1996). Estimating Negative Binomial Demand for Retail Inventory Management with Unobservable Lost Sales. *Naval Research Logistics (NRL)*, *43*(6), 839–861. https://doi.org/10.1002/(SICI)1520-6750(199609)43:63.0.CO;2-5

Al-Ahmadi, S., & Yanikomeroglu, H. (2010). On the approximation of the generalized- distribution by a gamma distribution for modeling composite fading channels. *IEEE Transactions on Wireless Communications*, *9*(2), 706–713. https://doi.org/10.1109/TWC.2010.02.081266

Axsäter, S. (2015). *Inventory Control Third Edition*. Springer.

Boute, R. N., Disney, S. M., Lambrecht, M. R., & Van Houdt, B. (2006). A smoothing replenishment policy with endogenous lead times. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.814428

Boylan, J. E., & Johnston, F. R. (1994). Relationships between Service Level Measures for Inventory Systems. *The Journal of the Operational Research Society*, *45*(7), 838–844. https://doi.org/10.1007/978-1-84800-011-7{_}20

Cohen, M. A., & Ernst, R. (1988). Multi-item classification and genereic inventory stock contr. *Production and Inventory Management Journal*, *29*(3).

Croston, J. D. (1970). Forecasting and Stock Control for Intermittent Demands. *Source: Operational Research Quarterly*, *23*(3), 289–303. https://doi.org/10.2307/3007885

D.E. Nicol. (1989). An innovative two-bin application for floor stock. *Production & Inventory Management*, *30*(2), 27–29.

Dhoka, D., & Lokeswara Choudary, Y. (2013). ABC Classification for Inventory Optimization. *IOSR Journal of Business and Management*, *15*(1), 38–41. https://doi.org/10.9790/487X-1513841

Feeney, G. J., & Sherbrooke, C. C. (1966). The (S-1,S) Inventory Policy Under Compound Poisson Demand. *Management Science*, *12*(11), 391–411. https://doi.org/https://doi.org/10.1287/mnsc.12.5.391

Fluiten. (2023). Fluiten. https://www.fluiten.it/

Gutierrez, M., & Rivera, F. A. (2021). Undershoot and order quantity probability distributions in periodic review, reorder point, order-up-to-level inventory systems with continuous demand. *Applied Mathematical Modelling*, *91*, 791–814. https://doi.org/10.1016/j.apm.2020.09.014

Hautaniemi, P., & Pirttila, T. (1999). The choice of replenishment policies in an MRP environment. *International Journal of Production Economics*, *59*(1-3), 85–92. https://doi.org/10.1016/S0925-5273(98)00026-7

Heerkens, H., & van Winden, A. (2017). *Solving Managerial Problems Systematically*. Noordhoff Uitgevers.

Jain, G. C. (1971). A Generalized Negative Binomial Distribution. *Consul Source: SIAM Journal on Applied Mathematics*, *21*(4), 501–513.

Johnson-Fluiten. (2023). Johnson-Fluiten. https://johnson-fluiten.com/en/

Kenton, W. (2023). Leadtime. https://www.investopedia.com/terms/l/leadtime.asp

Kiran, A. ! S., & Loewenthal, A. (1985). An Integrated Decision System for Inventory Management. *Computers & lndustrial Engineering*, *9*(4), 379–386. https://doi.org/https://doi.org/10.1016/0360-8352(85)90025-7

Kotter, J. P. (1995). Leading Change: Why Transformation Efforts Fail Harvard Business Review. *Harvard Business Review, March-April*.

Kumar Dhoka, D., & Lokeswara Choudary, Y. (2013). "XYZ" Inventory Classification & Challenges. *IOSR Journal of Economics and Finance*, *2*(2), 23–26. https://doi.org/10.9790/5933-0222326

Larsen, C., & Thorstenson, A. (2008). A comparison between the order and the volume fill rate for a base-stock inventory control system under a compound renewal demand process. *Journal of the Operational Research Society*, *59*(6), 798–804. https://doi.org/10.1057/palgrave.jors.2602407

Larsen, C., Seiding, C. H., Teller, C., & Thorstenson, A. (2008). An inventory control project in a major Danish company using compound renewal demand models. *IMA Journal of Management Mathematics*, *19*(2), 145–162. https://doi.org/10.1093/imaman/dpm036

Law, A. M. (2015). *Simulation Modeling and Analysis, Fifth Edition*. McGraw Hill. www.averill-law.com

Mely Permatasari, P., Yanuar Ridwan, A., & Santosa, B. (2017). Inventory Policy Determination for Raw Materials in ILY Pharmaceutical using Periodic Review (R, s, S) and Periodic Review (R, S) Method to Minimize Total Inventory Cost. *MATEC Web of Conferences*, *135*. https://doi.org/10.1051/matecconf/201713500056

Mohamadghasemi, A., & Hadi-Vencheh, A. (2011). Determining the ordering policies of inventory items in class B using If-Then rules base. *Expert Systems with Applications*, *38*(4), 3891–3901. https://doi.org/10.1016/j.eswa.2010.09.050

Moinzadeh, K. (1989). Operating Characteristics of the (S-1,S) Inventory system with Partial Backorders and Constant Resupply Times. *Management Science*, *35*(4), 472–477.

Noblesse, A. M., Boute, R. N., Lambrecht, M. R., & Van Houdt, B. (2014). Lot sizing and lead time decisions in production/inventory systems. *International Journal of Production Economics*, *155*, 351–360. https://doi.org/10.1016/j.ijpe.2014.04.027

Olhager, J. (2010). The role of the customer order decoupling point in production and supply chain management. *Computers in Industry*, *61*(9), 863–868. https://doi.org/10.1016/j.compind.2010.07.011

Pandya, B., & Thakkar, H. (2016). A Review on Inventory Management Control Techniques: ABC-XYZ Analysis. *REST Journal on Emerging trends in Modelling and Manufacturing*, *2*(3).

Partovi, F. Y., & Burton, J. (1993). Using the Analytic Hierarchy Process for ABC Analysis. *International Journal of Operations & Production Management*, *13*(9), 29–44. https://doi.org/10.1108/01443579310043619

Presutti, V. J., & Trepp, R. C. (1970). More ado about economic order quantities (EOQ). *Naval Research Logistics Quarterly*, *17*(2), 243–251. https://doi.org/10.1002/nav.3800170211

Scholz-Reiter, B., Heger, J., Meinecke, C., & Bergmann, J. (2012). Integration of demand forecasts in ABC-XYZ analysis: Practical investigation at an industrial company. *International Journal of Productivity and Performance Management*, *61*(4), 445–451. https://doi.org/10.1108/17410401211212689

Silver, Bischak, D., & De Kok, T. (2012). Determining the reorder point and order-up-to level to satisfy two constraints in a periodic review system under negative binomial demand. *Source: The Journal of the Operational Research Society*, *63*(7), 941–949. https://doi.org/10.1057/palgrave.jors.2602655

Silver, Pyke, D., & Thomas, D. (2017). *Inventory and Production Management in Supply Chains Fourth Edition*. Taylor & Francis Group.

Silver & Rahnama, M. R. (1987). Biased Selection of the Inventory Reorder Point when Demand Parameters are Statistically Estimated. *Engineering Costs and Production Economics*, *12*(1-4), 283–292. https://doi.org/10.1016/0167-188X(87)90089-9

Song, J. S. (1998). On the order fill rate in a multi-item, base-stock inventory system. *Operations Research*, *46*(6), 831–845. https://doi.org/10.1287/opre.46.6.831

Syntetos, A. A., & Boylan, J. E. (2005). The accuracy of intermittent demand estimates. *International Journal of Forecasting*, *21*(2), 303–314. https://doi.org/10.1016/j.ijforecast.2004.10.001

Tamilarasu, V. (2012). Change Management. *International Journal of Management Prudence*, *3*(2), 26–31.

Teunter, R. H., Syntetos, A. A., & Babai, M. Z. (2017). Stock keeping unit fill rate specification. *European Journal of Operational Research*, *259*(3), 917–925. https://doi.org/10.1016/j.ejor.2016.11.017

Teunter, R. H., Babai, M., & Syntetos, A. A. (2010). ABC Classification: Service Levels and Inventory Costs. *Production and Operations Management Society*, *19*(3), 343–352. https://doi.org/10.3401/poms.1080.01098

T.H. Willis & J.D. Shields. (1990). Modifying the ABC inventory system for a focused factory. *Industrial Engineering*, *22*(5), 38–39.

Tibben-Lembke, R. S. (2006). Average and Standard Deviation of Demand over Replenishment Lead Time.

Tyworth, J. E., Guo, Y., & Ganeshan, R. (1996). Inventory control under gamma demand and random lead time. *Journal of Business Logistics*, *17*(1), 291–301.

Van der Heijden, M. (2021a). 04c. Forecasting - Fast movers – Estimating standard deviation of forecast error over lead time.

Van der Heijden, M. (2021b). 08c. Undershoot.

Van der Heijden, M. (2021c). 09b. The (s, S) inventory control system.

van der Heijden, M. (2024). Simple approximation Order Fill Rate when parts are managed individually. *Unpublished manuscript*.

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

Wild, T. (2002). *Best Practice in Inventory Management*. Elsevier Science Ltd.

