Reducing Inventory Levels in the WKZ by Proposing a Revised Inventory Model

Author
J.H. Visscher
s2853248

Supervision

University of Twente
Dr. P.B. Rogetzer
Dr. ir. A.G. Leeftink

UMC Utrecht
Dr. ir. W.F.J.M. van den Oetelaar
O.I. Kooy

BSc Industrial Engineering and Management
Department of IEBIS

University of Twente
Drienerlolaan 5
7522NB Enschede
Tel. 053 489 9111

This thesis was written as graduation assignment as part of the Industrial Engineering and Management bachelor's program.

5 July 2024

Acknowledgements

Dear reader,

I am proud to present to you my bachelor thesis: "*Reducing Inventory Levels in the WKZ by Proposing a Revised Inventory Model*". This journey has been intense and challenging, but simultaneously extremely exciting, educational and satisfying. I would like to express my sincere gratitude to everyone who assisted me in writing this thesis.

To begin with, I am deeply grateful for the support of my academic supervisors: Patricia Rogetzer and Gréanne Leeftink. Their feedback was consistently sharp, detailed, and invaluable. I greatly appreciated the encouragement and support I received from them. They were highly approachable and always willing to help, something that I do not take for granted.

Additionally, undertaking this project would not have been possible without my supervisors at UMC Utrecht: Miranda van den Oetelaar and Odile Kooy. They provided me with the right tools and insights to conduct this research effectively. Their in-depth knowledge of the organization's processes was impressive, yet their excellent explanations ensured I never felt overwhelmed. I am also very thankful for the autonomy and trust they granted me throughout the research.

Finally, I would like to thank my friends and family for their unwavering support during the thesis-writing process. Whether it was adjusting the formatting of the front page or discussing the fundamental assumptions underlying my research, their feedback was always useful and much appreciated.

I hope that my research and findings prove to be practically applicable. I am proud to have laid the groundwork for a more quantitative approach to inventory management at UMC Utrecht.

Job Visscher

Enschede, July 2024

Management Summary

This thesis aims to reduce inventory levels at the Wilhelmina Kinderziekenhuis (WKZ) in Utrecht. The hospital, part of UMC Utrecht, is faced with renovation and construction plans, where the warehouses of several departments will be aggregated into one central warehouse. The available space for this warehouse is 26% lower than anticipated, meaning that a reduction in inventory levels is necessary. In addition, the frequency of shortage of some consumable products is too high and should be lowered. In this thesis, we focus on single-use consumable products. We solve the problems by providing the hospital with an inventory model to determine an appropriate order quantity Q and reorder point s for the consumable products that will be included in the new warehouse.

The WKZ currently uses an (R, s, Q) policy, meaning that employees review inventory levels every R periods, and place an order of size Q once the inventory level reaches s. We recommend the hospital to keep this policy for the new situation. There are two types of consumable products: those that are kept on stock and those that are ordered directly from the manufacturer. We focus on stock products in this thesis.

We apply a method of categorization to classify products based on their demand frequency (ADI) and variance (CV). This results in four categories: smooth, erratic, intermittent, and lumpy demand. We use a different optimization technique to compute the optimal ordering parameters for each category. Analysis shows that the majority (61%) of products is ordered infrequently ($ADI < 1.32$ weeks). Out of the 165 products that are ordered frequently, 152 have a high variance in demand ($CV > 0.5$). This indicates that there is a lot of variance in demand for the consumable products at the WKZ.

For determining the optimal Q, literature suggests the use of an economic order quantity EOQ for products with smooth and erratic demand. For products following intermittent and lumpy demand, no specialized technique could be found in literature. Therefore, we use average demand as parameter in determining the order quantity for these products. Applying EOQ to smooth and erratic demand gives some seemingly inconsistent results, where new order quantities deviate significantly from the current situation, both positively and negatively. Also using average demand for smooth and erratic demand gives more consistent results. We therefore propose this as an alternative method. Using the original method gives an increase in inventory levels of 8%, and an increase in number of orders of 12.3%. However, this results in a decrease in inventory costs of 28%. The alternative method results in a decrease of 9.4% in inventory levels. However, the number of orders increases by 18.7%. Combined, this results in a decrease in inventory costs of 9.8%. These results are not directly applicable yet, given that service level is not considered yet.

Literature suggests that for computing a reorder point s, modelling demand through a statistical distribution is appropriate. Smooth and erratic demand are modelled by a normal and Gamma distribution, respectively. Both intermittent and lumpy demand is modelled through a zero-inflated Poisson (ZIP) distribution, but lumpy demand gets an additional safety stock given its higher variance in demand. As a measure of service level, we use fill rate (P_2). We work with a target fill rate of 99%. Applying this model in combination with the original method of determining Q gives a reduction in inventory levels of 1% and a cost reduction of 23%. Using the alternative method of determining Q gives a reduction in inventory levels of 30% and a cost reduction of 35%.

The resulting reorder points for intermittent and lumpy demand show inconsistencies. This is likely due to the available data which only reflects orders, and not actual demand. In order to prevent inaccuracies in stock levels from occurring due to this assumption, we recommend the WKZ to not apply the model to these demand categories until better data becomes available. When leaving irregular demand out of the calculation, a reduction in inventory levels of 25% is possible, in addition to a minor decrease in number of orders (0.2%). Given that there is a space shortage of 26%, this thesis sufficiently solves the action problem.

Concluding, we recommend the WKZ to implement our suggested inventory model for smooth and erratic demand items. A reduction in inventory levels of 25% is possible in this safe scenario, where intermittent and lumpy demand items are left out. Therefore, the action problem is solved sufficiently well with this model. In addition, we delivered a quantitative way of uniformly controlling service level, which should reduce the number of shortages. Additionally, we recommend the hospital to gather data on actual demand instead of only tracking orders. We recommend employees to gather this data during the daily check of inventory levels. After doing this, the model proposed by this thesis should be run again, resulting in even more accurate results. Then, intermittent and lumpy demand items can also be included. This thesis' scientific contributions lie in developing a model in which each of the ADI-CV categories use a specialized technique for managing demand, and extending and validating the use of a ZIP model for managing irregular demand.

Contents

List of Variables and Abbreviations

Variable or abbreviation	Description
ADI	Average demand interval
BPMN	Business Process Model Notation
CV	Coefficient of variation
EOQ	Economic order quantity
ERP	Enterprise resource planning
ESC	Expected shortage per replenishment cycle
IC	Intensive care
MFA	Multi-functional area
MPSM	Managerial problem-solving method
NICU	Neonatal intensive care unit
OAV	Over-admitted variability
P_1	Cycle service level
P_2	Fill rate
PICU	Pediatric intensive care unit
Q	Order quantity
R	Review period
s	Reorder point
SPOZA	Spare part oriented ZIP approach
UMC	Universitair medisch centrun (academic hospital)
WKZ	Wilhelmina Kinderziekenhuis
ZIP	Zero-inflated Poisson

Chapter 1
Introduction

This chapter gives an introduction to the problem that this thesis tackles, and defines the way in which the problem will be tackled. The aim is to provide the reader with a motivation for the research, as well as the different steps that need to be taken to give an answer to the main research questions.

In Section 1.2, the action and core problem are identified. Next, in Section 1.3 the research questions that need to be answered in order to give a satisfactory solution to the core problem are given. For each question, we give a research design of how to answer that question. Section 1.4 discusses the outline and structure of the thesis. Most chapters serve to answer one of the research questions introduced in Section 1.3.

1.1 Background

This research is performed at the Wilhelmina Kinderziekenhuis (WKZ). The WKZ is a children's hospital located in Utrecht, The Netherlands. The hospital is part of a large academic hospital, UMC Utrecht. The mission of the hospital is to *"make and keep children healthy through innovative treatment"* (Wilhelmina Kinderziekenhuis, 2024). The hospital actively promotes conducting research and innovating in a wide variety of departments. Over the past years, the logistics department has grown significantly in size. Before, every department placed order independently. Many departments are now leaving these decisions up to the logistics department, which uses a centralized workflow of scanning barcodes to place orders.

Currently, the hospital is working towards a large renovation and expansion. Many departments, including the pediatric intensive care unit (PICU), neonatal intensive care unit (NICU) and obsetric department are gradually being moved to different locations in the hospital. With this change, many processes are being innovated. The aim is to provide the patient and their family with an increased amount of privacy and rest; so called "family-centered care". Whereas hospital rooms now contain several beds, the new rooms are completely individual. An impression of the new situation is shown in Figure 1.1. For this thesis, the changes in the logistics department are most relevant and will thus be focused on.

In 2009, the intensive care (IC) department of UMC Utrecht introduced a new logistical concept. Inventory is separated into four quadrants, based on frequency of usage and amount of required inventory per usage. Based on these quadrants, different ways of supplying are used. For example, the products with the highest amount and frequency of usage are kept on carts on every room, and replaced every 24 hours. Other products, which are used less frequently and in lower amounts, are only kept in a central storage location. The idea is to introduce a similar concept on the departments in the WKZ. However, given the larger variety in patient characteristics, especially weight, adjustments need to be made.

In the current situation, each department has a separate warehouse from which hospital beds and cupboards are stocked. The new situation will introduce a central warehouse, called the "multifunctional area" (MFA). One of the reasons for this is the high overlap in consumable assortments between departments. Furthermore, moving more facilitating processes and functions away from the department gives more focus and room for patient care. The fast-moving items that are kept on carts are stocked in the MFA, before being transported to the appropriate department and patient room.

(a) *New hospital wing that will house the PICU (EGM Architecten, 2024)*

(b) *New room type at PICU (4Building, 2024)*

Figure 1.1: *Artistic impression of new situation at WKZ.*

1.2 Problem identification

This section identifies the problems that the WKZ is facing regarding the renovation and construction of the new hospital wing. These problems are reduced to one core problem: outdated quantitative decision making in the inventory management process.

1.2.1 Action problem

The hospital is operating in an environment where personnel is becoming increasingly scarce, and costs are significantly rising. Thus, the urge to spend resources in a more efficient way is becoming increasingly pressing.

This becomes clear when considering the current expansion plans. The hospital board allocated significantly less space to the logistics department than logistics managers had hoped for. The current plans allocate 89 shelving units for the MFA, 26% less than the 120 shelving units that were asked for. This means that space has to be assigned more efficiently to make sure that all healthcare processes are not compromised. As the MFA also includes various healthcare-related processes such as preparing the beds, this induces a problem. If the hospital maintains the current stock levels and thereby the current space requirements, this reduces the quality of other processes. We formulate the action problem:

"Warehouses for consumables in the WKZ should go from decentralized to centralized, whilst adhering to space restrictions and without compromising service levels."

1.2.2 Problem cluster

In discussion with logistics employees and managers of the WKZ, possible underlying problems became visible. Employees have the experience that a lot of stock is kept of a variety of consumables. Many products are ordered infrequently - around once every month - and in high lot sizes. This leads to higher space usage, but also to situations where products are wasted. Given the limited insight into product waste, this is kept out of scope. In addition to the high stock of many products, having a shortage of products is a common challenge in the hospital. In conversation with employees, we observed that having a shortage of products occurs relatively irregularly. However, the impact of a shortage is very high, given that many products in the hospital could be life-saving. When finding solutions to reduce space usage, we therefore need to keep in mind whether no additional risk of incurring a shortage is caused. As such, we include the problem of having a high shortage rate in the scope of this research. The corresponding problem cluster can be found in Figure 1.2. An explanation of further problems is given in Section 1.2.3.

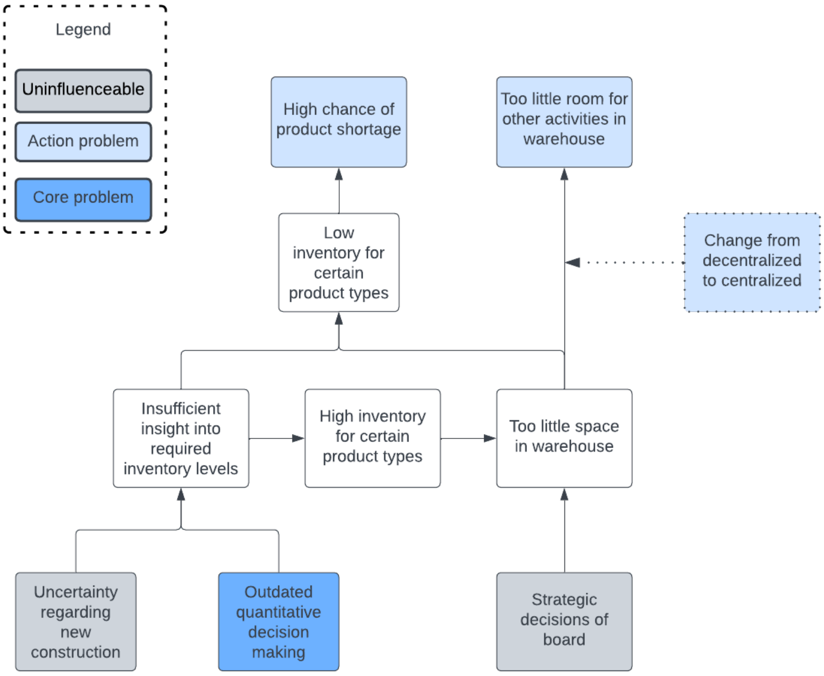

Figure 1.2: *Problem cluster for reducing space usage in the WKZ.*

1.2.3 Core problem

As can be seen in the problem cluster, there are further causes to the problem that inventory levels are suboptimal. In general, there is insufficient insight into the necessary inventory levels. This is partly since there is a lot of uncertainty with the construction of the new warehouse. The assortments are still up to change, and the introduction of the new logistical concepts means that the location of some products is uncertain. Healthcare processes are also not finalized, increasing this uncertainty. These decisions have to be made through collaboration of logistics and healthcare and are kept out of scope of this research, given time restrictions.

However, there is another underlying problem that can be influenced. The current inventory policies are outdated: the calculation is largely based on averages and the data that is used is not up to date. The first aspect means that variability in demand is not considered in a systematic way, leading to inconsistent rates with which product shortages occur. Outdated data means that new demand patterns are not considered by the current model. Healthcare is a sector in which new developments and trends are constantly occurring, leading to high variability in demand patterns. The hospital would therefore benefit from a model which can easily be updated with new data. Given that this problem leads to the other, previously identified problems and has no other currently identified root-causes, we can consider it a core problem.

Summarizing, we formulate the core problem as follows:

"Inventory policies at the WKZ are outdated and should be more quantitatively motivated and calculated in the new warehousing situation."

We elaborate on the problem owners for this problem. The main problem owner is the logistics department: they are responsible for the logistic processes in the MFA. In addition, the construction department acts a part in defining processes and setting priorities. Finally, the healthcare providers play a major role in defining priorities and assessing the quality of inventory policies. Logistics facilitates the healthcare department. Therefore, for every decision the impact on quality of care should be considered. An example is space usage. Every additional m^2 in the warehouse means one fewer m^2 that can be spent on providing care.

1.3 Problem solving approach

This section discusses the approach of researching the possibilities for the WKZ to improve their inventory policies. First, the main research question and sub-questions are given, which give structure to the research. For each question, the steps that need to be taken in order to give a well-motivated answer are explained. Next, limitations of the research are discussed. This includes a discussion on data availability, but also validity and reliability of results.

1.3.1 Research design

The main research question is tasked with finding an appropriate solution to the core problem as identified in Section 1.2. We formulate this question as follows:

How can the WKZ best revise its inventory management policy in order to reduce inventory levels and without compromising service level for the MFA?

Here, the assumption is made that decreasing inventory levels also leads to lower space usage. The question is answered by walking through the different steps of the managerial problem solving method (MPSM), defined by Heerkens and van Winden (2017), given its good fit to applied research. The first two steps - problem identification and solution planning - are tackled in this chapter. The next steps are each tackled with a separate sub-question in their own chapter. We identify the sub-questions with the necessary steps below.

Problem analysis

This step of the MPSM is tasked with analyzing the problem in a more detailed way. We will explore what the causes and effects of the problem are, in order to obtain a better idea of a possible solution. For this, the current process needs to be identified. We formulate the first sub-question:

Question 1. *How is consumable inventory managed in the WKZ?*

This question is answered by looking at the current and new supply process, and by looking at the different products and identifying how they can be categorized. In addition, the quality of the necessary data in the WKZ should be checked. This leads to three sub-questions:

a) *How does the order and supply process in the current situation work?*

This question will be answered by walking along and talking with logistics employees that are tasked with performing the ordering process. The most important focus point is the process of ordering materials to the decentralized warehouses.

b) *How will the order and supply process in the MFA work?*

This question serves to understand the extent to which processes will change in the MFA, and which of these changes is relevant for this research. Even though the entire process has not been finalized yet, some decisions have already been made and need to be adhered to in the research. For example, the review period of stock has already been determined to be one day.

The next sub-question is as follows:

c) *How can consumable products in the WKZ best be categorized?*

This question requires studying literature for best practices on categorizing inventory. For example, one could look at how much the product is ordered, or the variability of demand. The different types of inventory will likely also have different techniques of computing optimal policies. The final sub-question is as follows:

d) *How is demand data gathered in the WKZ, and what is the quality of this data?*

It is important to get insight into the data gathering process, to see where possible data cleaning needs to be done.

Solution generation

In this step, the possible techniques for inventory management are considered through a literature review. Given that the WKZ currently uses an (R, s, Q) policy, which is not likely to change in the near future, we restrict ourselves to this policy. This choice is motivated more in Section 1.3.3. We refer to Section 3.1 for some more elaboration on what an (R, s, Q) inventory policy entails. We introduce the following research question:

Question 2. *What techniques exist for computing the optimal parameters in an* (R, s, Q) *policy?*

We performed some preliminary analysis, which made it clear that some products can be considered as having 'regular' demand, and some 'irregular' demand. They will be separated in the research, given the different frameworks for computing optima:

a) *What techniques exist for managing inventory subject to regular demand?*

b) *What techniques exist for managing inventory subject to irregular demand?*

Solution choice

Here, the different techniques from the previous step are compared on feasibility and effectiveness. The corresponding research question is as follows:

Question 3. *Which techniques should be used for managing inventory in the new warehouse in the WKZ?*

The choice will be based on what literature says, and discussions with the company and university supervisors. Given the time span of 10 weeks, the most complex techniques are likely not feasible to implement. In addition, a complicated technique is difficult to explain to the stakeholders, which might put them off from using it. The balance between this and the possible effect of the model should be found.