Yan, D., Ahmad, S. Z., & Yang, D. (2013). Matthew effect, ABC analysis and project management of scale-free information systems. *Journal of Systems and Software*, *86*(2), 247–254. https://doi.org/10.1016/j.jss.2012.08.013

Appendices

Appendix A Order process explained

The majority of the need for parts comes from arriving customer orders. To understand the flow from an order to a finished product, the scheme visualised in Figure 23 is described.

The process starts when the client sends an email to Fluiten, which the sales office receives. This mail can be a simple order of a standard existing product or an offer request. If it is an order, the sales office checks whether the information is complete and enters it in the ERP. If the customer has requested an offer for an existing product, the sales office will determine the quotation based on existing prices. They might also request a single spare part. However, if it is a tailored request, the design department will first evaluate the feasibility of the application. A quotation is determined with the sales office, which is sent back to the client. If they accept the quotation, the tailored product is designed and entered in the ERP. The accepted quotation for an existing order can be put in the ERP right away.

The production office then checks the order. Depending on the availability of the parts required for the order, the due date requested by the customer is confirmed or postponed. If missing parts are not in the pipeline yet, they create a production order for the workshop, or the purchasing team will place an order. If the urgency of the order is high, they might prioritise already planned production of the parts or ask the purchasing team to try to push the delivery date of purchased parts closer. The final scheduled due date is communicated to the client by the sales office.

After checking the order, the production office will create a pick-and-assemble form. Every morning, these papers are given to the warehouse team. The picking department orders the forms by date. If the order's due in about ten days, they start collecting the parts from the warehouse. Only the orders for special customers receive high priority and are always started immediately. The shelves with the relevant parts are automatically presented by the warehouse one by one, from which the picking employee collects the parts in a blue box. If not all parts required for the order are in the warehouse, the box is placed in a waiting station with the pick-and-assemble paper. The box is brought to the assembly team if the order is complete.

Incoming parts from the sub-suppliers, parts-suppliers or the workshop are all collected at the same station in the warehouse. The warehouse employee scans the barcodes corresponding to the parts, after which the MRP notifies which orders in the waiting station need these parts. The employee then adds the new parts to the blue box and brings it to the assembly team when the order is complete. If the MRP does not show any waiting orders, or if there are more incoming parts than are needed right now, the parts are added to the warehouse.

Every morning, the assembly manager determines the order in which the waiting products will be assembled based on their due date. The parts are combined into finished products and then tested using high-pressure tools. If a product does not perform well, the assembly employee tries to discover the problem, which might require improvements from the workshop team. After the problem is found and solved, the product is brought to the shipping department, which packs the product and provides relevant labelling. The client is notified that the product is ready, after which a transportation company can pick it up. The process is then finished.

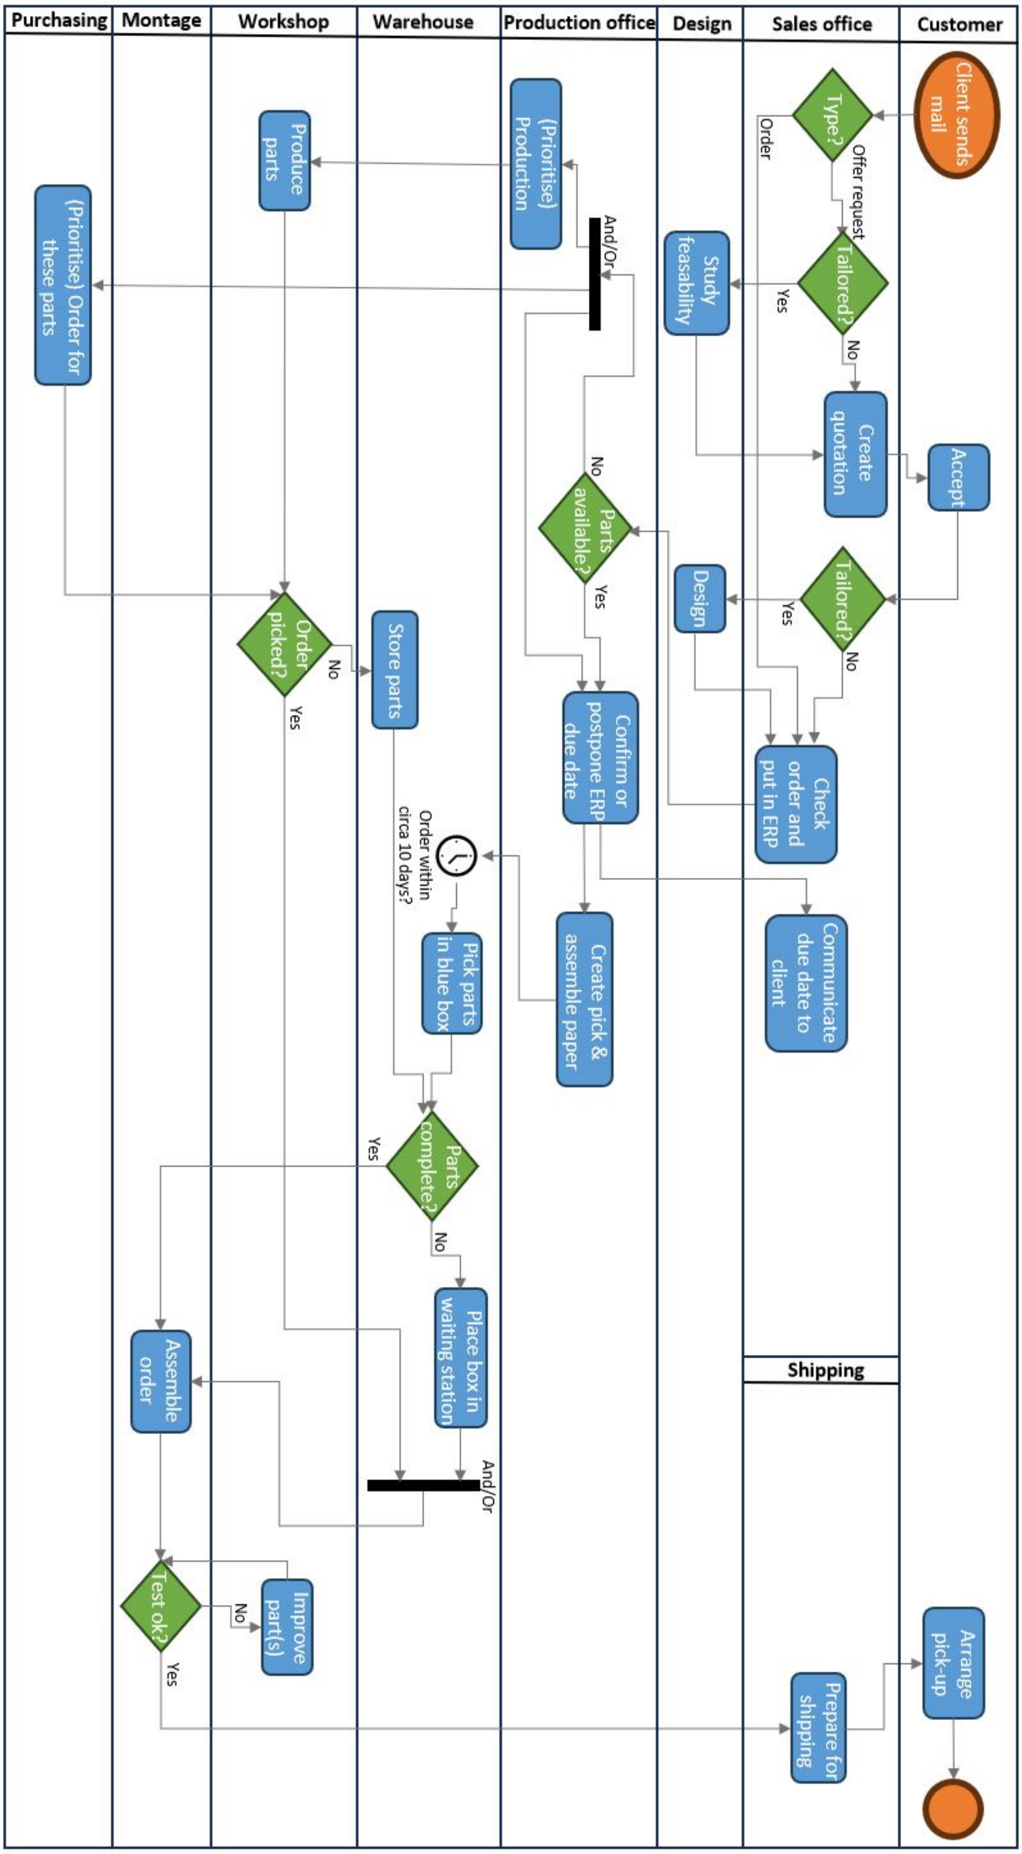

Figure 23: Order Process

Appendix B Data adjustments

The data from the ERP system has been used to make the desired analyses. The data selection has been adjusted to represent the current situation the best.

Regarding the selection of SKUs, all standard products and parts which have had any movement since the set-up are considered. As the SKUs have different types of backgrounds, it is hard to confidently assign a cost value to each part. Fluiten works with three different types of costs in the ERP, material, labour and service costs. In the case of purchased parts, the costs can easily be determined and stored under the material category. They have tried to determine the costs of labour and service for produced parts, adding the original raw material costs as well. In this research, the sum of the three categories has been taken as cost value. Although this calculation might not be fully correct, it will at least provide an insight into the value.