In addition to choosing the most appropriate model, the performance of this model should be assessed. This introduces another research question:

Question 4. *How does the chosen model perform?*

By calculating the necessary ordering quantities and frequencies, an estimate of the costs and space usage can be made. For this, data of all the departments that are included in the MFA is needed. The overlap of assortments needs to be taken into account with this data to give recommendations based on the aggregated demand. The situation with the implemented recommendations will be compared to the current policy, so the savings can be calculated.

One can observe that the final steps of the MPSM, evaluation and implementation, are not included in the research design. This is because it is highly unlikely that the new policy is implemented within the 10 weeks of conducting research. This makes it unfeasible to give a good evaluation on how the model performed after implementation. However, some recommendations for implementation will be given.

1.3.2 Limitations

This section discusses limitations that might occur during the research. This includes data availability as well as validity and reliability issues. Chapter 6 also reflects on these issues.

Data availability

Data on historical orders is available through a functionality of the ERP system of UMC Utrecht. Given that there was a large change in placing orders on some departments in 2022, only data of 2023 will be considered. The data is of high quality and needs little cleanup. However, only the data on orders is available (date and quantity of order), and not the actual demand. Therefore, some assumptions need to be made in order to 'rewire' the data so it reflects demand patterns. This will be done through the support of literature.

Validity and reliability

Validity refers to the extent to which one measures the thing that needs to be measured, whereas reliability is the extent to which one measures the same thing consistently. Both points are discussed here.

A possible validity error occurs with the quality of data. As described above, the data is not fully what one would want in inventory management research. The variability in demand does not become fully apparent through data on historical orders. This could lead to situations where we underestimate the required stock. This will be solved by introducing an additional safety factor in the research. Once the model is verified with new data from current year, the necessary safety factor can then be adjusted.

In addition, a validity error could occur with the fact that healthcare processes are also changed in the renovation. If this change also happens in what products are required for certain treatments, this changes the required stock levels. Once these changes become clear, they should be implemented in the model. This can be done by multiplying demand data by a certain factor, for example. Evaluating and re-running the model with data gathered after the implementation will also reduce these errors.

A reliability error is that we can only use data from 2023. Whilst one full year of data seems like a lot, general trends or seasonal effects are difficult to capture with this. As such, they should be kept out of scope. Statements about the true distribution of demand can also be given with less confidence given the limited time frame. Furthermore, the lack of insight into demand trends could mean that next years show different patterns. Again, this should be researched by verifying the model on the data of 2024 to see how it would have performed. This way, the hospital does not run the risk of keeping incorrect stock levels due to the model.

1.3.3 Scope

Several references to the scope of this research have already been made in the preceding sections. Here, they are expanded and elaborated upon. For theoretical background, we refer to Section 3.1.

An important consideration with regards to scope is which products will be considered. In principle, any consumable that is currently used in one of the departments that will be supplied from the MFA falls within scope. There are a few exceptions:

- Any product that belongs to one of the following categories: linen, medication, medical instruments, and medical gasses.
- Products that are introduced in the assortment after the start of 2023.
- Products that are not currently managed by the logistics department.
- Products that are categorized as 'ZNLA'. For more on this, one can refer to Section 2.1.

Although many different policies exist in inventory management, the current (R, s, Q) policy will be the main focus. The hospital recently changed to this policy and is now implementing it on most departments. This means that changing to a new policy would cost a lot of time and effort. The systems need to be developed that can handle a new policy, and all employees need to be educated to work with it. Therefore, we expect that the implementational effects of changing to an entirely new policy outweigh the possible gains. The largest short-term improvement is expected to be made by recomputing the optimal s and Q and should therefore be focused on.

Many parameters of this model are uncertain and can be tweaked and experimented with during the research. However, review period R has already been defined to be 1 day.

Finally, possible trends and patterns in demand are left out of scope. There is and insufficient amount of data to be able to make confident statements about such trends. Products that display high seasonality effects will therefore likely have sub-optimal inventory levels. However, given their high variance in demand and therefore safety inventories, there should not be a larger risk of running out of stock for these products.

1.4 Outline

The next chapters are structured according to the sub-questions that we identified in Section 1.3. Each chapter answers one such question. We outline the resulting structure below.

Chapter 2 discusses the situation at the WKZ. We begin by elaborating on the current situation of ordering consumable product to the decentralized warehouses, and how this process will be extended to the MFA. Next, we discuss three ways of categorizing products. Products with different categories have different suitable techniques for managing their inventory levels. Finally, we discuss some implications regarding data availability and quality.

Chapter 3 discusses the theoretical framework of inventory management by elaborating on some key concepts. In addition, we elaborate on our findings from literature regarding inventory management techniques for regular and irregular demand. We aim to apply a different technique from literature to each category.

Chapter 4 discusses how we integrate the findings from Chapter 3 are into one model. For this, we develop some extensions to the existing models in literature.

Chapter 5 discusses the results that are generated by applying the model from Chapter 4 to the data on orders from 2023. We slightly adapt the model by proposing an alternative method of computing the order quantity, given some initial results.

Chapter 6 provides concrete advice to the WKZ based on the results from Chapter 5. In addition, we elaborate on several objects of further research that would further improve the solution to the core problem of this thesis. Finally, we discuss the scientific contributions of this thesis.

We conclude this thesis by summarizing all findings and answers to the research questions in Chapter 7.

Chapter 2
Situation Analysis

This chapter discusses the current situation at the WKZ in more detail. We elaborate on the current process of ordering consumables to the departments in the current situation, and in the new situation where the MFA is used (Section 2.1). In addition, we explore the different consumable products and how they can best be categorized (Section 2.3). We focus on how categorization is currently done, and how it can be done to best fit this thesis. Finally, we discuss the available data on products and demand (Section 2.4). The availability and quality of the data is discussed, in addition to the way in which data cleaning is performed.

2.1 Current ordering process

This section discusses the way in which stock is currently managed at the WKZ. First, it describes the process in general, after which it focuses on the most important sub-process for this thesis. This is the process of ordering from the central warehouse to the decentral warehouses.

2.1.1 General overview

There are four general locations that play a role in the supply process of the WKZ: the manufacturer, the central warehouse, the decentral warehouses and hospital bed itself.

The biggest location where inventory is stored is the central warehouse located in the outskirts of Utrecht, called the "Hudsondreef". Here, products that are ordered from manufacturers arrive and get stored. A separate team of the logistics department operates in this location. The process of ordering to the Hudsondreef falls outside the scope of this research.

In the current situation, every department has separate decentralized warehouses. For example, the PICU has two wings, each with its own storage facility. A logistics employee is responsible for keeping track of inventory levels in these warehouses and places an order once the level is too low. We elaborate on this process in Section 2.1.2. The logistics employee also distributes the products to appropriate cabinets, beds, etc. For example, on the PICU, every bed has a tray of products that are used more frequently. The employee replaces this tray twice per day.

2.1.2 Ordering to decentralized warehouses

In this section, we elaborate on the process of stocking the decentralized warehouses. A system with digital scanners is used to make the process more efficient and less prone to errors compared to a system where everything is done manually. Every morning, on weekdays, a logistics employee starts with checking the inventory levels of each product in a certain warehouse. One to two employees assigned are to a department per day. Once the employee enters a warehouse, a QR code needs to be scanned to link the system to the right location. Every product has a reorder point (s) and a fixed order quantity (Q), as identified in Section 3.1. If the employee observes that the inventory level has fallen below the reorder point, they place an order with the size of the fixed order quantity. This is done by scanning the QR code on a card on the shelf of each product in the decentral warehouse. This card also mentions the reorder point, so the employee can quickly observe whether an order needs to be placed. The fixed order quantity can also be found on the card, but the system automatically places an order of that size once the card is scanned.

There are two types of cards: white (also referred to as 'ZROH') and yellow (also referred to as 'ZNLA'). ZROH products have high usage rates and are kept on stock in the Hudsondreef and can therefore be delivered quickly to the decentralized warehouses. If an order for a ZROH product is placed in the morning around 10 am, the order will be delivered in the afternoon on the same day. ZNLA products are used less frequently and therefore do not have a permanent location in the Hudsondreef. They have to be ordered directly from the manufacturer, leading to higher lead times of one to four weeks. To make sure that a ZNLA product that is placed on order is not ordered twice, the yellow card can be turned around to become a red card. This indicates that the product has been ordered at the manufacturer. This does not always go right: employees occasionally forget to turn the card back once the order arrives. This means that the product will be ordered again, leading to a surplus in inventory. Once a ZNLA product is delivered to the Hudsondreef, it is sent together with the first upcoming shipment that will travel to the WKZ. A shipment from the Hudsondreef to the WKZ first arrives in a central transport hub in the UMC. Then, it is transported by a transportation team to the right location in the WKZ. This transportation team also includes the people responsible for shipping from the Hudsondreef to the UMC. The order arrives with a packing slip so the employee can verify whether everything has been delivered in the right quantity. Finally, the employee that placed the order in the morning stocks all the products in the right location in the warehouse. A graphical representation of the process can be seen in Figure 2.1. The diagram is made using *Bizagi Modeler* (Bizagi, 2024), according to Business Process Model and Notation 2.0 (BMPN 2.0) conventions.

Remark. ZNLA products are left out of scope for this research, given the low average weekly demand of 0.7 product batches compared to 8.5 for ZROH products. As such, the impact of determining a policy for these products is lower. In addition, data on lead times is unavailable and the product prices are infeasible to retrieve in the way they are currently stored. Furthermore, the WKZ is bound to supplier agreements when considering the order quantity, meaning they are difficult to change.

It is important to note that on weekend days, no products are ordered from the Hudsondreef to the decentralized warehouses. Therefore, on Fridays, the logistics employees have to make sure that the warehouses contain enough inventory to last until Monday. We aggregate data on orders on weeks in order to prevent misrepresentations of variance in demand due to this effect. Section 2.4 discusses these considerations in more detail.

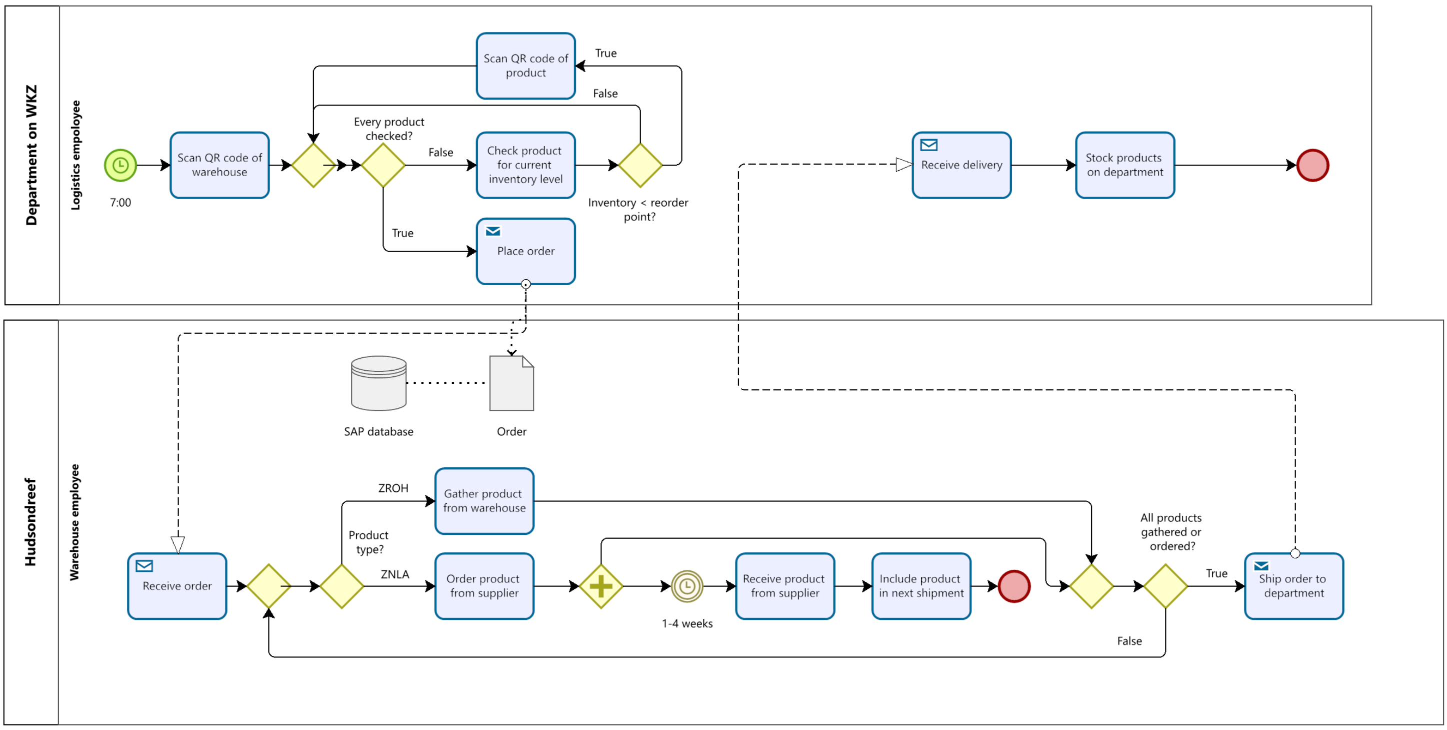

Figure 2.1: *BPM representing the current process of ordering to the decentralized warehouses.*

The policy that employees apply to decide whether or not to place an order can be considered an (R, s, Q) policy. We elaborate on this policy in Section 3.1. In the current situation, for most of the decentralized warehouses that will be transitioned to the MFA, the review period R is 1 day. In some warehouses that will also move to the MFA, R is equal to about 3 days, where the inventory levels are reviewed twice per week.

2.2 New situation

As mentioned in Section 1.1, the situation after the renovation will move the decentralized warehouses to a central warehouse, called the multi-functional area (MFA). This means that the demand of all separate departments is aggregated. Many consumables are used on several departments. There are 1,800 instances where a unique product is used for the departments that are being moved to the MFA. This includes products that are included on multiple departments which are therefore also counted multiple times. There are 900 different consumables in total, indicating that the overlap in products is significant. 550 of these products are kept on stock in the central warehouse (ZROH). By aggregating the demand of multiple warehouses, the hospital benefits from a relative reduction in variance of demand. This means that safety stocks can likely be reduced in the new situation.

The ordering process in the MFA remains similar to the current situation for ordering products to a decentralized warehouse. The same (R, s, Q) policy will be used, and the products are still scanned using the same procedures. The review period has been determined to be 1 day by the logistics team of the UMC. This is the same for the current situation in most departments, but some departments will thus move to shorter review periods. The process of delivering the consumables to all departments will change more, but this falls outside the scope of this research. As such, one can refer analogously to Figure 2.1 for an indication of the process of ordering consumables to the MFA. As identified in Section 1.2 and 2.1, there are things that need to be improved before moving to the new situation where the hospital has to deal with reduced space availability. The lack of a sophisticated mathematical method to compute reorder point s and order quantity Q gives an opportunity to deal with this. A new mathematical method will increase the benefits from aggregating the overlapping products that are now stored on multiple departments and gives more certainty for product availability.

2.3 Product categorization

As stated in Section 2.2, the departments that will be supplied from the MFA have in total an assortment of approximately 550 different ZROH products. Given that these products show different characteristics, it is not possible to apply the same technique for calculating Q and s everywhere. For example, some products are ordered frequently (once per week), whilst others are orders infrequently (once per month). Therefore, we aim to find a way to categorize the products that share similar characteristics effectively. We can then apply the same technique to products that fall in a certain category. In addition, categorization gives deeper insight into the types of products that we are dealing with, and possibly allow for prioritization on the most important products. First, we pay attention to the method that WKZ and UMC use to categorize and prioritize products (Section 2.3.1). Next, we elaborate on ADI-CV, which is used during the remainder of this thesis to distinguish products on their demand patterns (Section 2.3.2). Finally, we elaborate on ABC classification, which can be used to categorize products based on their average value (Section 2.3.3).

2.3.1 UMC Matrix

UMC Utrecht has developed their own framework of categorizing products. This method is currently used on the IC department to identify the relative importance of products, and which logistical implications this has. The products are categorized based on the demand that occurs when a treatment takes place. The analysis that we perform can play a role in classifying products according to the framework of the UMC. For example, products with high demand can be identified and put in a certain category. In addition, the categorization can be used later to decide which products to focus on in our analysis.

Products are distinguished according to two criteria: how many patients need the product (i.e. frequency of usage) and how much of the product one needs when demand occurs (i.e. size of demand). This introduces four categories, developed by UMC Utrecht:

- **24-hour stock.** These products have the highest frequency of usage, and once usage occurs, the demand is high. Because of this, these products are kept as closely as possible to the hospital beds. The products are stored in special carts in patient rooms and replaced every 24 hours from the decentralized warehouse where a replacement is also stored. As an example, one can think of gauzes or syringes as products in this category.
- **Procedure trays.** Here, we are dealing with products that are used frequently, but when they are needed one does not need a lot. As such, they are gathered in procedure trays. The aim is to collect the necessary items for a certain treatment on such a tray. An example is the C-PAP set, which is used when a recently born baby needs to be transported. This is used often, but the set itself is small and does require a lot of products.
- **Theme carts.** Some products have a low frequency of usage, but when they are used one needs a high quantity. Products that belong together once such a relatively rare event occurs are gathered in theme carts. An example of a theme cart is the burns cart, which is needed for burn victims.
- **Specials.** These products are used rarely and used in low quantities. Currently, they are stored in the decentralized warehouses, and therefore require a longer walking distance than the other products before they can be used. Because of the relatively low urgency of these products, this does not have a high impact on the healthcare processes. An example is printing paper, which does not need to be stored in the decentralized warehouses given its low usage.

One can view the categories in a graphical representation in Figure 2.2. The words in italics indicate the location where these products will be stored in the new situation. Currently, the decision on the category in which a product belongs is made in collaboration with logistics and care providers. Logistics proposes a decision based on the data that they observe, whereas the healthcare providers have experience with actually using the products.

2.3.2 ADI-CV

A popular method of categorizing products according to their demand patterns is the so-called ADI-CV classification method. For example, Finco et al. (2022) use the framework as basis for the analysis of irregular demand items.

Two values need to be computed: the average demand interval (ADI) and the coefficient of variation (CV). ADI refers to the average number of periods in which we observe no demand between two periods of positive demand (Finco et al., 2022). As such, it gives an indication for how frequently we observe periods with demand. It can be computed as follows:

$$ADI = \frac{1}{N}\sum_{i=1}^{N} t_i \tag{2.1}$$

where N is the number of periods with positive demand and t_i the interval between positive demand period $i-1$ and i.

The coefficient of variation refers to how variable the periods with positive demand are. To normalize the value to not be subject to scaling, we divide by the mean-positive demand.

$$CV = \frac{\sigma}{\mu} \tag{2.2}$$

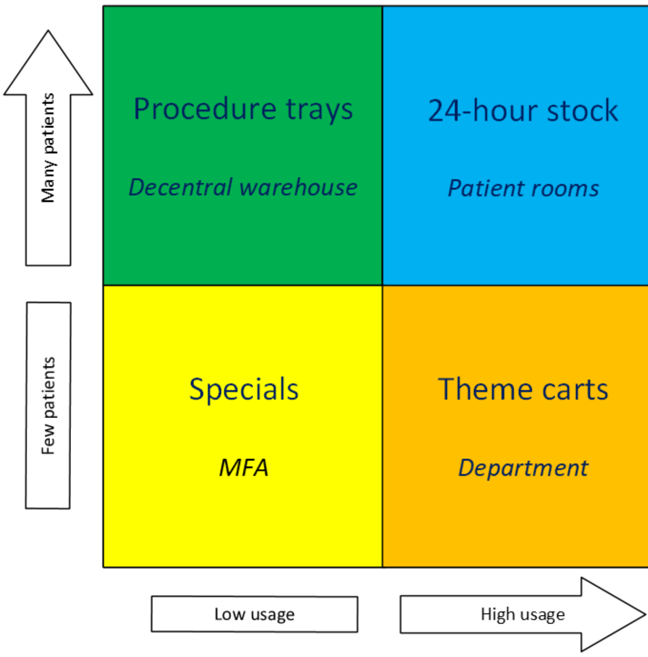

Figure 2.2: *Matrix that UMC Utrecht currently uses to categorize products.*

where σ is the standard deviation and μ the mean of the demand per positive period. Note that in this calculation, the values where we observe no demand are not considered. When dealing with a dataset where the true distribution is unavailable, we have the maximum likelihood estimators:

$$\hat{\mu} = \frac{1}{N}\sum_{i=1}^{N} X_i \tag{2.3}$$

$$\hat{\sigma} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(X_i - \hat{\mu})^2} \tag{2.4}$$

where X_i denotes the i-th period of positive demand. Once these values are computed, we categorize the demand as follows (Finco et al., 2022):

- $CV < 0.49, ADI < 1.32 \rightarrow$ Smooth.
- $CV < 0.49, ADI > 1.32 \rightarrow$ Intermittent.
- $CV > 0.49, ADI < 1.32 \rightarrow$ Erratic.
- $CV > 0.49, ADI > 1.32 \rightarrow$ Lumpy.

In this thesis, we call smooth and intermittent demand 'regular' and erratic and lumpy demand 'irregular'. A graphical representation of this categorization can be found in Figure 2.3. A visualization of what the four demand types look when plotted over time is shown in Figure 2.4. If the demand is smooth, most common inventory management techniques can be applied. An example of this is assuming that demand follows a normal distribution. In other cases, where the demand is intermittent, erratic or lumpy, more specialized techniques need to be applied. Sections 3.2 and 3.3 go into the techniques for regular and irregular demand in more detail. Table 5.1 shows how the demand types are distributed in the assortment of the WKZ. Note that we compute these values with a week as time unit, since some departments do not review inventories daily and no department places orders in the weekend. Given that the lead and review times are both 1 day in the new situation, we should compute the numbers back to daily demand. This is elaborated upon in the model development section.

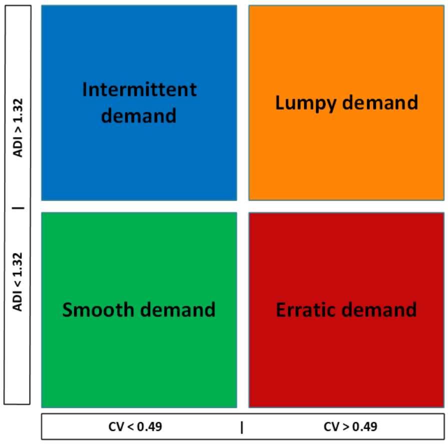

Figure 2.3: *Graphical representation of ADI-CV categorization.*

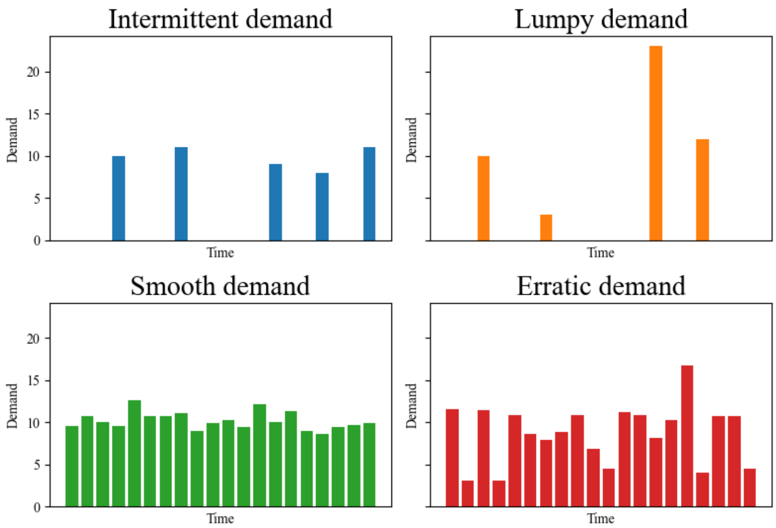

Figure 2.4: *Visualization of the four types of demand according to ADI-CV categorization (fictional data).*

2.3.3 ABC analysis

The final categorization method that we elaborate on is ABC analysis. This system ranks products based on their usage value: the usage rate times the individual product value (Slack and Brandon-Jones, 2019). ABC analysis then reasons through the Pareto principle that 80% of the value comes from roughly 20% of the products. As such, the products are arranged based on their inventory value. Shenoy and Rosas (2017) identify three classes:

- **Class A** includes the first 20% of products that account for 80% of the total inventory value.
- **Class B** refers to the next 30% of products that account for roughly 15% of the inventory value.
- **Class C** is made up of the final 50% of products, again taking op approximately 5% of the total inventory value.

Ranking products according to these classes gives an indication of how important a product is. In the WKZ, this is a bit more difficult compared to commercial companies, given that any product could be life-saving. However, the ABC analysis still gives insight into where costs can potentially be reduced the most. In Hasachoo et al. (2019), the ABC method is extended to an ABC-VEN approach, where products are also ranked according to criticality. Here, 'V' refers to a product being vital, 'E' stands for essential, and 'N' stands for nonessential.

In the assortment of the WKZ, the first 110 out of 524 most valuable products make up 80% of the total inventory value, the next 150 products 15%, and the final 264 the last 5%. This is 21, 29 and 50 percent of the total number of products, respectively. As such, the inventory of the WKZ adheres to the Pareto principle. One can see the products plotted against their cumulative value in Figure 2.5.

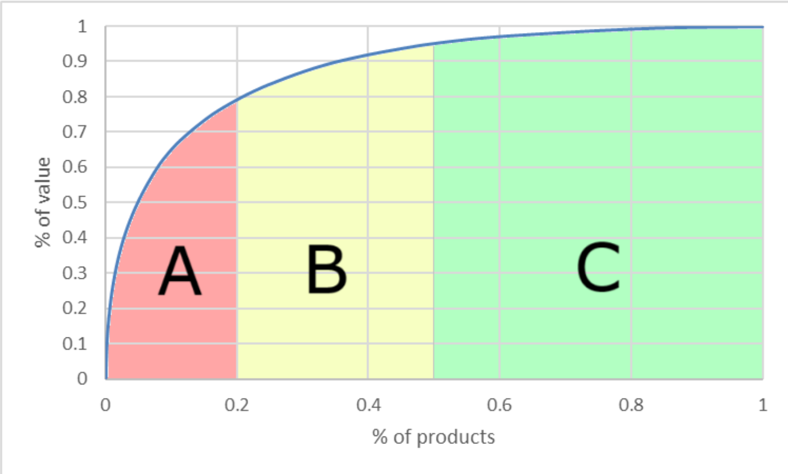

Figure 2.5: *Cumulative percentage of value (y-axis) of x% most valuable products (x-axis) in the WKZ.*

2.4 Data availability and cleaning

This section discusses the available data in more detail. Given that the WKZ will enter an entirely new situation where stock is located in a central warehouse, there is no direct demand data available. First, we elaborate on the way in which we work around this, so the view on what demand will look like in the new situation is as accurate as possible. In addition, this section mentions the cleaning techniques that we use in order to correct errors in the data.

2.4.1 Data availability

Once a logistics employee scans the products and an order is sent to the Hudsondreef, the order is stored in an ERP system that WKZ uses for inventory management. The quantity of the order is stored, in addition to the location of the warehouse that the order is for. This data can be accessed through an ERP tool, called "SAP Analysis". Through this tool, all departments and their warehouses can be selected, and their orders are automatically added up. Given that many products are ordered once per week or less, we aggregate the data per week. The result is a table with all consumables that will, in the new situation, be supplied from the MFA and their weekly orders. An example of three such consumables with their weekly order quantities can be seen in Table 2.6.

We believe that given the available data, this is the best way to identify the (variability in) demand. Since the hospital applies an (R, s, Q) policy, an order is only placed once demand falls below a certain threshold. Therefore, once an order is placed, the preceding period can be considered to have had the demand of the order quantity. This gives a quick translation between orders and

Table 2.6: *Weekly orders for three products, aggregated over all relevant departments.*

Artikel	Week 1	Week 2	Week 3	Week 4	(...)	Week 50	Week 51	Week 52
3000001	15	25	19	12	(...)	13	15	15
3000002	366	400	300	1066	(...)	400	0	500
3000003	5	1	2	0	(...)	0	3	3

demand. However, there are several risks with taking this approach that could limit the validity of the results. We discuss them below.

High order quantities

The first issue is that orders are placed in fixed quantities. When these quantities are relatively low, meaning that the corresponding products are ordered often, a problem is not induced. If WKZ places an order on day 1 and another one on day 4, we can reasonably assume that the order from day 1 was translated to actual demand in the department. However, high order quantities lead to situations where demand is misrepresented and variability is overestimated. We explain this below. The data shows that some products are ordered very infrequently, but once an order is placed, it is of a large quantity. These products are classified as intermittent or lumpy when using the ADI-CV analysis. An example of two such products can be found in Table 2.7. For the first product, it is unlikely that the actual demand during weeks 9 through 14 was zero, and in week 15 equal to 80. Discussing the details and experiences that an employee in the WKZ on the product reinforces this hypothesis. Given the high order quantity compared to average weekly demand, the reorder point is only reached after a few weeks.

Table 2.7: *Weekly orders for two products with a relatively high fixed order quantity.*

Artikel	Week 5	6	7	8	9	10	11	12	13	14	15	16	17	18
3000129	0	0	0	80	0	0	0	0	0	0	80	0	0	0
3000182	0	0	8	0	0	0	8	0	0	0	0	0	0	8

However, for some products, it could be the case that demand is actually fluctuating. For example, products placed in the theme carts, identified in Section 2.3.1, are used infrequently but in high quantities. Therefore, it is necessary to assume that for all products, these orders reflect the actual demand. One could consider spreading the demand out over the preceding weeks to reduce the variance in demand per week. However, this will cause the variance to be underestimated, since demand is not equal in every week before an order is placed. In some weeks, WKZ still observes higher demand than in others.

Waste

Products with a short expiration period are thrown away without being used to meet demand once they become obsolete. However, with the suggested method, we incorrectly identify the wasted products as demand. We assume that the number of wasted products due to expiry dates is low, and is thus kept out of scope.

Susceptibility to error

Another issue is that once an order is placed, the underlying reason for placing this order is not always clear. We implicitly assume that the ordering policy is perfectly adhered to. However, given that the policy is executed by humans, there are likely some deviations. The possible deviations are outlined below, and identified in conversation with employees and managers.

- **Precautionary orders**

  In some cases, employees order more of a certain product because they foresee a high usage. This could be due to the fact that the weekend is coming up where no orders are placed, or a care provider notices that more of a certain product is being used than usual. For these orders, it cannot be verified whether the demand was actually higher in the end. If not, it also leads to situations where no orders are placed given the already high inventory.

- **Incorrect delivery**

  Deliveries are not always completely correct, and might contain more or less of certain products. This is difficult to identify, since it is not workable to count every order for correctness.

- **Incorrect count**

  For products with a high reorder point, it can be difficult to count whether it has been reached exactly. It would cost too much time for an employee to count every item, so an estimation is made.

It is unknown how often these errors occur and what the impact is. Therefore, we keep it out of scope and assume that everything goes according to what the policy describes.

Summarizing, we showed some assumptions that need to be taken regarding the available data. The leading assumption is that the aggregated data on orders can be translated directly to demand. Working with these assumptions means that we often overestimate demand. Preliminary analysis shows that only 14 products have low enough variance of lead time demand to be able to use regular inventory management methods. Because of the high variances, safety stock level recommendations will likely be higher than necessary to achieve a certain service level. However, the methods that we will employ for making the calculations remain valid. We will show the usage for products that are likely to have more accurate demand data. Since smooth and erratic products are ordered frequently, the overestimation of variance due to high order quantities is lower than for intermittent or lumpy products. As such, the recommendations for these products are most applicable. The recommendation will be to gather data on daily demand of the other products for a few months, and apply the models again to obtain more accurate results.

2.4.2 Data cleaning

In general, the quality of data managed by the logistics department is high. However, there are issues when dealing with the batch sizes and price of products. We elaborate on these issues in this section.

Batch sizes

Many products in the WKZ are not purchased in units, but rather larger batches such as boxes or trays. How many units are contained in such a package is stored inconsistently. Most product descriptions are superseded by parentheses which include the amount. However, in some cases, the parentheses are not closed, the number is stored in text, or multiple parentheses are included. There are some additional special cases that need to be taken out. Most of the data could be cleaned by the Python script shown in Appendix B.1. Some final cleaning was done by hand. Given the large number of products, not every product could be verified for correctness. It is also unclear whether there are products that are purchased in batches but for which this is not mentioned in the used dataset. We recommend the WKZ to use a more consistent system of storing the number of products in a batch. For example, by introducing an additional attribute in the database instead of writing the batch size in product descriptions.

A second issue with batch sizes occurs when looking at the weekly order data. Sometimes order quantities are shown on unit level, and sometimes on batch level. We aim to have the data on orders shown on batch level for consistency. Whether a product is shown in units or batches cannot be found in the dataset, so a guess needs to be made. We employ the reasoning that if all weeks' order quantities can be divided by the batch size of a certain product, the numbers likely indicate units. Therefore, we divide by the batch size for these products. Table 2.8 shows an example of this process.

Table 2.8: *Example product where the demand is divided by the batch size.*

Artikel	Batch size
3000001	100

Week	1	2	3	4	(...)	50	51	52
Before	1500	2500	1900	1200	(...)	1300	1500	1500
After	15	25	19	12	(...)	13	15	15

Prices

In addition, data on product prices needs some cleaning. This data is again gathered from the ERP tool. Similar to weekly demand, prices are inconsistently given per batch size or per unit. Again, no clear distinction is possible to identify for which method is used where. However, a method which works well for most products is checking whether the price of a product is smaller than 2 euros, it is too cheap to be for an entire batch. Therefore, we assume that the price is for one unit which means it needs to be multiplied by the batch size. This method was verified for the first 50 products on the list and no obscurities were found. This does not mean that there are no errors in this process. There are likely still some edge cases for which the price is incorrect. This has impact on the determination of the EOQ, since it takes holding costs as input which are expressed as percentage of the cost price. Therefore, outliers on EOQ compared to the current ordering quantity need to be checked on a case-by-case basis. Again, we recommend the WKZ to store prices more consistently for future usage.

Daily data

For some analyses, data on daily orders is necessary. This is also available in the ERP tool of the WKZ. However, only days where an order is placed for at least one product, are included. For example, no orders are placed in weekends. We assume that all demand in such days is equal to zero, which is added in a new column per date. This assumption is reasonable, since the products that we will use it for are either intermittent or lumpy. This means that demand occurs infrequently, so it does not significantly matter whether the demand occurred in the weekend or on the closest Monday.

2.4.3 Experimental data

Given the lacking data on actual demand, we conducted an experiment to gather this data for certain products. In this experiment, the inventory level of twelve products in one warehouse of the PICU was tracked every day for two weeks. The daily demand was computed by subtracting the placed orders. Three products of each of the ADI-CV categories was included in the experiment, to get an overview of how the different categories behave. Each product also lies in the 'A' category of ABC categorization (Section 2.3.3), so the impact of changing the parameters for these products is the largest.

The data was gathered by two different employees. We noticed that this caused some difficulties, since they used a different way of denoting the number of products. Sometimes, we observe that an employee denotes the number of batch sizes (boxes, packages, etc.), and in other cases the total number of units. In addition, counting the number of products did not always go accurately, as sometimes the inventory increased between two days whereas no order was placed.

Because of these inaccuracies, we do not use this data to compute new order quantities Q and reorder points s. However, we still have a relevant observation, elaborated upon in Section 5.1. In addition, the above remarks regarding difficulties with gathering data are useful for determining how to gather demand data in the future. We discuss our proposed way of doing this in Section 6.1.

2.5 Conclusion

This chapter analyzed the situation in which the WKZ is operating. First, we elaborated on the current process of ordering consumable products to the warehouses. Each department has one or multiple of these warehouses. Employees regularly check the inventory levels and place an order of size Q once the level is below s. In this thesis, we focus on the consumable products that are kept on stock in the central warehouse: so-called 'ZROH' products.

Next, we elaborated on the situation in the new multi-functional area (MFA). Many products are currently stocked in multiple warehouses that will be moved to the MFA. This means that the hospital will benefit from a reduction in variance of demand by aggregating these products. The ordering process remains the same.

We identified three ways to categorize products. The first one is a matrix developed by UMC Utrecht, categorizing products on frequency and amount of usage of a product. ABC analysis classifies in three categories based on average inventory value. In this thesis, we mainly use ADI-CV classification, where products are put into 4 categories based on average demand interval (ADI) and coefficient of variation (CV).