Distribution By Value
In the DBV analysis, the year 2020 is not analysed as detailed as the surrounding years. In 2020, COVID influenced the company. Some markets had less demand (naval industries), while in others, there was an increase in the number of orders (pharmaceutical industry). This also impacted the demand for each product and thus part. As the company did not fully close or completely change its way of working, the year 2020 is not entirely left out of the research, but it does not add value to intensively analyse it with a DBV.

Inventory Coverage
In the IC analysis, the expected usage rate (in units/year) greatly influences the final value of the coverage. In this general analysis, this usage rate was determined by taking the average of the years 2017-2023.

2016 has been left out, as only half of the year is recorded in the current ERP system. The records of these six months are already eight years old and are thus not essential in determining the usage rate. By removing this year, we also removed 476 SKUs whose latest demand was in 2016. However, none of these SKUs have any inventory currently, so it will not influence their coverage.

The year 2023 is critical, as it represents the most recent data, and some (for example, tailored) parts are only introduced this year. However, at the time of the analysis, only the demand until September 4th (the 247th day of the year) is known. Based on a year of 365 days, the quantity has been increased by 365/247. As the summer break has been included in the known data, the upcoming Christmas break will be somewhat covered when pasting the past months into the next few.

Appendix C Distribution By Value

	2023 (until 4-9)		2022		2021		2019	
	Material	*MLS*	*Material*	*MLS*	*Material*	*MLS*	*Material*	*MLS*
With demand	37%	37%	51%	51%	41%	41%	39%	39%
Without demand	63%	63%	49%	49%	59%	59%	61%	61%
With cost	91%	94%	91%	94%	91%	94%	91%	94%
Without cost	9%	6%	9%	6%	9%	6%	9%	6%
Demand no cost	3,3%	3,0%	3,6%	3,0%	3,3%	3,0%	3,4%	3,0%
No demand/cost	5,1%	3,4%	5,1%	3,4%	5,1%	3,4%	5,1%	3,4%

Table 22: Analysis DBV 2019-2023

Multiple aspects are considered to obtain a complete overview of the DBV. The analysis has been performed for 2019, 2021, 2022 and 2023 (up until September 4th), as shown in Table 22. The year 2020 is left out due to COVID; see Appendix B. Going back in time, the share of SKUs with demand decreases. This makes sense for the way the analysis is set up. As we consider all SKUs which have had any form of 'movement' in the last seven years, we include the SKUs introduced recently. If a part was introduced in 2021, it did not have any demand yet in 2019.

A difficulty is determining the value of a SKU. In the ERP system, a monetary value can be entered for the material, labour and service costs as described in Appendix B. These values are not always very trustworthy anymore. Some SKUs do not have a cost value. The division between SKUs with or without a cost value is based on the current ERP and thus constant in the analysis over the years. It is, however, essential to use it to determine the share of SKUs which did have demand in a specific year but no cost value, as their demand is not contributing to the annual usage. We find a share of 3.0%, consistent over the years as these SKUs are precisely the same ones over the years. Thus, they do not influence the comparison between the years. Further investigation and review with the production manager tells us that they are all small parts, such as screws or springs.

	2023 (until 4-9)		2022		2021		2019	
	Material	*MLS*	*Material*	*MLS*	*Material*	*MLS*	*Material*	*MLS*
SKUs for 80%	5.2%	6.1%	4.9%	7.2%	5.0%	9.2%	5.2%	6.2%
20% of SKUs	98%	98%	97%	95%	98%	93%	98%	98%

Table 23: 80/20 values DBV 2019-2023

Throughout the years, a very low percentage of the total number of SKUs already contributes a substantial share to the total annual usage (as can be concluded from Table 23 and Figure 24). We now consider the year 2022 - Material Labour Service for the full DBV analysis, as this is the most recent year with complete data and provides a cost value to more SKUs. The 3.0% of SKUs with demand but without a cost value, as mentioned earlier, only contribute 0.8% to the total demand in parts sold. As they are small parts, for this current DBV, it is not a problem to ignore the fact that they have no cost.

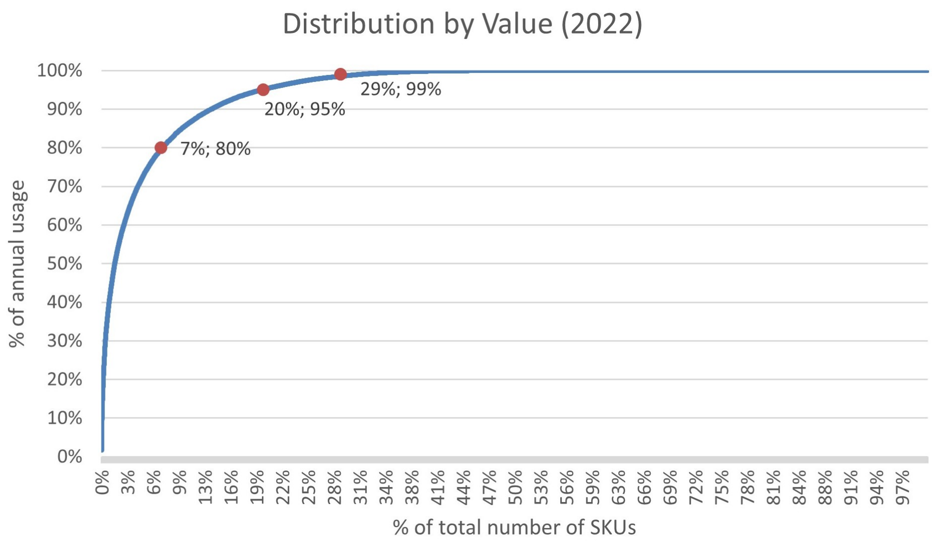

Figure 24: Distribution By Value 2022 - Material, Labour and Service costs

In 2022, 51% of the SKUs had demand. We then find that with 7.2% of the SKUs, 80% of the annual usage is reached. With 20% of the SKUs, we already reach 95% of the annual usage, whereas 29% is enough for 99% of the usage. Another 19% (reaching 48% of the total, missing the previously mentioned 3% with demand but without cost value) of SKUs represent the final 1%. These 28% of SKUs (from 20-48%) are either low- to medium-expensive parts of which only a small quantity is sold, or inexpensive parts of which demand was high. There were, for example, 273 SKUs with a cost value of less than 2 euros, of which the demand was higher than 500 units, contributing to the annual usage with 2.20%. These inexpensive parts could easily be kept in inventory to avoid backorders without resulting in too many costs.

This DBV confirms the observed problem of having many parts with many different characteristics and demand patterns. We can use this analysis to make a division between the SKUs and suitably approach them.

Appendix D Inventory Coverage

At the time of these calculations, only the on-hand inventory of Tuesday, September 26th, 2023, was available. Although an average inventory level would be more representative of the inventory throughout the year, the snapshot will suffice for this general analysis. In Appendix B, the data processing is described. The expected usage rate is determined using the average demand of the last seven years as a starting point. When building the improved inventory model, a more proper forecasting method should be used, but for now, it will do.

Of the total 19.610 SKUs in scope, 6,786 had inventory at the time of the analysis (35%). The IC analysis gives the result which can be found in Figure 25.

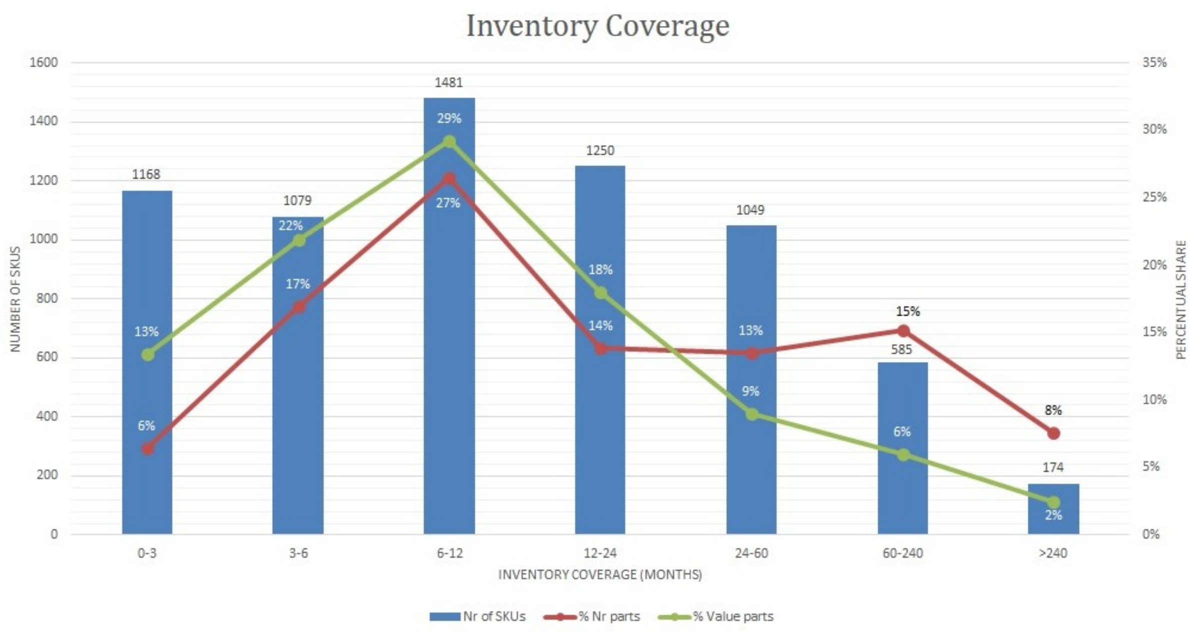

Figure 25: Inventory Coverage

On the x-axis, the IC in months is represented in relevant buckets. The blue columns represent the number of SKUs with coverage within this bucket. The red line represents the share of the number of parts in the inventory in the corresponding bucket, while the green one represents the share of the value of the parts. Looking at the figure shows that the majority of the parts have a long IC. This is confirmed by the median being 10.2 months. Many SKUs are covered for at least the next two years. The percentual calculations provide more insight into the characteristics of the SKUs. When looking at the division of the number of parts, we find a peak (27%) in the 6-12 months bucket, which is a long time, but compared to the long lead times, still reasonable. The peak of the value of the parts is in the same bucket (29%).