Finally, we made some remarks regarding data availability and quality. Only data on historical orders is available through the ERP system of the UMC. This incurs some issues where we overestimate variability when orders are placed in high quantities. Furthermore, data on batch sizes and product prices is occasionally stored incorrectly. We solved this by doing some data cleaning.

Chapter 3
Literature Review

This chapter presents the findings from conducting a literature review with as goal to get insight into the available techniques that exist for quantitatively managing inventory. Section 3.1 introduces the discipline of inventory management, in which this thesis operates. Sections 3.2 and 3.3 discuss the techniques that can be used for computing the optimal decisions for regular and irregular demand, respectively. Note that we focus on techniques that can be directly or almost directly applied to an (R, s, Q) policy. For the techniques for computing a new order quantity Q and reorder point s, some general concepts in inventory management are necessary. This includes the EOQ, some statistical distributions used to model demand and two measures for service level. We refer to Appendix A for the reader that is unfamiliar with these concepts. We will use EOQ as the leading way to compute order quantities for regular demand.

3.1 Theoretical framework of inventory management

This research is conducted within the discipline of inventory management. The key concepts that belong to this discipline are elaborated upon here. We focus on the conceptual principles, and not the mathematical formulation and derivation. This is treated in later sections of this chapter.

3.1.1 General concepts

Showkat and Ali (2020) define inventory as being "*composed of assets that will be sold in the future in the normal course of the business operations*". This definition does not fully align with the inventory at WKZ. Items are not meant for sale, but rather to provide a service. However, once the product is used by a healthcare provider for treatment, this can be considered a sale.

Inventory management refers to making decision on three issues, as identified by Silver et al. (2021). They are as follows:

1. How often to determine the inventory status.
2. When to place the replenishment order.
3. How large the replenishment order should be.

The first two issues belong to *timing decisions*, as identified by Slack and Brandon-Jones (2019). The third issue is part of so-called *volume decisions*. Finally, the authors also mention *control decisions*. This refers to the models and methods that are used in making the timing and volume decisions.

3.1.2 Inventory types

Timing and volume decisions have to be based on a lot of factors. To improve the focus of the decisions, Silver et al. (2021) present six different types of inventory. For this research, three types are most important:

1. **Cycle inventory**
   Cycle inventory is the inventory that results from "*ordering batches instead of one product at a time*" (Silver et al., 2021). It is the amount of inventory that is ordered in a so-called *replenishment cycle*. Depending on the type of policy, replenishment cycles are constant or variable.

2. **Safety inventory**
   This is the inventory aimed to "*counter uncertainty in demand*" (Shenoy and Rosas, 2017). Often, demand cannot fully be determined for the upcoming replenishment cycle. In the case that demand turns out to be higher than expected, safety stock serves to reduce the chance of running out of stock. Safety stock is often denoted by SS.
3. **Anticipation inventory**
   Anticipation inventory arises when an organization foresees a rise in demand for the coming period. This could be due to seasonal effects or promotional periods. The organization orders extra inventory to compensate for this rise in demand. In the WKZ, anticipation inventory could occur during the flu season when more children get sick, for example.

An illustration of these types of inventory can be seen in Figure 3.1.

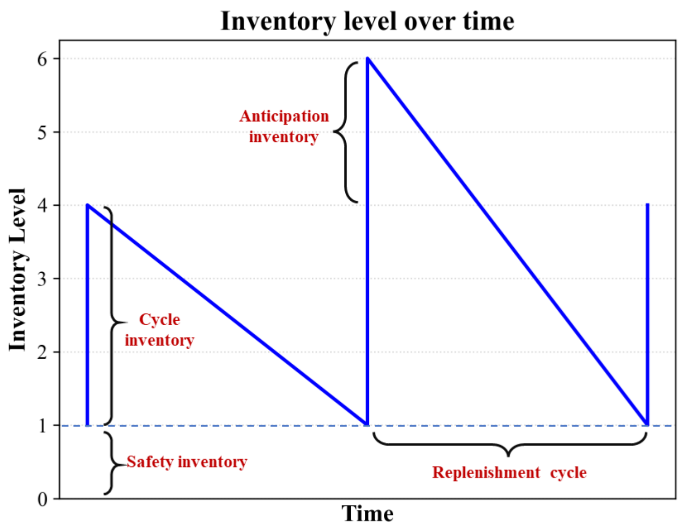

Figure 3.1: *Illustration for the key inventories identified by Silver et al. (2021).*

3.1.3 Variables and performance measures

A wide variety of decisions can be made when considering an inventory policy. These decisions serve as input to an inventory model, which have an impact on the resulting performance measures. They are discussed in this section. We refer to Silver et al. (2021) and Shenoy and Rosas (2017) for more detailed elaboration on all concepts.

Order quantity refers to the size of an order that is placed, and is often denoted by Q. An optimal order quantity can often be computed. If this optimum is reached, we speak of an economic order quantity (EOQ).

Lead time refers to the time that it takes for an order to arrive after it has been placed. We denote lead time by L. Lead times can be deterministic (known with certainty) or stochastic (unknown to some extent). In the WKZ, they are deterministic for products stored in the central warehouse.

Demand is the amount of product that a customer requests over a certain period of time. Demand, denoted by D, is almost always stochastic in nature. Another important concept is **lead time demand**. This is the demand that occurs during the lead time L, and denoted by D_L.

Reorder point is the inventory level at which an order is placed, denoted by r. A part of the reorder point is always the lead time demand D_L. The remainder is safety stock SS (i.e. $r = D_L + SS$). A **Stockout** occurs once there is not enough inventory available of a certain product.

Service level refers to the extent to which agreements with customers are being met. There are several ways to measure service level, which is elaborated upon in Section A.3.

Inventory costs refer to the total monetary impact of maintaining an inventory policy for a unit of time. We mention a few important different types of costs. **Unit costs** is the cost of one unit of inventory of a certain product. **Holding costs** (h), refer to the costs of holding one unit of stock for one unit of time. They are often expressed as percentage of the unit costs. **Ordering costs** (K) are the costs of placing one order. We elaborate on these costs in Section A.1. We define inventory costs as holding costs plus ordering costs.

3.1.4 Inventory policies

An **inventory policy** is the set of rules that an organization applies to decide when to place their orders, and in which quantity. We distinguish between **continuous review** and **periodic review** policies, in addition to **order-up-to** and **fixed order quantity** policies. Continuous review policies assume that we always have insight into the current inventory level, and place an order immediately once the inventory level is equal to r. In periodic review policies, we place an order at fixed times, with T units of time between orders. In order-up-to policies we order that amount of inventory to end up with an inventory level of S (order-up-to-level) once the order arrives. In fixed order quantity policies, we always place an order of Q units.

Popular policies that arise from these concepts are the (r, Q) policy, in which an order of Q units is placed once the inventory reaches r units, and the (T, S) policy, in which we review every T units of time and order up to S units. The policy that is used in the entire UMC Utrecht is an (R, s, Q) policy. This means that every R time units, the inventory is reviewed. An order of size Q is placed if the inventory is lower than s units. For most warehouses, inventory levels are reviewed every day.

We also distinguish between **backorder** and **lost-sales** models. In the former, we place an item on backorder if we run out of stock, and the sale can still occur later. This is can be interpreted as a negative inventory level. In the lost-sales model, the demand is lost for good once the customer is faced with a stockout.

3.2 Techniques for modelling regular demand

This section explores the techniques that exist for modelling inventory subject to regular demand. We define regular demand as demand for which $ADI < 1.32$ (Section 2.3), so this includes both smooth and erratic demand. The goal is to find an optimal value for s and Q, which minimizes costs with a constraint on service level. For this demand, the EOQ is a relatively good way to determine a value for Q (Shenoy and Rosas, 2017). Therefore, we focus on computing an appropriate level for the reorder point s. Johansen and Hill (2000) develop a cost-oriented approach of solving this. Although this thesis works with a service level-oriented approach, several concepts from Johansen and Hill (2000) can be applied.

3.2.1 Smooth demand

When demand is smooth ($CV < 0.49$) we can reasonably assume that demand is normally distributed (Silver et al., 2021). Note that this value is computed on demand during lead and review time. The identification of categories in Section 2.3.2 used weekly demand to obtain an estimation, which therefore still needs to be transformed to the right time unit. This can be done by using that if demand in one period $D \sim \mathcal{N}(\mu, \sigma)$, lead time demand $D_L \sim \mathcal{N}(L\mu, \sqrt{L}\sigma)$. This property also holds for non-normal demand distributions, and will also be used for erratic demand.

To obtain a good safety stock, we need to consider the demand that occurs after an order has been placed. For the remainder of this section, we assume that $D_L \sim \mathcal{N}(\mu_L, \sigma_L)$. For normally distributed demand, ESC can be computed using the standard normal loss function $G(z)$. The following relation holds:

$$ESC = \sigma_L \cdot G(z) \tag{3.1}$$

with

$$G(z) = \phi(z) - z \cdot [1 - \Phi(z)] \tag{3.2}$$

where σ_L is the standard deviation of lead time demand (Silver et al., 2021). Here, z refers to the 'safety factor', indicating how many standard deviations of demand we want to include in the safety stock. Thus, $z \cdot \sigma_L$ gives the safety inventory.

Substituting ESC in Equation A.8 for $\sigma_L G(z)$ in Equation 3.1 and rewriting leads to the following relation:

$$G(z) = \frac{(1 - P_2) \cdot Q}{\sigma_L} \tag{3.3}$$

with P_2 being the fill rate that we want to achieve. We can compute the appropriate z by filling the right-hand side of Equation 3.3 in the inverse normal loss function. No closed form of this equation exists, so root-finding algorithms have to be used to get arbitrarily close to the right z. In this thesis, we use a bisection algorithm. Once z is computed, we compute s by adding the mean lead time demand to the safety inventory $z \cdot \sigma_L$:

$$s = \mu_L + z \cdot \sigma_L. \tag{3.4}$$

Remark. Johansen and Hill (2000) develop a way of considering the impact of having a nonzero review time. An undershoot is incurred when demand has fallen below the reorder point, but the inventory has not been reviewed yet. Concepts from renewal theory are used in this paper. Given that there is a gap in literature for modelling undershoot for erratic, intermittent and lumpy demand, it will not be used in this thesis. Chapter 4 discusses how the review period will be dealt with instead.

3.2.2 Erratic demand

In case of demand with an erratic nature, a normal distribution cannot be applied. Hence, the method of deriving an appropriate safety factor should be changed. Babiloni et al. (2010) argue that it is appropriate to use a Gamma distribution to model lead time demand when $CV > 0.49$. By knowing the distribution of demand, we can compute ESC for a given reorder point s. Using this, we can iteratively get closer to the s which gives an appropriate value for ESC and thereby P_2. First, the general formulation of ESC is necessary, formulated by Silver et al. (2021):

$$ESC(s) = \int_s^\infty (x_0 - s) f_x(x_0) dx_0. \tag{3.5}$$

Here, f_x refers to the PDF of the considered distribution. The equation can intuitively be interpreted as considering all demand during lead time greater than the reorder point, and multiplying the shortage that this demand incurs by the probability of it happening.

Silver et al. (2021) rewrite this equation for the Gamma distribution in a more tractable form:

$$ESC(s; \alpha, \beta) = \frac{\alpha}{\beta}\big(1 - F(s; \alpha + 1, \beta)\big) - s\big(1 - F(s; \alpha, \beta)\big). \tag{3.6}$$

where F is the CDF of the Gamma distribution and s the reorder point. After this, Equation A.8 can be used to derive P_2 for a given s. To find s for a predetermined P_2, an inverse of $ESC(s)$ would be required. Again, no closed form of this inverse exists, meaning we again use a bisection algorithm to iteratively get closer to the correct s. Given that Equation 3.6 is monotonically decreasing, this algorithm can easily be applied. $ESC(s)$ is monotonically decreasing in s since:

$$\forall s : \forall \epsilon > 0 : ESC(s) = \int_s^\infty (x_0 - s) f_x(x_0) dx_0 = \int_s^{s+\epsilon} (x_0 - s) f_x(x_0) dx_0 + \int_{s+\epsilon}^\infty (x_0 - s) f_x(x_0) dx_0$$
$$\geq \int_{s+\epsilon}^\infty (x_0 - s) f_x(x_0) dx_0 \geq \int_{s+\epsilon}^\infty [x_0 - (s+\epsilon)] f_x(x_0) dx_0 = ESC(s+\epsilon).$$

3.3 Techniques for modelling irregular demand

This section explores the techniques that exist for modelling inventory subject to irregular demand. This area of inventory management has seen a lot of development in recent years, but is less mature than that of regular demand. To get an overview of the available techniques, we first discuss our findings of a literature review (Section 3.3.1). Next, we elaborate in detail on the SPOZA framework proposed by Finco et al. (2022), which will be used for this thesis (Section 3.3.2).

3.3.1 Overview

First, we provide an overview of the recent developments in the field of managing irregular demand. This area of research is relatively new, with many different approaches and techniques currently being researched. We aim to find a technique that is both effective and easy to implement.

A popular method of managing intermittent demand that is used by Babiloni et al. (2010), Strijbosch et al. (1999), Hasachoo et al. (2019), Xu et al. (2012) and Kalaya et al. (2019) is the so-called "Croston's method". This technique aims to forecast the demand for upcoming periods by providing an estimation of the average time between two positive demands and the average positive demand. The general idea of the process is to run through the historical demand data, and updating the parameters if we observe positive demand. The parameter is updated based on how different our observation is from what we expect. To make sure that the updates are not too radical, we multiply the update by a smoothing constant $\alpha \in (0, 1]$.

Another technique that many authors use is that of fitting the data to a statistical distribution. The idea here is the same as with the techniques given in Section 3.2, but now using different distributions than the Normal or Gamma distribution. Babiloni et al. (2010) mentions the negative binomial, compound Poisson and compound Bernoulli distribution as suitable for intermittent demand products. The article refers to several other articles where these methods are applied. For example, Strijbosch et al. (1999) is mentioned for using a Compound Bernoulli method to manage an (s, Q) policy. Dunsmuir and Snyder (1989) develop a way of using the Gamma distribution in a general way to determine re-order levels. Lolli et al. (2017) uses a wide variety of density functions to fit the demand distribution, through a statistical software package. Wang and Xiao Xia (2015) give a case study where the Normal, Gamma and compound Poisson distributions are used, depending on the type of product. Simulation is used afterwards to study the effectiveness of the pre-determined policies. Simulation is also used as a method by Strijbosch et al. (1999) to "*study the performance of the* [...] *inventory control procedures*". Finco et al. (2022) develop a framework in which a zero-inflated Poisson distribution is used, which can be used for both lumpy and intermittent demand.

Finally, some authors use more advanced and specialized techniques. Mak et al. (1999) use genetic algorithms to iteratively improve the policy based on its previous performance. Lolli et al. (2019) apply a neural network to classify products and find the respective optimal policy accordingly. Hasachoo et al. (2019) use Integer Linear Programming (ILP) and Kalaya et al. (2019) use a combination of exponential and Poisson distributions to model the demand.

We choose to work with the technique identified by Finco et al. (2022) in this thesis. This is the so-called 'SPOZA' framework. Finco et al. (2022) perform a case study with the framework and note that implementing it gives promising results with lower inventory costs than other, more classical approaches such as Croston's method. In addition, the framework is relatively easy to implement (Finco et al., 2022). We focus on this framework for the remainder of this section. For further research, more techniques can be implemented and compared to one another.

3.3.2 SPOZA framework

This section discusses the SPOZA framework proposed by Finco et al. (2022) for managing intermittent and lumpy demand. The authors use a zero-inflated Poisson (ZIP) distribution to model lead time demand. SPOZA stands for "spare part oriented ZIP approach". The reason for this is that spare parts are a "*typical example of irregular demand items*" (Finco et al., 2022).

Zero-inflated Poisson regression

The central distribution in the SPOZA framework is a zero-inflated Poisson distribution. The process generates values using a combination of a binomial and Poisson distribution. The binomial distribution generates the number of periods in which we observe positive demand during lead time, whilst the Poisson distribution generates the size of this demand. Therefore, the ZIP distribution has two parameters: p and λ. Here, p denotes the probability of getting a period with zero demand and λ the rate of the Poisson distribution. These parameters are estimated for a single time period. If the lead time demand consists of multiple time periods, we can use the parameter n of the binomial distribution as the number of periods that are included in the lead time.

In r.v. X follows a ZIP distribution, we write $X \sim ZIP(p, \lambda)$. The density function of the ZIP distribution is as follows:

$$\mathbb{P}(X = i) = \begin{cases} p + (1-p)e^{-\lambda} & i = 0 \\ (1-p)\frac{\lambda^i e^{-\lambda}}{i!} & \text{otherwise.} \end{cases} \tag{3.7}$$

Beckett et al. (2014) derive that the moment estimators of the ZIP distribution are as follows:

$$\hat{\lambda} = \frac{\hat{\sigma}^2 + \hat{\mu}^2 - \hat{\mu}}{\hat{\mu}} \tag{3.8}$$

$$\hat{p} = \frac{\hat{\sigma}^2 - \hat{\mu}}{\hat{\sigma}^2 + \hat{\mu}^2 - \hat{\mu}} \tag{3.9}$$

where $\hat{\mu}$ denotes the sample mean and $\hat{\sigma}$ denotes the sample standard deviation of demand during a single period.

Beckett et al. (2014) also indicate that when the mean $\hat{\mu} \geq \hat{\sigma}^2$, the estimators can become negative. In that case, they argue that $\hat{p}$ should be adjusted to 0 and $\hat{\lambda}$ to $\hat{\mu}$.

Determining reorder point

Once the parameters of the ZIP distribution are determined, the reorder point for a certain P_1 can be determined. If lead time consists of $n > 1$ periods that the parameters are estimated for, Finco et al. (2022) show how to compute the distribution of this demand. Note that we use the notation $D(i; X)$ for the probability that distribution D under parameters X has realization i.

$$\mathbb{P}(X = 0) = B(0; n, p) + \sum_{i=1}^{n} B(i; n, p) \cdot \text{POIS}(0; i \cdot \lambda) \tag{3.10}$$

$$\mathbb{P}(X = j) = \sum_{i=1}^{n} B(i; n, p) \cdot \text{POIS}(j; i \cdot \lambda). \tag{3.11}$$

The authors mention that the equations use an approximation by *"considering the ZIP method as the sum of two independent steps in more than one unit of time"* (Finco et al., 2022). However, they mention that this is acceptable. Once this distribution is determined, the reorder point can be determined by using the inverse cumulative distribution function. Since the distribution is discrete, this can be done analytically using a computer.

Additional safety inventory

Since a Poisson distribution has the characteristic that its mean is equal to the variance, the reorder point can be underestimated for products with demand with a high variance. This makes the distinction between intermittent and lumpy products less apparent. Finco et al. (2022) mention the option to include additional safety stock for certain products. This is done by computing the over admitted variability (OAV) if the policy were applied on historical data. The OAV for a certain period refers to the amount that would have been short in that period. Computing this value gives an empirical indication of how the policy performs.

3.4 Conclusion

This chapter identified several literature streams that are relevant to the main research problem of reducing inventory levels. The first stream is computing an order quantity for products. We identified one common way of doing this: the economic order quantity EOQ, elaborated upon in Appendix A.1. This concept can be reasonably applied for smooth and erratic demand. We found no similar way of determining an order quantity for intermittent and lumpy demand, which might indicate a gap in literature. The next stream is computing a reorder point for regular demand. The techniques for doing this can be found in many sources, but the most extensive ones used in this thesis are Silver et al. (2021), Slack and Brandon-Jones (2019) and Shenoy and Rosas (2017). Johansen and Hill (2000) apply the techniques to the more specific (R, s, Q) inventory policy. The final literature stream is computing a reorder point for irregular demand. The techniques here are less developed and newer than the generally well-known and accepted techniques for managing the reorder point of regular demand. Finco et al. (2022) develop a technique that uses a ZIP distribution, which we will use in this thesis.

This thesis fills the gap of combining all these techniques into a central body of work and applying it to a practical situation. Using ADI-CV classification, we take specific models and adjust them for each demand category. We could not find many previous attempts at this in literature. Finco et al. (2022) also use ADI-CV classification, but apply the same technique to each demand category. Babiloni et al. (2010) define a more specific technique per category, but do not apply it in a case study.

Chapter 4
Model Development

In this chapter, we develop the framework that we use to compute s and Q for the products in scope. These are the consumable products on departments that will be moved to the MFA and are kept on stock in the large warehouse. We do this by combining the findings obtained from literature in Chapter 3. We first explain how the items are categorized, using the ADI-CV method identified in Section 2.3.2. Next, we define a method for managing the reorder point regular demand which requires a distinction in the use of distribution that is adapted from Silver et al. (2021). Finally, we develop a method for managing the reorder point irregular demand based on the model from Finco et al. (2022). There, intermittent and lumpy demand are largely managed using the same technique, but we introduce an additional safety factor for lumpy demand.

4.1 Determining product categories

This section discusses how categorization can best be done to distinguish products for the remainder of the research and apply different techniques to them. We use the ADI-CV method for this, given that preliminary analysis showed that the data contains a lot of weeks with zero demand (Section 2.4). This suggests that quite a lot of products follow irregular demand patterns.

We first compute the estimated mean and standard deviation of weekly positive demand using Equations 2.3 and 2.4. We then calculate this back to lead time demand (with $L = \frac{2}{7}$ weeks), by using the property that if mean and standard deviation of demand are μ and σ for 1 period, the mean and standard deviation of demand are $L\mu$ and $\sqrt{L}\sigma$ for L periods, respectively. Note that we restrict ourselves to positive demand only for determining the appropriate categories. Once reorder points need to be computed, we also consider the zero demand values to estimate μ and σ. We compute ADI according to Equation 2.1. The code that we use for computing CV and ADI can be found in Appendix B.2.

Remark. Even though actual lead time in the WKZ is 1 day, and often even shorter, we use a lead time of 2 days. This is since there is also a review period of 1 day that needs to be considered. Models exist that include this review period in a more accurate way, but we could not find them for erratic, intermittent or lumpy demand. Therefore, we use a risk-averse scenario where the entire review period is included in the lead time. Therefore, we make sure that we do not underestimate demand which could lead to shortages.

4.2 Inventory model for products with regular demand

This section combines the findings from the literature review in Chapter 3 to make a combination of models with the purpose of managing regular (smooth and erratic) demand. The models are implemented in Python, for which we refer to Appendix B.4.1 and B.4.2.

For both smooth and erratic demand, we first need to compute the estimated mean μ and standard deviation σ of lead time demand. Here, every period of demand is considered, so also periods with zero demand. In our specific case, we have $n = 52$ periods of one week, so the standard estimators are:

$$\hat{\mu} = \frac{1}{52}\sum_{i=1}^{52} X_i \qquad \hat{\sigma} = \frac{1}{51}\sum_{i=1}^{52}(X_i - \hat{\mu})^2 \tag{4.1}$$

where X_i denotes the demand in period i. Next, we use the holding cost rate h and ordering costs K to compute the EOQ for the product using Equation A.2. In this thesis, we use that holding cost rate is equal to 0.4 and ordering costs are €2.5. For the determination of these parameters, we refer to Section 5.2. After this, we use the previously computed CV to distinguish smooth and erratic demand.

4.2.1 Smooth demand

If $CV < 0.49$, the demand can be considered as smooth. This means that a normal distribution of is reasonable to assume. The parameters of this distribution are estimated using the moment estimators of the normal distribution. Afterwards, a desired fill rate needs to be determined. This research uses a fill rate of 99%, which was discussed in collaboration with stakeholders from the WKZ. Next, we use Equation 3.3 to determine for which safety factor z this fill rate is achieved. This is done by first computing $\frac{(1-P_2)\cdot Q}{\sigma_L}$ and looking for the right z that is mapped to this value under the normal loss function G. We call the desired fill rate the 'target' (t).

Given that no inverse of G exists, this is done iteratively through using a bisection algorithm. The main idea is to split the graph of $G(z)$ up into two parts and check in which part the correct z lies. This part is then again split up into two parts, etc. An estimate of the correct z, denoted by m, is then computed as the middle of the line that we know z has to lie in. We keep on bisecting the graph until the difference between our current value and the value we are looking for, $|G(z) - G(m)|$, is smaller than a certain tolerance ϵ. Algorithm 1 shows the pseudocode of this algorithm. We largely adapt the bisection algorithm from Atkinson (2008) and make some minor changes to fit to the properties of $G(z)$. For example, in the formulation of Atkinson (2008), the goal is to find a root of the function. However, in our case, we are not looking for a root but a value for which $G(z) = t$. Since G is monotone, this value for z is unique. In addition, we can check in which of the two parts z lies by simply evaluating whether $G(m) > t$ or $G(m) < t$.

Algorithm 1 Bisection method for smooth demand

```
Function: G (normal loss)
Tolerance: ϵ
Target fill rate: t
Left bound of region: a
Right bound of region: b                    ▷ with G(a) ≤ t ≤ G(b)

m ← (a + b)/2

while |G(m) − t| > ϵ do
    if G(m) > t then
        a ← m
    else                 ▷ Then certainly G(m) ≤ t, i.e. t ∈ [G(a), G(m)] given monotonicity
        b ← m
    end if
    m ← (a + b)/2
end while
return m
```

This algorithm results in a safety factor z for which we attain the correct fill rate P_2. With this, a safety inventory and reorder point can be computed using Equation 3.4.

4.2.2 Erratic demand

This section discusses how erratic demand is handled in this thesis. Demand is erratic when $CV > 0.49$. In this case, the distribution that can be used to model lead time demand is a Gamma distribution. We can estimate the parameters of this distribution using the moment estimators of a Gamma distribution (Section A.2). Again, a desired fill rate is determined (99% in this research).

Equation A.8 gives a way to translate between ESC and this desired fill rate. Together with Equation 3.6, this can be used to determine ESC for a certain reorder point s, after which the fill rate corresponding to s can be determined. This can be done using an algorithm similar to Algorithm 1. However, whereas safety factor z could be directly filled in to the normal loss function there, we now have to first compute ESC and then fill that in to Equation A.8. Algorithm 2 shows the pseudocode of the algorithm. Again, we adapt the general structure of the algorithm from Atkinson (2008). We adjust the algorithm by adding some extra steps that translate between the Gamma loss function G and fill rate P_2.

Algorithm 2 Bisection method for erratic demand

```
Function: G (Equation 3.6)
Tolerance: ϵ
Target fill rate: t
Order quantity: Q
Left bound of region: a
Right bound of region: b                                   ▷ with a ≪ b

m ← (a + b)/2
ESC ← G(m)
P2 ← 1 − ESC/Q

while |P2 − t| > ϵ do
    if P2 > t then
        a ← m
    else
        b ← m
    end if
    m ← (a + b)/2
    ESC ← G(m)
    P2 ← 1 − ESC/Q
end while
return m
```

Figure 4.1 shows a graphical representation that summarizes the steps that need to be taken for deriving the optima for regular demand.

4.3 Inventory model for products with irregular demand

This section discusses the framework that we use for managing irregular demand, so when $ADI > 1.32$. We adopt a part of the framework suggested by Finco et al. (2022).

The fact that product demand is classified as intermittent or lumpy indicates that the products are ordered with a low frequency. As explained in Section 2.4, this can be since products are ordered in high quantities. Given that this leads to high average inventory levels that significantly affect the space usage, we take Q as the average demand in two weeks. This means that products will be, on average, ordered once every two weeks. In conversation with stakeholders at the WKZ, this is considered to be a good balance between order frequency and quantity. In case of obligations related to minimum order quantities from suppliers or packaging sizes, the order quantity needs to be rounded to the nearest multiple of this amount. Data on obligations from suppliers is out of scope for this thesis and will therefore not be considered.

Next, we use data on daily instead of weekly orders to obtain a more accurate insight into when orders are placed exactly. We compute the estimates of the mean and variance over this data. Section 3.3 mentioned that the ZIP distribution is appropriate to use for products with irregular demand. Using the estimated mean and variance we can estimate the parameters of the distribution using Equations 3.8 and 3.9.

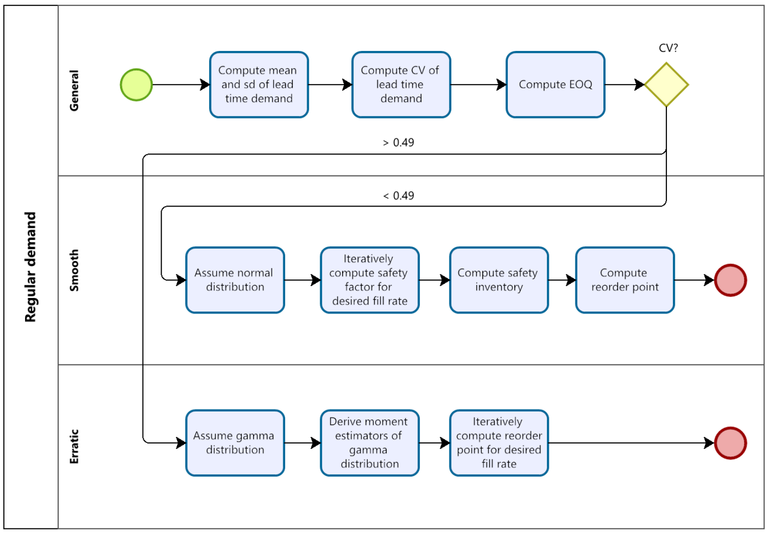

Figure 4.1: *Process for computing optima for regular demand.*

Once the parameters of the distribution are estimated, Equations 3.10 and 3.11 can be used to compute the probability of seeing a certain demand during lead time (so with $n = 2$). We want to find the smallest value j for which the probability that our demand is smaller or equal to j is greater than the desired P_1. If we place an order at this point, the fraction of replenishment cycles that incurs no backorders will be at least equal to P_1. Therefore, this is the appropriate reorder point. Algorithm 3 shows the pseudocode of how we compute this reorder point.

Algorithm 3 Computing reorder point using ZIP distribution

Desired cycle service level: P_1
$p \leftarrow 0$
$j \leftarrow 1$

Compute $\mathbb{P}(X = 0)$ ▷ using Equation 3.10

while $p < P_1$ **do**
  Compute $\mathbb{P}(X = j)$ ▷ using Equation 3.11
  $p \leftarrow p + \mathbb{P}(X = j)$
  $j \leftarrow j + 1$
end while

return j

4.3.1 Lumpy demand

In the case of lumpy demand, the variance of demand is underestimated as the Poisson distribution does not have a separate parameter for including variance. To mitigate the risk of working with an underestimation, we use an empirical method to determine possible additional safety stock. By using data on daily orders (assumed to be analogous to demand), we can assess the cases in which a backorder would have occurred when applying the policy that was previously used. It is valid to use data on daily orders, since the products are ordered very infrequently. It therefore matters whether an order was placed in the beginning of the week, or at the end. We apply a method similar to the method of computing over-admitted variability (OAV) in Section 3.3.2. However, we develop a method to derive an empirical estimator of ESC, consistent to the performance measures used for regular demand. For this, we first use the order quantity Q equal to the mean demand per week that is recommended in Section 4.3. This gives an indication of the expected fill rate if the entire model in this chapter is applied. For completeness, the Q that the WKZ currently uses will also be used. We assume that the inventory at the start of the demand period (2023) is equal to $Q + s$. Given that review period is now already included in the model, we consider lead time to be equal to one day. The process of finding the empirical ESC is shown in Algorithm 4. The algorithm runs through the historical orders as if they were demand, and computes the inventory level on each day if an (R, s, Q) policy were applied using our suggested parameters. If a shortage occurs or inventory is replenished, ESC is updated according to the size of the shortage.

Algorithm 4 Computing empirical ESC for ZIP distribution

```
Demand on day i: D_i                                ▷ i ∈ {1, 2, ..., 365}
Incoming order on day i: O_i                        ▷ i ∈ {1, 2, ..., 365}
Order quantity: Q
Reorder point: s
Inventory level at time i: x_i

x ← Q + s
i ← 1
ESC ← 0
n ← 0

while i ≤ 365 do
    x_i ← x_{i−1} − D_i + O_i
    if x_i ≤ s then
        O_{i+1} ← Q
        if x_i ≤ 0 then
            ESC ← n/(n+1) · ESC + |x|/(n+1)
        else
            ESC ← n/(n+1) · ESC
        end if
        n ← n + 1
    else
        D_{i+1} ← 0
    end if
    i ← i + 1
end while
```

The resulting ESC can be used to compute an empirical fill rate. It is computationally infeasible for this thesis to compute the exact reorder point which would lead to the desired empirical fill rate. However, as an additional safety measure, we add the empirical ESC to the reorder point if the fill rate is lower than desired. In addition, we recommend the WKZ to study the lumpy products in more detail, to see where the high variance is coming from.

Summarizing, Figure 4.2 visualizes the process of obtaining s and Q for irregular demand.

4.4 Conclusion

This chapter developed the model that is used to determine the optimal order quantity Q and reorder point s for the stock consumables in the WKZ. We take data on weekly orders to estimate the mean $\hat{\mu}$ and variance $\hat{\sigma}$ of weekly demand of every product in scope. In addition, we compute ADI and CV to put products into four categories: smooth, erratic, intermittent, and lumpy. Depending on the category, we apply a different technique to compute the optima.

For regular (smooth and erratic) demand, we use EOQ as the order quantity. The estimation of the necessary parameters is elaborated upon in Section 5.2. For irregular (intermittent and lumpy)

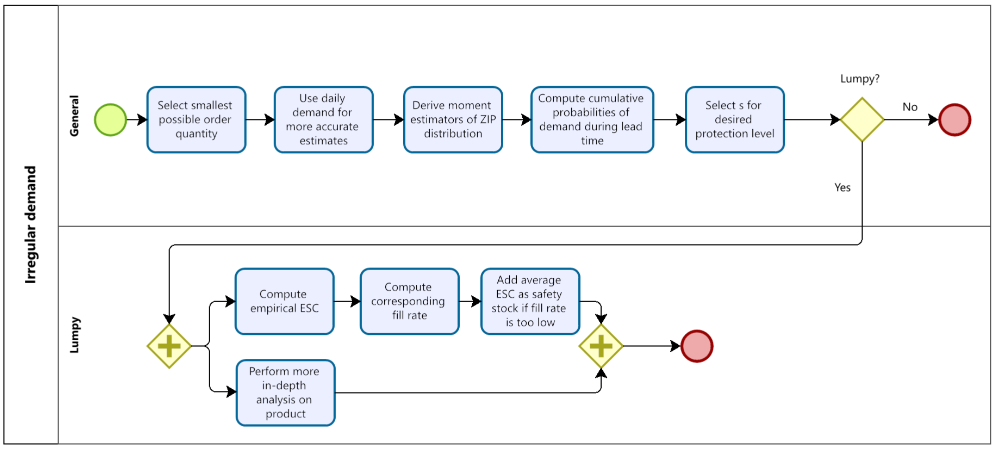

Figure 4.2: *Process for computing optima for irregular demand.*

demand, we use the average bi-weekly demand as order quantity. These values for Q are used as input for determining s. This goes as follows:

If demand is smooth, we can model it by using a normal distribution. We estimate the parameters of the distribution using the standard estimators $\mu = \hat{\mu}$ and $\sigma = \hat{\sigma}$. Using a bisection algorithm, we look at the right safety factor k for which we attain a desired fill rate P_2 (we use $P_2 = 0.99$ in this thesis). This is done using the normal loss function.

If demand is erratic, we can model it by using a Gamma distribution. We estimate the parameters of the distribution with $\alpha = \hat{\mu}^2/\hat{\sigma}^2$ and $\beta = \hat{\mu}/\hat{\sigma}^2$. Using a slightly adjusted bisection algorithm, we can determine the reorder point s for which we obtain $P_2 = 0.99$. This is done using the Gamma loss function.

If demand is intermittent, we use the CDF of the ZIP distribution and determine the smallest integer value of j for which it is greater than the desired cycle service level P_1. This j is equal to the desired reorder point. We developed an empirical way to compute ESC based on historical data, to verify whether using a certain P_1 approximately leads to $P_2 = 0.99$.

For lumpy demand, we use the same model as for intermittent demand. However, we also add the empirical ESC to the reorder point to deal with the higher variance in positive demand.

In Chapter 5, we apply this model to the order from 2023 data of the 527 products that fall in scope of the research.

Chapter 5
Results and Discussion

This chapter discusses the overall results and findings from applying the model from Chapter 4 to the case of the WKZ. The goal is to find a new value for order quantity Q and reorder point s in order to reduce inventory levels and shortage rates. We refer to Section 1.3.3 for the definition of which products fall in scope. As described in Section 2.4, we use data on weekly and daily orders from 2023 to perform the analysis. We discuss and interpret the results in this chapter as well.