There are 274 SKUs with an IC of less than three weeks (left tail). The division between purchased or fabricated SKUs is pretty equal (116 to 158) and thus does not have an influence on this low IC. This coverage could be considered too low depending on the individual characteristics, such as lead time or lot size.

On the other tail, we find that 23% of the parts correspond to SKUs with enough inventory for at least five years (15% for 60-240 months and 8% for over 240 months). When translating this to the inventory value, these parts contribute with 8%. This could be due to low value SKUs or an inventory level resulting in a high IC for the SKU but which is relatively low compared to other parts. The total value is then still minimal. A reason for this high IC but low inventory level might be because a SKU was only introduced last year. Then, the current expected usage rate determined by using the average might be far too low. For this reason, the parts with inventory for over 240 are analysed further. Of the 174 parts, 58 have only had demand in 2023. However, if we adjust the expected usage rate to the demand of 2023 (modified for an entire year), all SKUs still have an IC of at least three years, with the median being 11.3 years. Even with this adjustment, their IC values are very high and belong in the top two buckets. Checking the ten individual SKUs with the highest IC confirms that they are dead stock. They have all had one input, the date on which the new ERP is set up, and minimal usage throughout the years.

In conclusion, the presumption that many parts have a too high inventory level can be confirmed by this analysis.

Appendix E Occurence of backorders

We analysed the picking orders with a due date from January 2023 to September 2023, the most recent complete months. Important to note is that the parts for these orders need not be picked in this exact period but could be collected in the month(s) before. The analysis is done on both picking order-level and SKU-level.

SKUs
Of the 19,610 SKUs in scope, 8,642 were needed for the work- (picking) orders in the selected time frame. Figure 26 shows the fill rate of each SKU with demand. The fill rate represents the share of the demand, when a certain SKU is required for a workorder, which was present in inventory at the initial moment of collection. We find that 11% had a fill rate of zero, meaning that these SKUs always had to be added to the rest of the order at the moment of their arrival. Luckily, these SKUs only represent 1% of the total units picked. On the other side of the S-curve, we find 61% of the SKUs with a fill rate of 100%. They represent 49% of the units picked. The average over all SKUs is 81%.

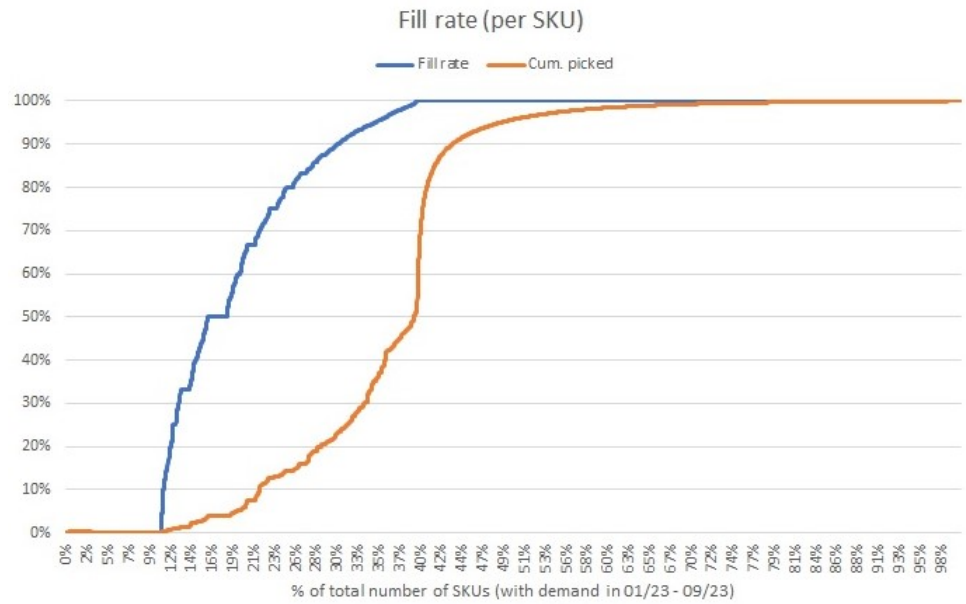

Figure 26: Fill rate on SKU level

The steep curve in the number of picked parts around 40% of the SKUs can be explained by the ordering of the parts, which all had a 100% fill rate. The SKUs with high demand in the period are placed first, resulting in a Pareto analysis within the 100% fill rate category. This is not disruptive to the analysis, as they all have a fill rate of 100%.

Work orders

In total, there were 15,299 picking orders in the selected nine months. 7,147 were picking orders for the workshop, 326 for service requests and 7,826 for the assembly department. As shown in Table 24, these orders vary significantly in size, as the picking orders for assembly consist of entire products (on average 11.4 parts), whereas those for service or the workshop are of specific parts only (1.6 and 1.4 parts on average, respectively).

Type of picking order	Average number of SKUs
Workshop	1.6
Service	1.4
Assembly	11.4

Table 24: Average number of SKUs in a picking order

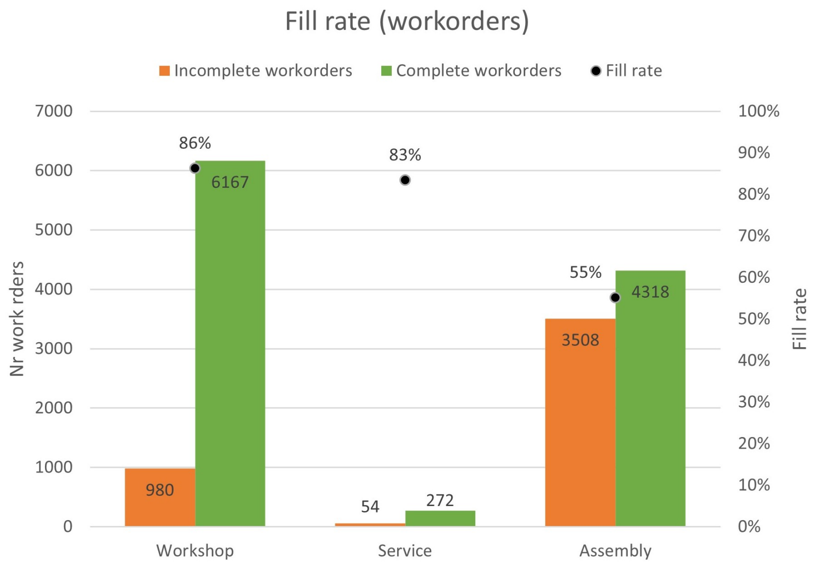

Figure 27: Fill rate on workorder level

In the first analysis, we determine the fill rate per workorder, indicating which share of the orders of each department is filled directly from inventory. We find a fill rate of 86% for the workshop, 83% for service and 55% for assembly picking orders (see Figure 27).

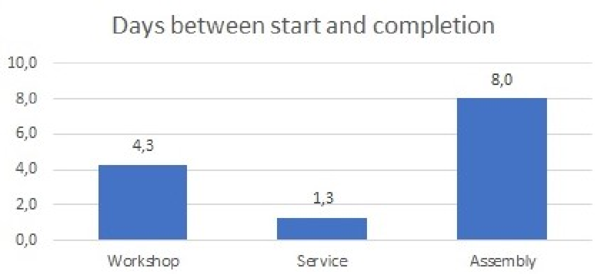

Figure 28: Average number of days between start and completion of picking per type of order

On average, it takes 4.3 days for a workshop order to go from the start of picking until completion, as Figure 28 shows. Service orders take 1.3 days, whereas assembly orders, on average, stay 8.0 days waiting for the last SKUs.

An analysis of each order provides insight into the level of incompleteness of each waiting order (see Figure 29). We find that if an order is incomplete, the fill rate for the order is, on average, 51% and 52%, respectively, for workshop and service orders. These are low rates, which the small size of these orders can explain. We find that, on average, only 1.2 and 1.0 SKUs are missing for the workshop or the service department, respectively. With a small total required number of parts, even this one missing part has a high influence on the fill rate. For assembly orders, we find that, on average, 2.8 SKUs are missing, resulting in a 78% fill rate for incomplete orders.

If we combine incomplete orders with the 100% fill rate of complete orders and take the average for each type of order, we find a 93% rate for the workshop. This means that of all the picking orders for the workshop, 93% of the parts were in inventory at the initial picking moment. For service, we find an average of 92% and 90% for assembly. These percentages must be improved, especially for the assembly-department, to improve the fill rate of final workorders.

Figure 29: Fill rate within individual orders

Appendix F Labour costs of backordering

The individual events should be analysed to determine the costs of a backorder.

In this research, we focus on the backorders for the assembly department, meaning the moments in which the due date is set, the final assembly has started, and still not all parts are complete. The work of the purchasing and production offices and possible additional rush service costs are thus not included.

The final steps start with the picking team. They collect all available parts and get notified by the MRP if some parts are missing. They place the blue box with the incomplete parts on a separate shelf, with the missing part(s) marked on the work order paper. They contact the production manager in the rare situation that the part is not there and not ordered. The box with the already picked parts stays on the shelf until the missing part(s) arrives(-) at the warehouse department. The responsible employee links the arrived part to the waiting orders and completes them. The blue box is then carried to the assembly department.

This does not cost too much time per order. However, as many orders are picked daily, and a large share of those have missing parts, the wasted time adds up quickly. To give a monetary value to the labour of a backorder, the previously mentioned actions have been timed and quantified, as visualised in Tables 25, 26 and 27. With these numbers, a formula can be created to estimate the impact of a backorder.

Costs per hour	€35
Costs per minute	€0.58

Table 25: Labour costs picking team

For each backorder		Minutes	Cost
Picking employee	Place blue box in waiting station	2	€1.17
Incoming goods employee	Blue box to assembly	3	€1.75
		Total per backorder	€2.92

Table 26: Backorder costs: per backorder

For each backordered SKU		Minutes	Cost
Picking employee	Check if SKU is ordered	2	€1.17
Incoming goods employee	Match incoming SKU to incomplete order	5	€2.92
		Total per backorder	€4.08

Table 27: Backorder costs: per backordered SKU

This results in the following formula to calculate the labour costs of each backorder;

$$Backorder\ costs = €2.92 + €4.08 * Nr\ backordered\ SKUs \tag{42}$$

Appendix G Intermittent demand analysis

The demand of the SKUs in 2021, 2022 and 2023 up until September are analysed. First, the data is cleaned. 5,325 SKUs were left out as they had no demand, 3.655 had only one demand point and thus no intermittent periods and 12 are not represented well in the data due to internal procedures. This left 10,618 SKUs for the analysis.

The analysis is done on a monthly basis. If a SKU had demand in two consecutive months, we determined the time between demands to be zero (months). Additionally, the average demand of all SKUs placed in a specific bin during the selected period is calculated. We find the following Figure 30.

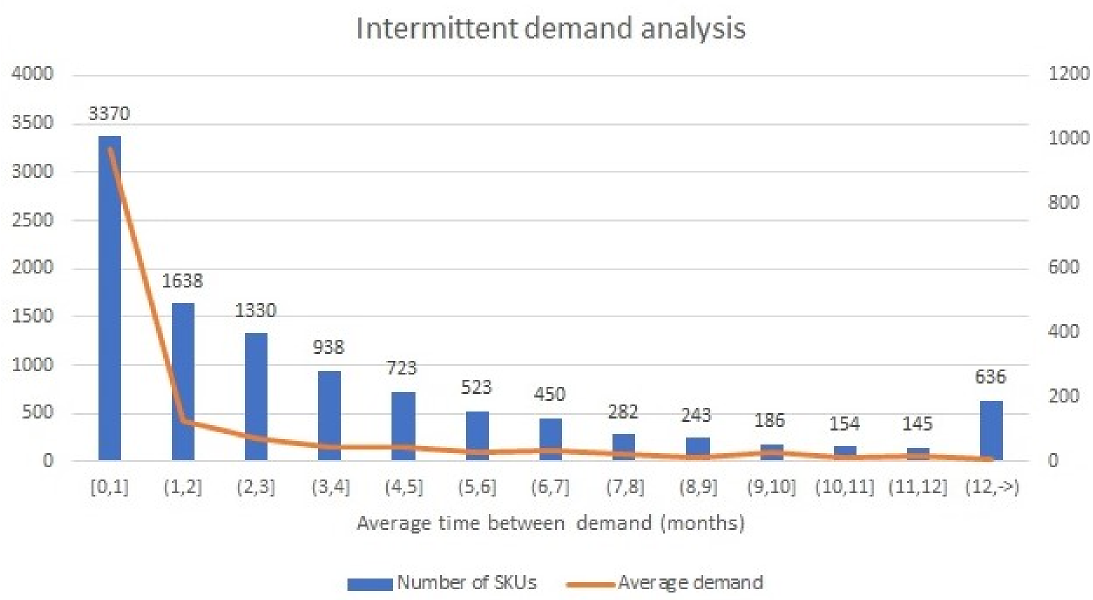

Figure 30: Intermittent demand analysis, years 2021, 2022 and 2023 (until September)

The histogram shows a large number of SKUs with, on average, long periods between two moments of demand. There are 4,280 SKUs (22% of all) with at least three months between two demand moments. However, the size of the occurring demand is minimal. On average, the monthly demand over all years is 0.80 for these SKUs. Counting the ones with an average demand size in the months when there was demand of at least 20 units, we find 215 SKUs. The policy regarding these parts are individually discussed with the production manager.

Appendix H Parts-suppliers delivery date performance

To determine the delivery date performance, we can compare the confirmed delivery date of the requested parts to the actual delivery date. The suppliers' performance and how well they respect the confirmed delivery date significantly influence the lead time's accuracy. In the case of the parts-suppliers, a date is requested by the buyers, which is confirmed or postponed by the supplier. By comparing the date confirmed by the supplier themselves, we do not penalise their lead time but the credibility of their promises. Negative values indicate that the order was delivered early. The order data from 2018 until October 2023 are included. Results are visualised in Figure 31.

There are many parts-suppliers. The 25 with at least 40 orders in the selected period have been analysed in detail. The top 10 suppliers contributed with 80% of the total ordered value. The weighted average performance, by their share of the total ordered amount, from the promised delivery date was a delay of 0.39 days. When ordering the differences between promised and actual delivery date, and taking the weighted average 3rd quartile, we find 3.39 days. This statistical measure is chosen to avoid penalizing early delivery (as, for example, the standard deviation calculation would).
The average value is excellent. The third quartile value is reasonable. High values with significant contributions can mainly be found for suppliers 218, 4814 and 617. It might be worth the effort to reduce this variance together with the supplier. The purchasing department has been and is working on improving the reliability of the suppliers. A recommendation for improvement regarding the late deliveries of some specific suppliers has been given to the purchasing department.

Considering the long supply lead times and flexible choice of supplier for many parts, the performance is not initially included in the model. An option to include extra lead time will be included in the model, but no detailed framework is required.

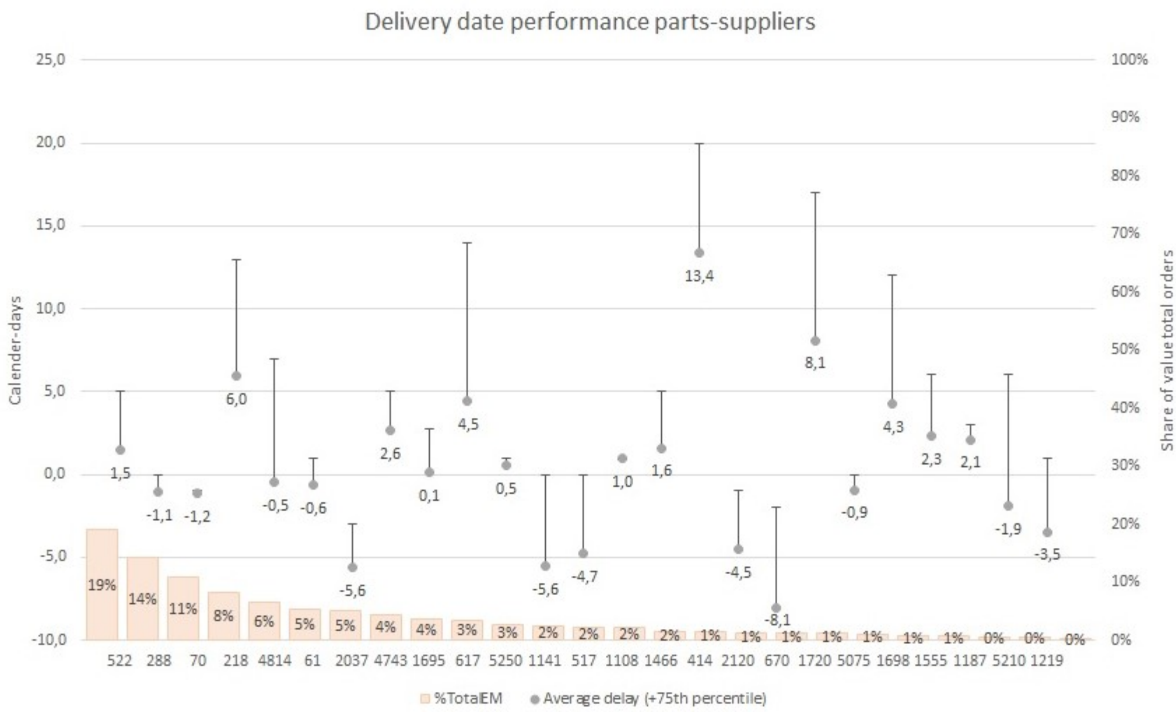

Figure 31: Delivery date performance parts-suppliers

Appendix I Data connections tool

The data connections required for the tool are visualised in Figure 32. The prototype of the inventory management tool is built around one main file, 'Master'. The general user input can be entered here, as well as the variable lower bound on the individual TFRs. The basic given information for each SKU, its value, supplier lead time, possible added extra safety lead time, MOQ and creation date are listed.