First, we characterize the products based on ADI-CV classification (Section 5.1). Next, we elaborate on the results from determining a new order quantity for all products (Section 5.2). Finally, we discuss the results of computing a new reorder point for all products (Section 5.3). We pay extra attention to the results for irregular demand items (Section 5.3.2).

Remark. It is important to re-iterate that the leading assumption is that data on historical orders can be directly translated to demand. Therefore, when we speak of the 'demand' of products, this assumption is made implicitly.

Remark. The models are programmed in Python, after which some analysis is done in Microsoft Excel. Figures and tables are constructed using R through the ggplot2 and Stargazer libraries.

5.1 Results of product characterization

The first step of applying the model is characterizing the various products. As a reminder, this thesis focuses on consumable products that are currently kept in the central warehouse (ZROH) and to be included in the MFA. In total, 550 products from the data of 2023 fit this characterization. However, 23 of the products listed here are not in the assortment at the moment of writing this thesis anymore. We leave these products out of scope. Contrarily, out of the products in the assortment at the moment of writing this thesis, 112 are not included in the data from 2023. Note that this includes both ZROH and ZNLA products, among which the division cannot be found. Once the MFA is implemented, we expect that a larger number of products in the assortment will be different. This means that for some products, advice on the appropriate s and Q cannot be given through the use of the model. This increases the relevance of regularly running the model with up-to-date data.

Next, we apply ADI-CV categorization. Table 5.1 shows the corresponding results. Approximately 61% of the products has an irregular demand pattern. This means that the majority of products is, on average, ordered less than once every 1.32 weeks. Section 5.2 discusses possible implications of this regarding order quantity. Of the regular demand items, 94% has erratic demand. This indicates that most products in the WKZ have a high variance in demand.

Table 5.1: *Division of products in the WKZ among the ADI-CV categories.*

Category	**Demand type**	**Count**	**Percentage**
Regular	Smooth	13	2.5%
Regular	Erratic	195	37.0%
Irregular	Intermittent	152	28.8%
Irregular	Lumpy	167	31.7%

Table 5.2 shows the five-number summary in addition to the median of ADI and CV when considering the entire dataset. It is noteworthy that the first quartile of CV is equal to zero, indicating

that at least 25% of the products has zero variance in positive demand. When zooming in on the different categories, Figure 5.3 shows that almost all of these zero-variance demand patterns come from intermittent products, apart from one outlier in the smooth category. In fact, for intermittent products, it is noteworthy that the third quartile is equal to zero, with all the nonzero values of CV being outliers. We hypothesize that this effect comes from the fact that orders are placed in fixed quantities, as Section 2.4 elaborates on. The very low variance of intermittent products indicates that predicting when an order is placed will be most important, for which ADI is useful. Note that the y-axis of the CV box plots of Figure 5.3 has been cut-off at 3 to improve the visualization. This means that a few outliers are left out of the picture. The most striking one falls in the erratic category, with a CV of 16. Overall, the coefficients of variation for the different items are quite close together, as Figure 5.3 visualizes well.

In the box plots of ADI (Figure 5.3), some things are noteworthy. The low ADI of the smooth and erratic categories are expected, as this is what defines this category. The distribution is spread out quite evenly in both categories. The mean ADI of erratic demand seems slightly smaller than that of smooth demand, but the result is not statistically significant ($p = 0.086$). A larger difference can be found between the ADI of intermittent and lumpy demand, which is equal to 16.8 and 4.9, respectively ($p < 2.2^{-16}$). The standard deviation of ADI of intermittent products is also significantly larger (14.9) than that of lumpy products (4.6), with a p-value smaller than 2.2^{-16}). This difference indicates that intermittent products in the WKZ have a higher mean time between periods of positive demand. In addition, the differences in this time are bigger between intermittent products than between lumpy products. This result is not entirely surprising: many products in the intermittent category are ordered only 2-5 times per year. This means that there is less opportunity for the observed demand to be variable, which would put the product in the lumpy category. Another reason for the much higher ADI of intermittent products could be large order quantities. If a product is ordered in high quantities, the same quantity will almost always be placed. After all, it will then almost never occur that the product needs to be ordered twice in a week. This reduces the variance of the product, so it falls in the intermittent category. Because the order quantity is high, there will be a longer time between positive demand periods, leading to a high ADI.

Table 5.2: *Summary statistics for ADI and CV.*

Statistic	Min	Q1	Median	Mean	Q3	Max
CV	0.000	0.000	0.655	0.633	0.895	2.763
ADI	0.000	0.444	2.059	6.564	7.667	51.000

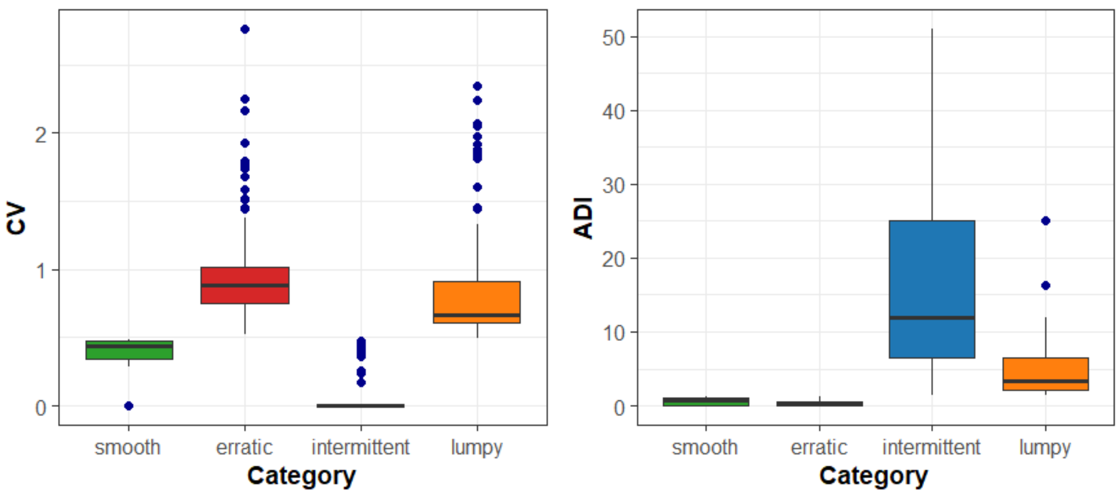

Figure 5.3: *Boxplots showing how CV and ADI are distributed amongst the categories.*

Remark. In Section 2.4.3 we mentioned the experiment that we conducted in order to gather daily demand data for twelve products over two weeks. Six of these products were classified as irregular using the weekly order data. However, when considering the gathered actual demand, we observe that five of these products were demanded at least once per week, meaning that they would have been classified as regular demand. This indicates that more products follow regular demand patterns than using order data suggests. This increases the relevance of gathering demand data in the WKZ.

5.2 Computing order quantity

Once the products are put into the appropriate categories, we calculate the new values for Q, according to the method identified in Chapter 4. We observe that using EOQ for regular demand gives radical changes compared to the current situation. As such, we formulate an alternative formulation of the optimal Q for smooth and erratic demand.

In order to get insight into the effect of using the models, we compute three statistics: the total number of yearly orders, the sum of average inventory levels, and the total inventory costs. Number of orders gives information about the impact that the model has on the amount of work that is required. Additional orders means that we incur additional handling costs. On the other hand, average inventory level gives information about the amount of space that is required. We compute the average number of orders by dividing the yearly demand by the order quantity. For average inventory, we assume that inventory decreases linearly over time, similarly to what is used in the derivation of EOQ. This means that average inventory can be obtained by computing $s + \frac{Q}{2}$. Here, s denotes reorder point and Q order quantity. Finally, total inventory costs are obtained by multiplying the average inventory levels by the product prices and the holding costs rate, and adding this to the order costs.

5.2.1 Original order quantity

First, we use the original model identified in Chapter 4, where EOQ is used for smooth and erratic demand and computed through Equation A.2. In the model, holding cost rate h is equal to 0.4. This is used since Winston and Goldberg (2004) mention that h is often chosen to lie between 0.2 and 0.4 in practice. Since space usage is very important in the WKZ and of direct influence to h, we choose the upper bound of this standard. We determine K to be equal to €2.50. This calculation is based on the average number of products included in a single truck drive, the average handling time of an order and transportation costs. Note that this is a rough estimate, and more accurate data on ordering costs could be looked for in further research. Intermittent and lumpy demand get assigned a Q of the average bi-weekly demand. Table 5.5 shows the summary statistics of how the new Q changes with respect to the current Q. Figure 5.4 shows that for smooth and erratic demand, the change is quite variable. This indicates that it is possible that EOQ does not work desirably in the WKZ. This could be since in the WKZ, space usage is a more accurate reflection of holding costs, rather than unit costs. Therefore, we propose an alternative method of deriving Q in Section 5.2.1. We use both methods in this thesis.

Intermittent and lumpy products both show a median relative decrease in order quantity of approximately $50 - 60\%$. Interestingly, the median for lumpy products is much smaller than the mean: -0.57 compared to -0.25. This indicates that there are some big outliers that heavily influence the mean. Figure 5.4 confirms this. Note that some outliers are not included for a better visualization of the boxplots. Because of the outliers, the median gives a better representation of the center of the distribution.

When using the original model (EOQ for regular and bi-weekly demand for irregular demand), the number of yearly orders increases by 12.3% to 9936. Surprisingly, average inventory levels also increase, by approximately 8%. When considering Figure 5.4, this result can be explained. Smooth and erratic demand is on average significantly higher than intermittent and lumpy demand (see Table 5.6). Using EOQ for smooth and erratic products results in on average higher order quantities, as can be seen in Figure 5.4. A small relative increase in Q still results in a high absolute increase in inventory levels. On the other hand, intermittent and lumpy products decrease on average in order quantity. Given the low annual demand of these products, the absolute decrease in inventory level is not very big. However, the number of orders that is placed still increases significantly for these products.

Even though both number of orders and average inventory level increases, the average inventory costs decrease by 27%. The cost-oriented focus of EOQ becomes clear in this result. The most expensive items are ordered more often whilst the cheaper items are ordered less, leading to a decrease in holding costs of 34%.

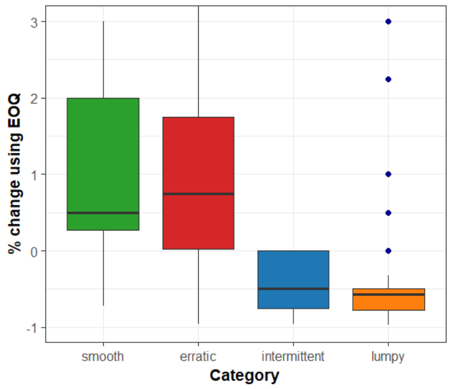

Figure 5.4: *Boxplots showing the relative change in Q when using EOQ for all demand categories.*

Table 5.5: *Summary statistics for the change in Q compared to the current situation, when using the original formulation of Q.*

Category	Min	Q1	Median	Mean	Q3	Max
smooth	-0.72	0.27	0.50	0.91	2.00	3.00
erratic	-0.96	0.02	0.75	1.44	1.74	25.50
intermittent	-0.96	-0.75	-0.50	-0.47	0.00	0.00
lumpy	-0.97	-0.78	-0.57	-0.25	-0.49	25.20

Remark. Note that average inventory level does not directly translate to space usage. Average inventory level refers to the total number of units on stock. However, some products are bigger than others. If the average inventory level increases, this could mean that products with high space usage decrease in inventory level, whereas the inventory level of products with low space usage increases, or vice-versa. It is not clear whether this is the case in this research, since data on product sizes is unavailable.

Table 5.6: *Average price per unit and yearly demand, split by category.*

	Avg price per unit	Avg yearly demand
smooth	1.19	3100
erratic	3.94	3849
intermittent	5.83	45
lumpy	4.73	174

5.2.2 Alternative order quantity

As illustrated in Section 5.2.1, the order quantity for products with smooth and erratic demand increases on average when using EOQ. This is undesirable for decreasing space usage. In addition, the variance in change of Q compared to the current situation is high, as can be seen in the wide inner quartile ranges in Figure 5.4. This could indicate that using EOQ yields inconsistent results for the WKZ. Therefore, we also use an alternative way to determine Q where average demand is the leading variable. For smooth demand, we use average weekly demand and for erratic demand, we use bi-weekly demand. These values are chosen so that the median decrease in Q is approximately equal for both demand types. Figure 5.8 and Table 5.7 show the results. Using this method of deriving an order quantity, we obtain a more consistent result with less variance in changes with respect to the current situation. The annual number of orders increases with 18.7% to 10501 compared to the current 8846 annual orders. The average inventory level decreases with 9.4%. When looking at total inventory costs, we observe a decrease of 9.8% compared to the current situation.

Summarizing, the alternative order quantity results in lower inventory levels, but higher order frequencies and inventory costs compared to the original order quantity. Both methods result in lower inventory costs compared to the current situation.

Table 5.7: *Summary statistics for the change in Q compared to the current situation, when using the alternative formulation of Q.*

Category	Min	Q1	Median	Mean	Q3	Max
smooth	-0.75	-0.50	-0.19	-0.23	-0.04	0.40
erratic	-0.88	-0.42	-0.17	0.02	0.32	2.91
intermittent	-0.96	-0.75	-0.50	-0.47	0.00	0.00
lumpy	-0.97	-0.78	-0.57	-0.25	-0.49	25.20

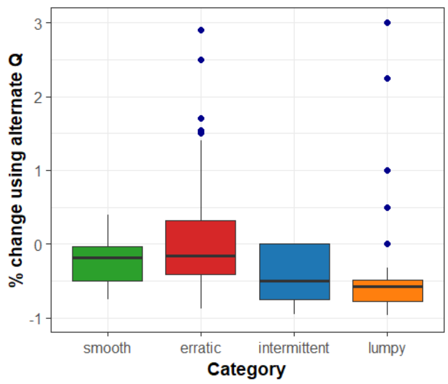

Figure 5.8: *Boxplots showing the relative change in Q when using the alternative Q for all demand categories.*

5.3 Computing reorder point

This section discusses the results of computing an appropriate reorder point s through the model developed in Chapter 4. Since the derivation of s depends on Q, we use the two variants of Q discussed in Section 5.2. Section 5.3.1 discusses the general results that we observe when running the model for both the original and alternative Q. We pay additional attention to some interesting results of irregular demand items in Section 5.3.2. Again, the results are evaluated on average inventory levels, number of orders and inventory costs.

5.3.1 Results for all categories

One can observe the results of the percentual change in s when using the original and alternative Q in Figures 5.9 and 5.10, respectively. In total, 34 products can be classified as outliers, of which the biggest are left out in Figures 5.9 and 5.10.

When using the original Q, we observe an average increase in s of approximately 15.9%. For the alternative Q, this is equal to 9.5%. However, the average inventory decreases in both cases. The reason for this is likely similar to the effect described in Section 5.2.1. The products that have

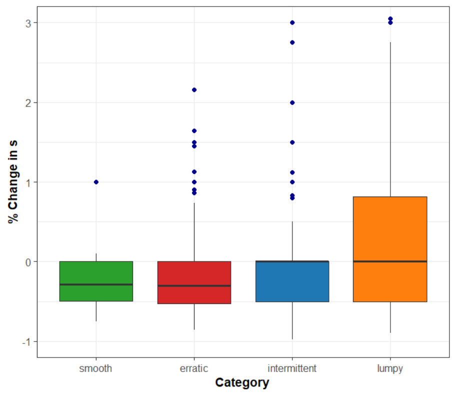

Figure 5.9: *Boxplots showing the relative change in s when using the original Q for all demand categories.*

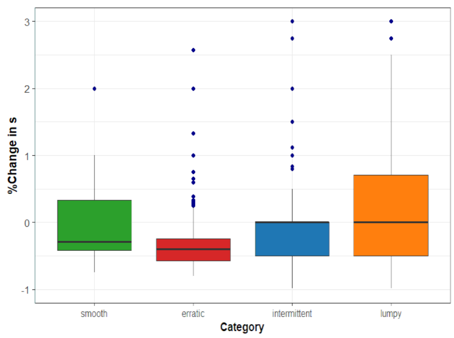

Figure 5.10: *Boxplots showing the relative change in s when using the alternative Q for all demand categories.*

high annual demand have a big impact on average inventory, even if s decreases with a small percentage. This offsets the products with low annual demand for which s on increases by a relatively big amount. When using the alternative method of deriving Q in addition to the new reorder point s computed with this Q, we see a relatively small increase in annual number of orders (18.7%), but average inventory levels decrease significantly (30%). Table 5.11 shows the changes in inventory levels in more detail.

In Figure 5.10 we observe that most products with erratic demand are concentrated around a decrease in reorder point of 40%, apart from some outliers. This indicates that either the WKZ currently consistently overestimates the necessary safety stock for erratic demand items, or the benefit of aggregating demand for these products is the largest.

Table 5.11: *Changes in average inventory level when using the new values for Q and s*

	Avg inventory (Q)	%change	Avg inventory (Q and s)	%change
Current	41,253	-	41,253	-
Original Q	44,524	0.079	40,916	-0.0082
Alternative Q	37,369	-0.094	29,048	-0.30

When looking at the change in average inventory costs, we observe a decrease in inventory costs of 22.8% when the model for s is used in combination with the original order quantity, and 34.5% when used in combination with the alternative order quantity. This means that using the model for s in combination with the original order quantity leads to additional costs compared to when only using the original order quantity. This is likely since the changes in order quantities would lead to situations where we do not attain fill rate $P_2 = 0.99$. Reorder point s has to increase in order to compensate this, increasing inventory costs.

5.3.2 Results for irregular demand

As Table 5.12 shows, the reorder point for irregular demand items increases on average. In this section, we elaborate on possible reasons for this. In addition, we use Algorithm 4 to derive an empirical fill rate for the products under the new model. By doing this, we can check whether the model would perform in a satisfactory way according to the same metric that we use for regular demand.

Table 5.12: *Summary statistics for the relative change in s for irregular demand items*

Category	Min	Q1	Mean	Median	Q3	Max
intermittent	-0.98	-0.50	0.16	0.00	0.00	7.00
lumpy	-0.90	-0.50	0.60	0.00	0.82	18.00

Increased reorder point

The mean change in reorder point for intermittent and lumpy products is an increase 16% and 60%, respectively. In Section 2.4 we identified a possible difficulty regarding the quality of data, which was caused by having fixed (high) order quantities. In this section, we elaborated on two products that exemplified this problem. For consistency, we take the same products in this section to illustrate the effect on the reorder point.