The second file, 'Forecasts', uses weekly historical demand starting from 2019 and the set exponential smoothing parameters α to calculate the expected lead time demand and its standard deviation. Together with the intermittent demand analysis, this enables the classification of the SKUs. The calculated parameters and division over the classes are given back to the main file.

In the file 'Target Fill Rate', the method by Teunter et al. (2017) is applied. Using the general SKU information and their parameters and the TFR input, the individual TFRs are set. These are loaded back into the main file.

Lastly, the file 'Classes' uses the provided ordering and holdings costs, as well as all SKU characteristics. Combined with the chosen policies and demand distributions for each class of SKUs, the parameters of the inventory policies are determined. These are added to the main file, which results in a complete overview of relevant information for each single SKU.

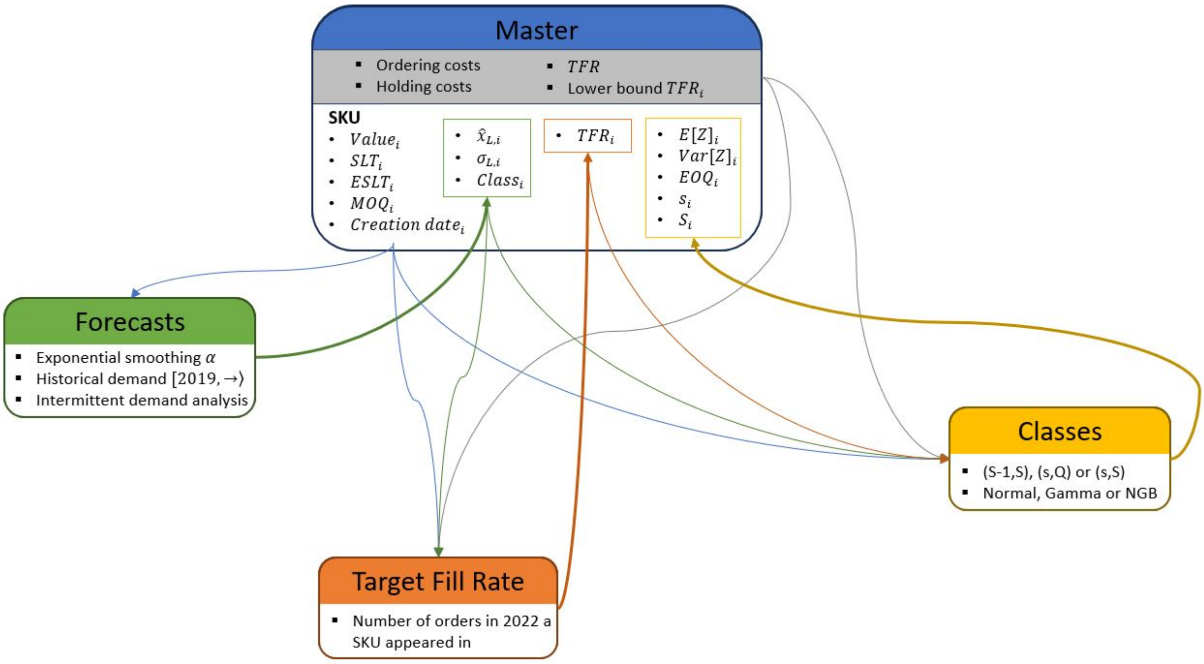

Figure 32: Data connections in tool

Appendix I.1 Loading historical data into 'Forecasts'

The historical data as given by the ERP system provides the files 'Consumi_202x_Data_x°Quad'. First, the movements have to be identified by week instead of date, by using the Excel function '=WEEKNUM(DATA)'. Combining the year and week in the following manner 'YEAR-WEEK' provides consistency over the years (e.g. 2022-12). This is stored as column 'YW'. Using the Power Query Editor, these three months of movements can be sorted by week and part ID. First, the rows of a specific part in a specific week should be grouped using the settings as in Figure 33.

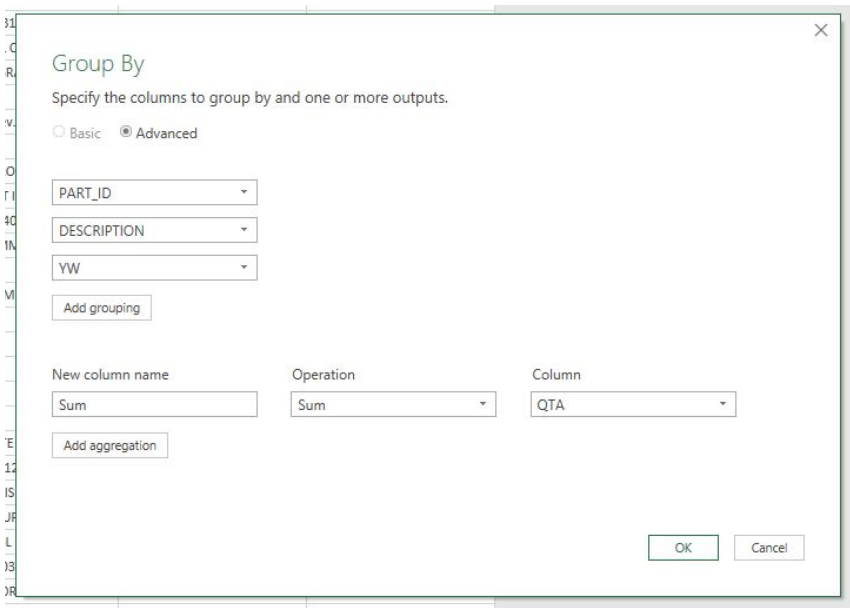

Figure 33: Grouping settings in Power Query Editor

To obtain the final table with the parts as rows and the weeks as columns, the settings as in Figure 34 are applied.

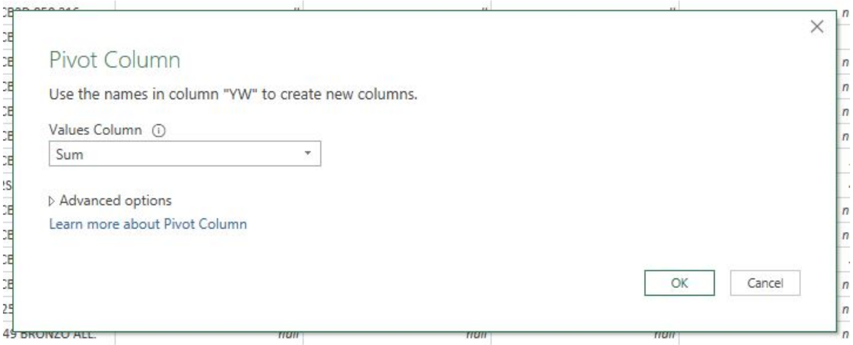

Figure 34: Pivotting settings in Power Query Editor

This final table can be easily added to the overall ‘Forecasts’ file using the Excel function ‘=XLOOKUP()’. After extending the cells in range to include the newly added data, mean and standard deviation parameters are up to date. These updated parameters are directly renewed in the ‘Master’ file.

Appendix I.2 Backorders analysis

The backorder analysis is based on the within the company well-known file ‘ANALISI PRELIEVI PER ODP’ (analyses pickings for end-products). Here, using a start and end date to determine the range of data to load, a direct connection with the ERP system is made and all pickings within this time range are selected. Essential for the backorder analysis is that the pickings are ordered from oldest picking date to newest (name of the column is *data_prelievo*).

The second page of the file contains a button connected to the first VBA code. After pressing this button, the code will run and all unique orders in the loaded list are analysed. The total number of parts picked for this order, the first picking data, the number of parts successfully picked on the first date, the number of parts later and the delay between the first and the last part are printed. With this, the order fill rate and the average delay of the orders can be determined and linked back to individual order codes.

The third page of the file concerns the individual parts which are backordered. Again, a button is present that activates a VBA code. This code prints all individual parts which are late, the quantity which was picked late and its product group. This enables the analysis of individual parts, to determine which are late most often.

Appendix J Simulation model

As described in Chapter 5, a simulation in Spyder - Python was created to determine the performance of the proposed inventory models. Three parts of this model are highlighted in this appendix.

The main part of the simulation determines the inventory movements for each part (Figure 35). On each day, the variables *Start Inventory*, *Incoming delivery* and *Pipeline* are read from the information already in the simulation. If there are older backorders waiting and there is an incoming delivery, the new pieces are used to fulfil this waiting demand. The actual quantity which can be added to inventory is reduced by this number. The occurring demand of the day is read from the demand information provided to the model. If there is enough in inventory, summing the starting inventory and incoming delivery, all demand can be filled. If not, the share which is available is reserved and reduced from inventory. The backordered demand is the rest of the initial demand.