Product 3000129 is categorized as intermittent and used on one department. The product currently has an order quantity of 80. However, the average demand per week is approximately equal to 12, meaning that around 8 orders are placed every year. Translating the order data directly to demand means that we expect a high variance in demand. Therefore, we have to take a high safety inventory to be prepared for the high demand peaks. Using the ZIP model identified in Chapter 4, we obtain $s = 75$. This is an increase of 275% compared to the current reorder point $s = 20$. Given the data and the assumption that orders can be directly translated to demand, the new reorder point makes sense. However, given that a reorder point of 20 has caused no problems in the past, it is clear that making the assumption gives an overestimation of demand.

For product 3000182, we observe a similar pattern. However, occasionally an order of 2 times the batch size is placed. The reason for this is not clear, since there are also many weeks in which no order is placed. This puts the product in the lumpy category. After running the model, we obtain a reorder point of 10 - an increase of 233% compared to the current situation.

Remark. Both products have an empirical fill rate of 1, both when using the current and newly determined value for Q. This indicates that the reorder point is appropriate to absorb the variability in demand.

Fill rate

In this section, we discuss the empirical fill rate that we compute for irregular products using a slightly adjusted version of Algorithm 4.

The new order quantity for intermittent and lumpy products is taken to be the average demand over two weeks. For some products, demand is so irregular that this means that a very low order quantity is obtained. For example, product 3000150 is ordered in quantities of 20, but with very low frequency. In 2023, a batch of 20 was ordered twice, leading to the new $Q = 2$. Using the ZIP distribution, a reorder point of 1 is obtained for this product. This is a surprising result, given that the two periods of positive demand were much higher. This becomes visible in the empirical fill rate: -3.05. Once demand occurs, the inventory suddenly decreases to $s + \frac{Q}{2} - 20 = 1 + 1 - 20 = -18$. Given that orders are placed in a quantity of 2 with the new model, it takes 9 periods to get back to a positive inventory level. This significantly increases ESC, leading to a negative empirical fill rate.

The average empirical fill rate when using the new value of Q is 0.29 for intermittent demand and 0.91 for lumpy demand, lower than the desired $P_2 = 0.99$. The fill rate for lumpy products is higher since we add the empirical ESC to the reorder point. We suspect that the low fill rate is due to the effect, described above, of having low order quantities. To still get some insight into the capacity of the reorder point to attain the desired fill rate, we also run Algorithm 4 using the order quantity that is currently used in the WKZ. This yields an empirical fill rate of 0.97 for intermittent and 0.98 for lumpy demand, which is quite close to the desired 0.99. This indicates that the model for determining reorder points for irregular demand items works well for reasonably high order quantities.

Because of the unstable results for irregular demand items at the WKZ, the hospital could consider ignoring the model's results for these items. If the current s and Q are used for irregular and the model output for regular demand, a reduction in average inventory level of 25% is achieved. In addition, the total number of annual orders decreases by 0.2%. We obtain a decrease in inventory costs of 10.5%.

5.4 Conclusion

In this chapter, we applied the model from Chapter 4 to the order data of consumable products in 2023. We observe that approximately 61% of products has irregular demand, meaning that they are ordered infrequently. Out of the regular demand products, 94% has erratic demand. For irregular demand items, the difference in ADI for intermittent demand is significantly larger than that of lumpy demand.

When only applying the model for calculating order quantity Q, we observe big changes compared to the current situation for regular demand. This indicates that the EOQ is likely not directly applicable to the products of the WKZ. On average, this leads to an increase in average inventory levels of 8% and an increase in number of orders of 12.3%. However, inventory costs decrease by 27%.

Since using EOQ leads to some inconsistent results, we proposed an alternative method of deriving Q for regular demand. There, we also use average demand as determining factor. This gives less radical changes when compared to the current situation. On average, we obtain a decrease in inventory levels of 9.4% and an increase in number of orders of 18.7%. However, inventory decrease less than when using EOQ: we obtain a reduction of 9.8%.

Applying the model for determining reorder point s results in a decrease in inventory levels of 1% when used in combination with the original order quantity, and a decrease of 30% when used in combination with the alternative order quantity. We also achieve a decrease in inventory costs of 22.8% and 34.5%, respectively.

Because of the effect described in Section 2.4 on high order quantities, we obtain some big increases in s for irregular demand. When using the current order quantities, we do obtain empirical fill rates of 0.97 and 0.98 for intermittent and lumpy demand, respectively. Not using the model for irregular demand and simply taking the current parameters leads to a decrease in average inventory levels of 25%, when used in combination with the alternative Q and new s for regular demand. This also leads to a decrease in inventory costs of 10.5%.

Chapter 6

Recommendations and Future Research

This chapter serves as a reflection on the findings of this thesis. First, it discusses the practical recommendations that we give to the WKZ based on the findings from preceding chapters (Section 6.1). The focus is put on implementability of the model that was developed in Chapter 4, in addition to further data gathering that would improve the effectiveness of the model. Moreover, this chapter discusses possible areas of future research that would extend the usability and effectivity of the model that we developed in Chapter 4 (Section 6.2). Finally, Section 6.3 discusses the scientific contributions of this thesis to the field of inventory management.

6.1 Recommendations

First, we give practical recommendations to the WKZ that can be given from this thesis. Some are already treated or introduced in preceding chapters, but still elaborated upon in this section.

6.1.1 Implementation of model

There are a few considerations and recommendations with regards to the implementation and usability of the model that we developed in this thesis. This section elaborates on them.

To begin with, we recommend the WKZ to integrate the model with the decision-making process for the MFA. In Chapter 5 we found that a reduction of at least 25% is possible with regards to average inventory level. Given that there is a space shortage of 26% (Section 1.2), this thesis solves the action problem sufficiently well assuming a direct translation between inventory levels and space usage. In addition, Section 1.2 identified the high shortage rate of certain products. Since the model works with a fixed fill rate for each product, these shortages should also be reduced.

We recommend that the WKZ uses the model to determine the reorder point and order quantities for smooth and erratic products. For these products, the data was of high quality and relatively directly translatable to actual demand. However, for intermittent and lumpy items, this translation could not be done with confidence. This led to higher reorder points. Given the high uncertainty with these products, we recommend the WKZ to not use the new order quantities and reorder points for irregular demand items. Even then, the average inventory levels can be reduced by 25% without increasing the number of annual orders and thereby workload. Since products with irregular demand still benefit from a reduction in variance through aggregated demand, the WKZ could consider reducing safety stock for these products by a fixed factor. The exact factor is something that should be determined through further research.

Remark. The recommendation that we give keeps a lot of safe assumptions in mind, which significantly reduces the probability that we recommend inventory levels that would results in shortages. However, given the vital importance of having enough inventory on stock in the hospital it is still important to check the new parameters for each product on a case-by-case basis. Given the statements that were already made regarding data quality, there is still a possibility that some errors occurred whilst running the model.

6.1.2 Data gathering

As explained in Section 2.4, there are some inconveniences with the available data at the WKZ. To begin with, the data refers to orders instead of demand. We recommend the WKZ to gather data on actual demand instead. The current data does not always reflect demand patterns accurately, as was observed for irregular demand items. We noted that this led to an increase in recommended reorder points for these products. By gathering data on actual demand, more accurate recommendations for reorder points can be given. Currently, the inventory levels of most warehouses are reviewed on a daily basis. We recommend the employees to track these levels using the scanning system elaborated on in Section 2.1. This system should be extended by allowing the user to scan every product and enter the inventory level in addition to placing an order when this inventory level is too low. This will lead to an increase in workload, since currently the employees must only scan the products that need to be ordered. If this increased workload turns out to be infeasible, gathering the data once or twice per week is likely also sufficient. Note that the data on inventory levels still needs to be translated to actual demand by subtracting the received orders.

If data on actual demand becomes available, the hospital should also use our proposed model for managing irregular demand. Due to the effects elaborated upon in Section 2.4, making the assumption that placed orders translate directly to demand led to unstable results for irregular demand. However, once the WKZ gathers demand data, this assumption does not need to be made anymore and the results will be applicable.

In addition, we recommend the WKZ to improve the storage of data on products. As explained in Section 2.4.2, data is occasionally stored inconsistently. We refer to this section for detailed recommendations. Making the adjustments mentioned there should decrease the likelihood of making errors when running the model. It also allows for a smoother analysis process in future projects, where less data cleaning will need to be done.

Finally, we recommend the WKZ to re-run the model once the MFA has been implemented for some time. It is possible that gathering data on the demand of all products turns out to be unfeasible. In this case, the lower order quantities that we recommend, in addition to the fact that all warehouses are integrated, should make it so that products are ordered more often. This will make it so a more detailed and representative overview of the demand of products is available when only tracking historical orders. In addition, given that products will be ordered more often, more demand will be identified as regular, meaning that the model can be applied in more cases.

6.2 Further research

There are several interesting objects of future research that are related to this thesis. Performing additional research on these areas would be beneficial for deriving a more effective and useful model for the WKZ. We identify and motivate these objects of future research below.

6.2.1 Investigate and experiment with parameters

To begin with, this thesis made assumptions on parameters such as desired fill rate P_2, holding cost rate h or ordering costs K. Although the estimation for these parameters was well-motivated, it would be beneficial to perform more in-depth research. We describe our vision on how to do this for the most relevant parameters below:

Service level

In this thesis, we worked with cycle service level (P_1) and fill rate (P_2). The leading indicator was P_2, given that it is the most interpretable for stakeholders. In conversation with stakeholders, we determined that 0.99 is an acceptable and desirable fill rate. Using trial-and-error, we found that a cycle service level of 0.975 worked well to achieve this (empirical) fill rate for irregular demand. One could experiment over several fill rates as a sensitivity analysis to determine the effect on inventory levels. In addition, several cycle service levels for irregular demand items could be tested, to see which works most desirably and achieves the empirical fill rate closest to the target.

EOQ parameters

In this thesis, we used a holding cost rate that is conventional according to literature. Given that space usage is the leading performance indicator for the WKZ, further research could alternatively use data on the size of products instead of the price. The product size would then be used as the leading variable in determining the holding cost rate. However, currently, product size is not recorded in the database of the WKZ. In this thesis, we used an estimation of order costs by separating the costs into several parts that could be calculated. Alternatively, one could consider obtaining the exact money spent on orders and dividing it by the number of orders. Further research could also look into the effects of adjusting the EOQ parameters, to see how big the effect on order quantities is.

Ordering frequency

In the original method for deriving Q for irregular demand and in the alternative method for all types of demand, we used average demand as a way to determine Q. For smooth demand we used average weekly, and for the others bi-weekly demand. These parameters result in desirable inventory levels for the new situation in the WKZ. However, more extensive experimentation could be done to determine a good order frequency specific to all types of demand. If it turns out that the reduction in space usage is more than necessary in practice, the order quantities could be adjusted upwards, and vice-versa.

6.2.2 Working with order data

The most desirable situation is having the most accurate data possible, in this case having access to exact product demand per day. However, given that gathering data is an expensive and time-consuming undertaking, it is likely unfeasible to gather this data for many organizations. In order to not have to make the strict assumption that orders translate directly to demand, more research should be performed on how to appropriately work with order data. One could consider looking at the time between orders as a leading variable instead of orders per week, for example. To the best of our knowledge, no literature on how to handle this could be found, meaning it is likely that a knowledge gap exists here.

6.2.3 Extend ZIP model

Using a ZIP model for inventory optimization is an interesting and recent development in inventory management, first proposed in 2022 by Finco et al. (2022). In the paper, the authors showed the effectiveness of the model through a case study. Applying an adapted version of the model in this thesis also seemed to result in appropriate reorder points. However, given the recent introduction, some improvements and additions to the model are possible.

To begin with, the ZIP model has no theoretical way of including fill rate as a constraint. Currently, only cycle service level can be used to determine an appropriate reorder point. This thesis developed a way to compute an empirical fill rate through historical data. However, with the unavailability of a long time horizon of data, a theoretical way of determining the fill rate becomes necessary.

Furthermore, the model currently does not allow for the immediate distinction between intermittent and lumpy demand. A Poisson distribution has one parameter which defines both its mean and variance. This means that either the variance of intermittent demand could be overestimated, or the variance of lumpy demand underestimated. This thesis worked around the second case by introducing an additional safety factor for lumpy demand. More research could be put into finding other ways of making this distinction. Other distributions could be extended with the binomial distribution, allowing for zero-inflation. For intermittent demand, one could consider using a zero-inflated normal distribution. For lumpy demand, a zero-inflated Gamma distribution could be appropriate.

6.2.4 Include ZNLA products

This thesis focused on ZROH products, which are stored in the central warehouse of the UMC Utrecht. However, the hospital also works with ZNLA products which are ordered from the manufacturer once inventory falls below the reorder point. This means that lead time is no longer deterministic and brief, but stochastic and long. One should research how to include this additional variance in the models. This means that gathering data on lead times per supplier is also necessary.

6.2.5 Include undershoot of reorder point

In an (R, s, Q) policy, which the WKZ uses, inventory is reviewed every R periods. Between these periods, an undershoot of reorder point s can occur. Ways of including this undershoot exist for normally distributed demand, but could not be found for erratic, intermittent or lumpy demand. As such, this thesis used a longer lead time to deal with the review period. A more accurate way of determining the distribution of the undershoot needs to be researched for non-normally distributed demand.

6.3 Scientific contribution

This section discusses the scientific contributions of this thesis. The main contribution is that we integrated several specialized techniques for managing demand to the ADI-CV categories. In addition, we extended the ZIP model and provided additional evidence that the resulting reorder points are accurate.

To begin with, we developed a model that applies a different technique for managing demand to the ADI-CV demand categories. We have found no previous such attempt in literature. Finco et al. (2022) develop a framework for ADI-CV categorization, but apply the same technique to each category. They only use ADI-CV categorization to determine whether applying the framework is necessary. Babiloni et al. (2010) and Dunsmuir and Snyder (1989) mention specific techniques to manage intermittent demand, but do not apply it in a case study. This thesis fills the gap of applying a variety of techniques in a real-life context.

Furthermore, we extended the use of the ZIP model proposed by Finco et al. (2022). The model is currently only suitable for use when cycle service level (P_1) is taken as the measure for service level. In this thesis, we used fill rate (P_2) given its higher interpretability. Therefore, we extended the ZIP model with an empirical method to determine an approximate fill rate. This is done through using historical data. This method can also be extended to other measures of service level.

Finally, we provided additional insights into the use of a ZIP model for managing intermittent demand. The model was first used in this context in 2022, by Finco et al. (2022). Given the recent introduction, it is relevant to provide some real-life evidence for the functioning of the model. We showed that, when making the assumption that our data reflects demand, using the model results in accurate reorder points that adhere to our desired fill rates when verifying them with historical data. This increases the credibility of the model.

Chapter 7
Conclusion

This chapter concludes this thesis by going through the sub-questions introduced in Chapter 1. For these questions, we summarize the findings from the corresponding chapter. We conclude with giving an answer to the main research question of how the WKZ can best reduce its inventory levels.

Question 1: How is consumable inventory managed in the WKZ?

Chapter 2 tackles this question. For most departments, the hospital uses an (R, s, Q) policy for managing consumable inventory. There, the inventory levels are reviewed every R units of time, and an order of size Q is placed once the inventory level reaches s. Depending on whether a product is defined as 'ZROH' or 'ZNLA' it is ordered from the central warehouse or directly from the manufacturer, respectively. ZROH products have a lead time of approximately one day and are the focus of this thesis. In the new situation where all warehouses are moved to the multi-functional area (MFA), the same inventory policy that is currently applied will be used.

The WKZ currently categorizes products based on how often they are demanded, and in which quantity once demand occurs. This introduces a 2x2 matrix, which the hospital uses to decide where to store their products. For this thesis, we used ADI-CV classification, which separates products based on the average interval between two positive demand periods, and the coefficient of variation of positive demand.

Question 2: What techniques exist for computing the optimal parameters in an (R, s, Q) policy?

Chapter 3 answers this question through a literature review. For determining order quantity Q, we identified the concept of an economic order quantity (EOQ), which finds the optimal balance between holding costs and ordering costs. Most techniques that we found for determining the optimal reorder point s involve modelling demand through a statistical distribution and finding the point for which a desired fill rate is attained. One such distribution is the zero-inflated Poisson (ZIP) distribution, which can be used for modelling intermittent and lumpy (irregular) demand.

Question 3: Which techniques should be used for managing inventory in the new warehouse in the WKZ?

Chapter 4 develops the model which includes all necessary techniques. For determining the optimal Q for smooth and erratic demand, we use EOQ. Given that in Chapter 5, using EOQ gave some inconsistent results, we also used average demand as an alternative to derive Q. This strategy was used per default for irregular demand items, since the necessary assumptions for EOQ are not sufficiently met.

We used three different statistical distributions to compute the optimal reorder point s. For smooth and erratic demand, a normal and Gamma distribution were used, respectively. Since no closed-form solution of the inverse of these distributions exists, a bisection algorithm was developed to numerically find the solution for a desired fill rate. For irregular demand, we used a ZIP distribution. The ZIP distribution does not work for fill rates, so we used cycle service level as constraint on desired service level. Lumpy demand items were given an extra safety stock, given their higher variance in demand. To test the attained fill rate for irregular demand, we developed and implemented an algorithm that uses historical data to empirically obtain a verification.

Question 4: How does the chosen model perform?

In Chapter 5, we observed that when only applying the original model for determining Q, an increase in average inventory levels of approximately 8% occurs. Simultaneously, the number of annual orders increases by approximately 12%. When also applying the model for determining s, a reduction in average inventory level of 1% occurs. Applying EOQ to smooth and erratic demand items gives inconsistent results, with significant deviations from the current order quantity. As an alternative, we therefore also consider using average (bi-)weekly demand as order quantity for these items. Applying this model gives an average inventory level reduction of 30%, when used in combination with the model for s. The results for irregular products often show a significant increase in s, which is likely due to the insufficient data quality. Leaving these products out of the model to prevent inaccuracies results in an average inventory level decrease of 25%. Furthermore, shortages are now acceptable for all products, since we homogeneously chose a fill rate of 99% in conversation with stakeholders.

Main question: How can the WKZ best revise its inventory management policy in order to reduce space usage and shortage rates for the MFA?