After the mutations, the final positions are calculated. The on-hand inventory is determined by starting with the initial inventory, adding incoming deliveries and subtracting fulfilled demand. The inventory position uses this on-hand inventory but includes the replenishment orders in the pipeline and subtracts older waiting backorders and new unfulfilled demand. The initial on-hand inventory for next week is the same as the current final inventory position.

Based on the inventory position, a replenishment order might be placed. This is determined in a separate procedure (Figure 36). If the procedure indeed returns an order, the pipeline for the corresponding future days is increased. The arrival of the replenishment order is always on the Monday of the expected week of arrival to include the existing flexibility in picking in the model. If the part has been given an extra safety lead time, this is added to the duration of the delivery time.

```
#Daily simulation
for part in range(0,NrParts): #NrParts
for day in range(1,NrDays+1):

    #Initialize existing variables
    StartInventory=Simulation[part][1][day]
    IncomingDelivery=Simulation[part][2][day]
    Pipeline=Simulation[part][8][day]

    #Fill waiting backorders
    WaitingBackorders=Simulation[part][3][max(1,day-1)]+Simulation[part][6][max(1,day-1)]
    SumBackorders=WaitingBackorders
    if WaitingBackorders>0:
        if IncomingDelivery>0:
            WaitingBackorders=max(0,SumBackorders-IncomingDelivery)
            IncomingDelivery=max(0,IncomingDelivery-SumBackorders)
        Simulation[part][3][day]=WaitingBackorders

    #Fill ocurring demand
    OcurringDemand=DailyDemand[part][day]
    Simulation[part][4][day]=OcurringDemand

    #Fill demand fulfilled
    if OcurringDemand<=StartInventory+IncomingDelivery:
        FulfilledDemand=OcurringDemand
        Simulation[part][5][day]=FulfilledDemand
    else:
        FulfilledDemand=StartInventory+IncomingDelivery
        Simulation[part][5][day]=FulfilledDemand

    #Fill demand backordered
    BackorderedDemand=OcurringDemand-FulfilledDemand
    Simulation[part][6][day]=BackorderedDemand

    #Fill on hand inventory after demand
    AfterInventory=StartInventory+IncomingDelivery-FulfilledDemand
    Simulation[part][7][day]=AfterInventory

    #Fill inventory position
    InventoryPosition=-BackorderedDemand-WaitingBackorders+AfterInventory+Pipeline
    Simulation[part][9][day]=InventoryPosition

    #Fill beginning inventory next week
    Simulation[part][1][min(730,day+1)]=AfterInventory

    #Fill placed order with or without undershoot
    if USEUNDERSHOOT=='without':
        PlacedOrder=ReorderPolicyWithoutUndershoot(InventoryPosition, Class, part, ROwithoutU, OUTLwithoutU, EOQbtMOQ)
    else:
        PlacedOrder=ReorderPolicyWithUndershoot(InventoryPosition, Class, part, ROwithoutU, OUTLwithoutU, ROwithU, OUTLwithU,

    if PlacedOrder>0:
        Simulation[part][10][day]=PlacedOrder
        #Orders for the week should always arrive on monday to include actual flexibility in the week
        ExactArrivalDay=day+7*(SLT[part]+ESLT[part])
        MondayArrivalDay= ExactArrivalDay- ((ExactArrivalDay-3) % 7 )

        #Fill pipelines in right days
        for pipday in range(day+1,min(730,MondayArrivalDay)):
            Simulation[part][8][pipday]=Simulation[part][8][pipday]+PlacedOrder

        #Fill incoming delivery at right moment
        Simulation[part][2][min(730,MondayArrivalDay)]+=PlacedOrder
```

Figure 35: Main part of simulation

```python
def ReorderPolicyWithUndershoot(InventoryPosition, Class, part, ROwithoutU, OUTLwithoutU, ROwithU, OUTLwithU, EOQbtMOQ):
    match Class[part]:
        case 2:
            if InventoryPosition<ROwithoutU[part]:
                PlacedOrder=OUTLwithoutU[part]-InventoryPosition
            else:
               PlacedOrder=int(0)

        case 3:
            if InventoryPosition<ROwithoutU[part]:
                PlacedOrder=EOQbtMOQ[part]
            else:
               PlacedOrder=int(0)

        case 4:
            if InventoryPosition<ROwithoutU[part]:
                PlacedOrder=EOQbtMOQ[part]

            else:
               PlacedOrder=int(0)

        case 5:
            if InventoryPosition<ROwithU[part]:
                PlacedOrder=OUTLwithU[part]-InventoryPosition
            else:
               PlacedOrder=int(0)

        case 6:
            if InventoryPosition<ROwithU[part]:
                PlacedOrder=OUTLwithU[part]-InventoryPosition
            else:
               PlacedOrder=int(0)

        case _:
            PlacedOrder=int(0)

    return PlacedOrder
```

Figure 36: Order size determination

The size of the order is determined in a separate part of code. The current inventory position is taken from the main section, as well as other part characteristics such as class, reorder point, order-up-to-level and optimal order quantity. A match class method checks which policy should be used for the specific part under analysis and determines the order size based on the EOQ or difference between the order-up-to-level and the current situation.

In the main part, all movements of all the SKUs have been modelled. Afterwards, the demand is not provided individually per part, but as the workorders it originates from. For each line in the order (each part which is required for the end product), the code checks whether the demand for this part was actually filled. In case the demand for a part on a day comes from multiple orders, the order with the closest deadline receives the highest priority in using the available pieces. This priority has been included in the model by scanning the orders on a day from the earliest to the latest deadline. For each of the orderlines can then be determined whether the picking was successful. If not all pieces were present, the day on which these waiting backorders were filled is identified, after which the delay is calculated.

```
for orderline in range(0,NrOrderLines):
    #Known from excel
    OrdersCheck[orderline+1][0]=Orders[orderline]
    OrdersCheck[orderline+1][1]=ArticoloCode[orderline]
    OrdersCheck[orderline+1][2]=QtyRequired[orderline]
    OrdersCheck[orderline+1][3]=DayFirstPrelievo[orderline]
    OrdersCheck[orderline+1][7]=EndProducts[orderline]

    #Filled parts
    PartIndex=np.where((Part_IDs==ArticoloCode[orderline]))[0][0]
    DayIndex=list.index(Days,DayFirstPrelievo[orderline])+1
    FilledParts=min(AvailableInventory[PartIndex][1][DayIndex],QtyRequired[orderline])
    OrdersCheck[orderline+1][4]=FilledParts

    #Remove from availability
    AvailableInventory[PartIndex][1][DayIndex]-=FilledParts

    #When, if not yet, would it be complete?
    MissingParts=QtyRequired[orderline]-FilledParts
    if MissingParts>0:
        for day in range(min(NrDays,DayIndex+1),NrDays):
            MissingParts-=AvailableInventory[PartIndex][1][day]
            if MissingParts<=0:
                OrdersCheck[orderline+1][6]=day-DayIndex
                break

    #Fill rate?
    OrdersCheck[orderline+1][5]=FilledParts/QtyRequired[orderline]
```

Figure 37: Orderline check

Afterwards, all orderlines of each order are checked (Figure 37) to see how many parts of the full order were present. In case the order was not fully complete, the latest delay of the individual parts is taken as the delay for the order.

Appendix K Sensitivity analysis results

Appendix K.1 Without undershoot

		FR SKUs	TFR met	VFR	OHI (€)	#BO pieces	#Orders
Class 2	*113*	-	-	-	-	-	-
Class 3	*3,449*	-	-	-	-	-	-
Class 4	*366*	-	-	-	-	-	-
Class 5	*262*	-1.0%	-8.4%	-1.0%	-16,029	3,294	7
Class 6	*3,276*	-2.6%	-6.0%	-0.9%	-303,904	4,784	-10,071
Overall	*7,466*	-1.2%	-2.9%	-0.7%	-319,933	8,078	-10,064

Table 28: App. - Changes in parts results: without undershoot

FR	#BO orders	FR each order	#parts late	Delay (day)
-3.5%	348	-0.8%	0.1	0.0

Table 29: App. - Changes in order results: without undershoot

	Unit costs	#BOs	Costs
Orders	€2.92	348	€1,016
Parts	€4.08	715	€2,917
			€3,933

Table 30: App. - Changes in backorder costs results: without undershoot

Appendix K.2 (s,S)-policy for C-items

		FR SKUs	TFR met	VFR	OHI (€)	#BO pieces	#Orders
Class 2	*113*	-	-	-	-	-	-
Class 3	*3,449*	4.1%	7.3%	2.1%	16,072	-1,116	-3,413
Class 4	*366*	0.1%	-0,5%	0.5%	463	-1,004	-193
Class 5	*262*	-	-	-	-	-	-
Class 6	*3,276*	-	-	-	-	-	-
Overall	*7,466*	1.9%	3.3%	0.2%	16,535	-2,120	-3,606

Table 31: App. - Changes in parts results: (s,S)-policy for C-items

FR	#BO orders	FR each order	#parts late	Delay (day)
1.6%	-160	0.4%	-0.0	-0.8

Table 32: App. - Changes in order results: (s,S)-policy for C-items

	Unit costs	#BOs	Costs
Orders	€2.92	-160	€-467
Parts	€4.08	-343	€-1,399
			€-1,867

Table 33: App. - Changes in backorder costs results: (s,S)-policy for C-items

Appendix K.3 Adjusted lower bound Target Fill Rate

Decreased lower bound of 70%

		FR SKUs	TFR met	VFR	OHI (€)	#BO pieces	#Orders
Class 2	*113*	-0.4%	0.9%	-0.3%	-1,495	85	-
Class 3	*3,449*	-0.6%	0.0%	-0.3%	-3,809	150	-29
Class 4	*366*	-	-	-	-	-	-
Class 5	*262*	-	-	-	-231	-	-2
Class 6	*3,276*	-0.7%	0.4%	-0.1%	-9,281	356	-18
Overall	*7,466*	-0.6%	0.2%	-0.1%	-14,816	591	-49

Table 34: App. - Changes in parts results: lower bound of 70%

FR	#BO orders	FR each order	#parts late	Delay (day)
-0.8%	75	-0.2%	0.0	0.3

Table 35: App. - Changes in orders results: lower bound of 70%

	Unit costs	#BOs	Costs
Orders	€2.92	75	€219
Parts	€4.08	157	€641
			€544

Table 36: App. - Changes in backorder costs results: lower bound of 70%

Decreased lower bound of 75%

		FR SKUs	TFR met	VFR	OHI (€)	#BO pieces	#Orders
Class 2	*113*	-0.3%	0.9%	-0.2%	-947	62	-
Class 3	*3,449*	-0.4%	0.1%	-0.2%	-2,535	88	-16
Class 4	*366*	-	-	-	-21	-	1
Class 5	*262*	-	-	-	-197	-	-
Class 6	*3,276*	-0.4%	0.2%	-0.1%	-56,571	428	-5
Overall	*7,466*	-0.3%	0.2%	-0.0%	-60,271	578	-20

Table 37: App. - Changes in parts results: lower bound of 75%

FR	#BO orders	FR each order	#parts late	Delay (day)
-0.5%	48	-0.1%	0.0	0.3

Table 38: App. - Changes in orders results: lower bound of 75%

	Unit costs	#BOs	Costs
Orders	€2.92	48	€140
Parts	€4.08	99	€404
			€544

Table 39: App. - Changes in backorder costs results: lower bound of 75%

Increased lower bound of 85%

		FR SKUs	TFR met	VFR	OHI (€)	#BO pieces	#Orders
Class 2	*113*	0.4%	0.9%	0.1%	885	-33	-
Class 3	*3,449*	1.1%	0.9%	0.4%	6,690	-220	62
Class 4	*366*	-	-	-	-	-	-
Class 5	*262*	0.0%	0.4%	0.0%	102	-3	-2
Class 6	*3,276*	0.5%	-0.4%	0.1%	73,685	-464	14
Overall	*7,466*	0.7%	0.3%	0.1%	81,361	-720	74

Table 40: App. - Changes in parts results: lower bound of 85%

FR	#BO orders	FR each order	#parts late	Delay (day)
0.8%	-78	0.2%	-0.0	-0.4

Table 41: App. - Changes in order results: lower bound of 85%

	Unit costs	#BOs	Costs
Orders	€2.92	-78	€-228
Parts	€4.08	-134	€-547
			€-774

Table 42: App. - Changes in backorder costs results: lower bound of 85%

Increased lower bound of 90%

		FR SKUs	TFR met	VFR	OHI (€)	#BO pieces	#Orders
Class 2	*113*	0.9%	0.9%	0.5%	2,663	-129	-
Class 3	*3,449*	2.0%	2.3%	0.9%	15,064	-495	35
Class 4	*366*	-	-	-	37	-	-
Class 5	*262*	0.0%	-	0.0%	428	-3	-2
Class 6	*3,276*	1.1%	0.2%	0.3%	377,736	-1,400	39
Overall	*7,466*	1.4%	1.2%	0.2%	395,929	-2,027	72

Table 43: App. - Changes in parts results: lower bound of 90%

FR	#BO orders	FR each order	#parts late	Delay (day)
1.8%	-183	0.4%	-0.0	-0.8

Table 44: App. - Changes in order results: lower bound of 90%

	Unit costs	#BOs	Costs
Orders	€2.92	-183	€-534
Parts	€4.08	-358	€-1,461
			€-1,995

Table 45: App. - Changes in backorder costs results: lower bound of 90%

Appendix K.4 Added safety lead time

Added safety lead time of 1 week

		FR SKUs	TFR met	VFR	OHI (€)	#BO pieces	#Orders
Class 2	*113*	-1.2%	-1.8%	-1.6%	-947	445	-
Class 3	*3,449*	0.2%	0.7%	0.1%	3,976	-49	9
Class 4	*366*	0.1%	0.5%	0.1%	574	-134	-10
Class 5	*262*	0.0%	-0.8%	-0.0%	7,019	18	-40
Class 6	*3,276*	0.2%	0.4%	0.3%	451,728	-1,566	20
Overall	*7,466*	0.1%	0.5%	0.1%	462,351	-1,286	-21

Table 46: App. - Changes in parts results: added lead time of 1 week

FR	#BO orders	FR each order	#parts late	Delay (day)
-0.2%	16	0.0%	-0.0	2.2

Table 47: App. - Changes in order results: added lead time of 1 week

	Unit costs	#BOs	Costs
Orders	€2.92	16	€47
Parts	€4.08	-20	€-82
			€-35

Table 48: App. - Changes in backorder costs results: added lead time of 1 week

Added safety lead time of 2 weeks

		FR SKUs	TFR met	VFR	OHI (€)	#BO pieces	#Orders
Class 2	*113*	-1.2%	-1.8%	-1.6%	-947	413	-
Class 3	*3,449*	1.1%	3.1%	1.3%	9.730	-685	61
Class 4	*366*	0.3%	3.0%	0.2%	1,391	-490	4
Class 5	*262*	0.5%	5.3%	1.3%	23.221	-4,241	-18
Class 6	*3,276*	0.8%	2.1%	0.6%	634.524	-3,345	41
Overall	*7,466*	0.9%	2.7%	0.7%	667,921	-8,348	88

Table 49: App. - Changes in parts results: added lead time of 2 weeks

FR	#BO orders	FR each order	#parts late	Delay (day)
3.6%	-354	0.7%	-0.1	3.3

Table 50: App. - Changes in order results: added lead time of 2 weeks

	Unit costs	#BOs	Costs
Orders	€2.92	-354	€-1,034
Parts	€4.08	-666	€-2,717
			€-3,751

Table 51: App. - Changes in backorder costs results: added lead time of 2 weeks

Appendix K.5 Best solution

		FR SKUs	TFR met	VFR	OHI (€)	#BO pieces	#Orders
Class 2	*Best*	89.6%	70.8%	88.8%	31,169	2,909	280
113	*Diff. Actual*	-2.8%	-15.0%	-5.6%	-5,621	1,451	-
	Diff. Proposed	0.4%	0.9%	0.1%	885	-33	0
Class 3	*Best*	91.9%	78.9%	87.3%	122,126	6,776	3,480
3,449	*Diff. Actual*	4.3%	3.1%	-3.5%	-230,851	1,893	-
	Diff. Proposed	4.9%	7.8%	2.4%	23,213	-1,290	-3,409
Class 4	*Best*	98.2%	86.6%	95.8%	18,140	8,771	859
366	*Diff. Actual*	6.6%	55.5%	4.9%	-20,780	-10,248	-
	Diff. Proposed	0.1%	-0.5%	0.5%	474	-1,004	-193
Class 5	*Best*	93.9%	76.7%	95.8%	144,302	13,623	4,443
262	*Diff. Actual*	2.7%	45.4%	3.4%	-218,840	-11,095	-
	Diff. Proposed	0.0%	0.4%	0.0%	102	-3	-2
Class 6	*Best*	93.7%	81.7%	96.6%	1,771,682	18,567	14,831
3,276	*Diff. Actual*	9.9%	27.2%	7.2%	-539,698	-40,072	-
	Diff. Proposed	-1.6%	-5.9%	-0.7%	-232,221	3,938	-10,056
Overall	*Best*	93.0%	80.3%	95.6%	2,087,420	50,646	23,893
7,466	*Diff. Actual*	6.8%	17.4%	5.0%	-1,015,789	-58,071	-
	Diff. Proposed	1.6%	1.0%	-0.1%	-207,547	1,608	-13,660

Table 52: App. - Parts results: Best solution

	FR	**#BO orders**	**FR each order**	**#parts late**	**Delay (day)**
Best	79.0%	2,087	96.3%	1.6	18.6
Diff. Actual	17.6%	-1,753	5.4%	-0.9	5.1
Diff. Proposed	0.3%	-32	0.2%	0.0	-0.6

Table 53: App. - Orders results: Best solution

		# BOs			**Costs**		
	Unit costs	*Best*	*Diff. Actual*	*Diff. Proposed*	*Best*	*Diff. Actual*	*Diff. Proposed*
Orders	€2.92	2,087	-1,753	-32	€6,094	€-5,119	€-93
Parts	€4.08	3,327	-6,217	-66	€13,574	€-25,365	€-269
					€19,668	**€-30,484**	**€-363**

Table 54: App. - Backorder costs results: Best solution