The WKZ can best keep using its (R, s, Q) policy to manage its inventory. We proposed a model that, through historical data and statistical distributions, gives a recommendation on an appropriate order quantity Q and reorder point s. We recommend the WKZ to apply this model to the products that are classified as having smooth or erratic demand, and to leave the parameters as-is for the irregular demand items. By doing this, a reduction in average inventory level of 25% is possible without increasing the number of annual orders, but even slightly reducing them with 0.2%. As such, we solved the action problem of having too little space in the new warehouse. Additionally, fill rates will be consistent between products at 99%. To further reduce inventory levels and obtain a more accurate model, we proposed areas of further research in Chapter 6. In addition, the WKZ should gather data on actual demand instead of tracking orders. We recommend the WKZ to keep using the delivered model in the future, when new and accurate demand data becomes available.

Bibliography

4Building. Nieuwbouw IC Wilhelmina Kinderziekenhuis — 4Building — 4building.nl. https://4building.nl/portfolio/nieuwbouw-ic-wilhelmina-kinderziekenhuis/#, 2024. [Accessed 05-04-2024].

K. Atkinson. *An Introduction to Numerical Analysis, 2nd Ed.* Wiley India Pvt. Limited, 2008. ISBN 9788126518500. URL https://books.google.nl/books?id=lPV8Fv2XEosC.

E. Babiloni, M. Cardós, J. Albarracín, and M. Palmer. Demand categorisation, forecasting, and inventory control for intermittent demand items. *South African Journal of Industrial Engineering*, 21(2):115 – 130, 2010. URL https://www.scopus.com/inward/record.uri?eid=2-s2.0-78650790554&partnerID=40&md5=0e65c0ceda3756e46fe5d6dadd83d729. Cited by: 6.

S. Beckett, J. Jee, T. Ncube, S. Pompilus, Q. Washington, A. Singh, and N. Pal. Zero-inflated poisson (zip) distribution: parameter estimation and applications to model data from natural calamities. *Involve, a Journal of Mathematics*, 7:751–767, 10 2014. doi: 10.2140/involve.2014.7.751.

Bizagi. Process Modeling Software - Bizagi Modeler — Bizagi — bizagi.com. https://www.bizagi.com/en/platform/modeler, 2024. [Accessed 29-04-2024].

S. Chopra and P. Meindl. *Supply chain management.* Pearson, Upper Saddle River, NJ, 6 edition, Dec. 2014.

C. Dridi. A short note on the numerical approximation of the standard normal cumulative distribution and its inverse. Computational economics, University Library of Munich, Germany, 2003. URL https://EconPapers.repec.org/RePEc:wpa:wuwpco:0212001.

W. Dunsmuir and R. Snyder. Control of inventories with intermittent demand. *European Journal of Operational Research*, 40(1):16 – 21, 1989. doi: 10.1016/0377-2217(89)90267-1. URL https://www.scopus.com/inward/record.uri?eid=2-s2.0-0024961882&doi=10.1016%2f0377-2217%2889%2990267-1&partnerID=40&md5=312da4c067f617c64513321c1aad7046. Cited by: 50.

EGM Architecten. Geboortecentrum, Wilhelmina Kinderziekenhuis — egm.nl. https://egm.nl/architecten/projecten/-wilhelmina-kinderziekenhuis-geboortecentrum-en-intensive-care/557, 2024. [Accessed 05-04-2024].

S. Finco, D. Battini, G. Converso, and T. Murino. Applying the zero-inflated poisson regression in the inventory management of irregular demand items. *Journal of Industrial and Production Engineering*, 39(6):458–478, Mar. 2022. ISSN 2168-1023. doi: 10.1080/21681015.2022.2041741. URL http://dx.doi.org/10.1080/21681015.2022.2041741.

N. Hasachoo, P. Sirisawat, and T. Kaewket. Inventory replenishment policy for medicines with non-stationary stochastic demand: The case of a newly opened hospital in thailand. In *2019 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, pages 89–93, 2019. doi: 10.1109/IEEM44572.2019.8978689.

H. Heerkens and A. van Winden. *Systematisch managementproblemen oplossen.* Noordhoff Uitgevers, 2017. ISBN 9789001878825.

S. G. Johansen and R. M. Hill. The (r,q) control of a periodic-review inventory system with continuous demand and lost sales. *International Journal of Production Economics*, 68(3):279–286, 2000. ISSN 0925-5273. doi: https://doi.org/10.1016/S0925-5273(00)00051-7. URL https://www.sciencedirect.com/science/article/pii/S0925527300000517.

P. Kalaya, P. Termsuksawad, and T. Wasusri. Forecasting lumpy demand for planning inventory: The case of community hospitals in thailand. In *2019 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, pages 686–690, 2019. doi: 10.1109/IEEM44572.2019.8978541.

F. Lolli, A. Ishizaka, R. Gamberini, and B. Rimini. A multicriteria framework for inventory classification and control with application to intermittent demand. *Journal of Multi-Criteria Decision Analysis*, 24(5-6):275 – 285, 2017. doi: 10.1002/mcda.1620. URL https://www.scopus.com/inward/record.uri?eid=2-s2.0-85034777806&doi=10.1002%2fmcda.1620&partnerID=40&md5=af70d963821b66130b4e51a188d35acd. Cited by: 14; All Open Access, Green Open Access.

F. Lolli, E. Balugani, A. Ishizaka, R. Gamberini, B. Rimini, and A. Regattieri. Machine learning for multi-criteria inventory classification applied to intermittent demand. *Production Planning and Control*, 30(1):76 – 89, 2019. doi: 10.1080/09537287.2018.1525506. URL https://www.scopus.com/inward/record.uri?eid=2-s2.0-85055894672&doi=10.1080%2f09537287.2018.1525506&partnerID=40&md5=7173570ff94c3d3ffc4b876f1cd75375. Cited by: 58; All Open Access, Green Open Access.

B. Maddah and N. Noueihed. Eoq holds under stochastic demand, a technical note. *Applied Mathematical Modelling*, 45:205–208, 2017. ISSN 0307-904X. doi: https://doi.org/10.1016/j.apm.2016.12.026. URL https://www.sciencedirect.com/science/article/pii/S0307904X16306795.

K. Mak, Y. Wong, and G. Huang. Optimal inventory control of lumpy demand items using genetic algorithms. *Computers & Industrial Engineering*, 37(1):273–276, 1999. ISSN 0360-8352. doi: https://doi.org/10.1016/S0360-8352(99)00072-8. URL https://www.sciencedirect.com/science/article/pii/S0360835299000728. Proceedings of the 24th international conference on computers and industrial engineering.

D. Shenoy and R. Rosas. *Problems & solutions in inventory management*. Springer International Publishing, Basel, Switzerland, 1 edition, Oct. 2017.

N. Showkat and K. Ali. Conceptual framework of inventory management: A useful guide towards the major aspects of the subject. *International Journal of Advanced Research*, 8(9):395–411, Sept. 2020. ISSN 2320-5407. doi: 10.21474/ijar01/11671. URL http://dx.doi.org/10.21474/IJAR01/11671.

E. A. Silver, D. F. Pyke, and D. J. Thomas. *Inventory and production management in supply chains*. CRC Press, London, England, 4 edition, Sept. 2021.

N. Slack and A. Brandon-Jones. *Operations Management*. Pearson Education, London, England, 9 edition, Apr. 2019.

L. Strijbosch, R. M. J. Heuts, and E. H. M. van der Schoot. A combined forecast—inventory control procedure for spare parts. *Journal of the Operational Research Society*, 51(10):1184–1192, 1999. doi: 10.1057/palgrave.jors.2601013.

F. Wang and L. X. Xiao Xia. Evaluation and selection of periodic inventory review policy for irregular demand: A case study. In *2015 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, pages 275–279, 2015. doi: 10.1109/IEEM.2015.7385651.

Wilhelmina Kinderziekenhuis. Strategie - Het WKZ — hetwkz.nl. https://www.hetwkz.nl/nl/strategie, 2024. [Accessed 04-04-2024].

W. L. Winston and J. B. Goldberg. *Operations research: Applications and algorithms*. Brooks/-Cole, Florence, KY, 2004.

Q. Xu, N. Wang, and H. Shi. Review of croston's method for intermittent demand forecasting. In *2012 9th International Conference on Fuzzy Systems and Knowledge Discovery*, pages 1456–1460, 2012. doi: 10.1109/FSKD.2012.6234258.

Appendices

Appendix A
General concepts for determining order quantity and reorder point

In this chapter, we discuss general concepts that give a starting point for most techniques for computing the optimal order quantity Q and reorder point s. Section A.1 explains the concept of an economic order quantity, which is used as a quick way to compute an appropriate value for order quantity Q. Sections A.2 and A.3 discuss concepts necessary for computing reorder point s. Specifically, Section A.2 discusses the necessary statistical distributions that can be used to model lead time demand. Section A.3 gives two ways of measuring the service level.

A.1 Economic order quantity

An important concept for finding the optimal order quantity Q is that of an economic order quantity (EOQ). Here, the aim is to find a balance between holding costs and ordering costs to derive the optimal value for Q. Even though the derivation of the EOQ requires relatively strict assumptions, it can still be used reasonably well in situations where these assumptions are not fully met. We elaborate on these assumptions, and briefly review the derivation of the formula.

A.1.1 Assumptions

Silver et al. (2021) define the following assumptions that are necessary for the derivation of the EOQ. The most relevant ones are as follows:

1. Demand D is deterministic and constant over time.
2. Order quantity Q can be any positive real number (i.e. $Q \in \mathbb{R}^+$).
3. There are no quantity discounts for ordering larger batch sizes.
4. Lead time L is deterministic.
5. An order is delivered in its entirety at once.

Relaxations of these assumptions exist. For example, Maddah and Noueihed (2017) have shown that the EOQ still holds under stochastic demand. Other assumptions that can be relaxed are discussed below.

A.1.2 Derivation

We adapt the general structure of the derivation from Silver et al. (2021). We use the notation introduced in Section 3.1.

EOQ is a cost-oriented notion, meaning the goal is to minimize costs. The is influenced by two cost components: holding costs and ordering costs. The total costs TC can be expressed as follows:

$$TC(Q) = \frac{KD}{Q} + \frac{Qh}{2}. \tag{A.1}$$

$\frac{KD}{Q}$ are ordering costs, since $\frac{D}{Q}$ refers to the total number of orders per period and K the costs of placing one order. The second part refers to average holding costs. Since inventory decreases linearly, the average inventory is $\frac{Q}{2}$, and the holding costs per item per period is h.

The total cost function is convex, meaning that the optimum for Q can be found by differentiating with respect to Q and equating to zero. This yields the following result:

$$\frac{\delta TC(Q)}{\delta Q} = \frac{h}{2} - \frac{KD}{Q^2} = 0 \implies \text{EOQ} = \sqrt{\frac{2KD}{h}}. \tag{A.2}$$

Silver et al. (2021) mention that *"total costs are rather insensitive to deviations from the optimal lot size"*. This means that even if we cannot exactly order Q items due to for example lot sizes, the impact is relatively small. It also implies that working with parameters that are not completely accurate is less impactful.

A.1.3 Estimating parameters

For computing the EOQ, the parameters K, D and h are needed. With demand data, D can be estimated relatively accurately. However, the other variables are more difficult to compute.

Holding costs

Shenoy and Rosas (2017) argue that holding costs are made up of cost of capital, cost of storage, cost of inventory risks and service costs. Cost of capital is usually the largest component, and refers to the investment opportunity that is lost when spending the money to keep items on stock. Cost of storage refers to the costs that are incurred because of the space that needs to be provided to stock the items. Inventory risk refers to the chance of an item breaking or becoming obsolete. Finally, service costs are insurances, taxes and personnel costs for handling the products. Winston and Goldberg (2004) mention that annual holding costs of 20% - 40% of the unit costs are often used in practice. This percentage is referred to as the "holding cost rate".

Ordering costs

Ordering costs can also be broken down into several components. Shenoy and Rosas (2017) identify them as administration costs, transportation costs and inspection costs. Administration costs are incurred because of the "*time and effort expended in preparing a purchase order*" (Shenoy and Rosas, 2017). With (semi-)automatic ordering systems, these costs reduce significantly. Transportation costs include fuel, but also (un)loading costs. Finally, inspection costs are incurred if items need to be checked for quality or correctness.

A.2 Statistical distributions

For this thesis, several statistical distributions are required to model the varying types of demand. We assume the reader to be familiar with basic concepts from statistics and probability, such as random variables, mean, variance, quantile functions and probability density functions (PDFs). We denote sample mean by $\hat{\mu}$ and sample standard deviation by $\hat{\sigma}$.

A.2.1 Normal distribution

The normal distribution is a frequently used distribution to model a wide variety of processes. Its importance is largely due to the central limit theorem, which states that under weak conditions, the mean of a sample converges in distribution to a normal distribution. Because of this, it can be observed in many natural events. In addition, it has convenient properties that make the distribution easy to work with. The PDF of a normal distribution with mean μ and standard deviation σ is as follows:

$$\phi(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}. \tag{A.3}$$

No closed form of the cumulative distribution function (CDF) of the normal distribution exists, but it can be closely approximated by using the Gauss error function (Dridi, 2003). Similarly, the quantile function needs to be approximated.

Notation If a random variable (r.v.) X follows a normal distribution with mean μ and standard deviation σ, we write $X \sim \mathcal{N}(\mu, \sigma)$. The CDF is denoted by Φ, with $\Phi(x) = \mathbb{P}(X \leq x)$.

Properties If $X \sim \mathcal{N}(\mu_X, \sigma_X)$ and $Y \sim \mathcal{N}(\mu_Y, \sigma_Y)$ are independent, and $c \in \mathbb{R}$, we have that:

- **Addition.** $X + Y \sim \mathcal{N}(\mu_X + \mu_Y, \sqrt{\sigma_X^2 + \sigma_Y^2})$
- **Scalar multiplication** $cX \sim \mathcal{N}(c\mu, c\sigma)$

These properties will be useful for computing the safety stock levels. For example, if demand during lead time and review period are both normally distributed, we also know the distribution of the sum of the two variables. Because of the second property, a normal distribution is infinitely divisible. This is required for some techniques that we will use.

A.2.2 Gamma distribution

The Gamma distribution is a broad family of distributions. Silver et al. (2021) argue its usage for modelling demand when $CV > 0.49$.

Notation If r.v. X follows a Gamma distribution with shape α and rate β, we write $X \sim \Gamma(\alpha, \beta)$. The density function is as follows:

$$f(x) = \frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)} \tag{A.4}$$

where $\Gamma(k)$ denotes the Gamma function:

$$\Gamma(k) = \int_0^\infty t^{k-1} e^{-t} dt. \tag{A.5}$$

Properties For the remainder of this thesis, several properties of the Gamma distribution are useful and can be seen below. Let $X \sim \Gamma(\alpha, \beta)$, then:

- **Moment estimators.** $\hat{\alpha} = \hat{\mu}^2/\hat{\sigma}^2$ and $\hat{\beta} = \hat{\mu}/\hat{\sigma}^2$
- **Scalar multiplication.** $cX \sim \Gamma(\alpha, \frac{\beta}{c})$

A.2.3 Poisson distribution

A Poisson distribution is a discrete distribution, meaning that the random variable can only take on integer values. Winston and Goldberg (2004) mention the use of Poisson distributions for slow-moving products.

Notation If r.v. X follows a Poisson distribution with rate λ, we write $X \sim \text{POIS}(\lambda)$. The density function is as follows:

$$f(x) = \mathbb{P}(X = k) = \frac{\lambda^k e^{-\lambda}}{k!}. \tag{A.6}$$

The Poisson distribution has the property that the mean is equal to the variance. That is, $\mathbb{E}(X) = var(X) = \lambda$.

A.2.4 Binomial distribution

Finally, we discuss the binomial distribution, which is a generalization of the Bernoulli distribution. A Bernoulli distribution has one parameter $p \in [0, 1]$ that can be interpreted as a "success probability". In this thesis, we speak of a success when the Bernoulli process results in a 0 and a failure if it results in 1. As such, a Bernoulli distribution outputs 0 with probability p and 1 with probability $1 - p$. This is different from most definitions, where p is the probability of having a 1 as realization. However, for zero-inflated distributions that we use later, this notation is convenient. Zero-inflated distributions are a combination of multiple distributions where one distribution generates zeros. As such, these distributions often have a lot of probability mass at zero. One realization of a Bernoulli distribution is also called a "Bernoulli trial".

The binomial distribution generalizes this distribution with an additional parameter n, referring to the number of Bernoulli trials that we perform. The random variable that follows a binomial distribution therefore outputs the number of successes in n Bernoulli trials with success probability p.

Notation If r.v. X follows a binomial distribution with parameters $n \in \mathbb{N}$ and $p \in [0, 1]$, we write $X \sim B(n, p)$. The PDF of a binomial distribution is as follows:

$$f(x) = \mathbb{P}(X = x) = \binom{n}{x} p^x (1 - p)^{n-x}. \tag{A.7}$$

A.3 Service level

We elaborate on what we use as criteria for determining the quality of an inventory policy. Given the importance of having products on stock in a hospital, the leading performance measure in this thesis should be related to service level or product availability. The minimum requirement on service level should be discussed with stakeholders, to which the parameters s and Q can be adjusted. We discuss two ways of measuring service level below.

A.3.1 Fill rate

Fill rate (also referred to as P_2) is an appropriate measure to use for evaluating product availability. It refers to the fraction of demand which is supplied from stock, so without incurring a backorder. Chopra and Meindl (2014) mention that the fill rate can be computed as:

$$P_2 = 1 - \frac{\text{Expected shortage per replenishment cycle}}{\text{Expected demand during replenishment cycle}} = 1 - \frac{ESC}{Q}. \tag{A.8}$$

Note that this definition depends on Q, so Q needs to be determined first. The determination of ESC depends on the type of distribution with which demand is modeled. This thesis mostly uses fill rate for items with regular demand, since ESC can be determined analytically there.

A.3.2 Cycle service level

Another way to measure service level is the so-called cycle service level (P_1). Chopra and Meindl (2014) define P_1 as "*the fraction of replenishment cycles that end with all the customer demand being met*". Note that the managerial implications of this measure are therefore dependent on the number of replenishment cycles and thereby Q. If Q is very small, P_1 can be high, but the total number of stock-outs still undesirable. Mathematically, we can define P_1 as the probability that lead time demand is lower than the reorder point:

$$P_1 = \mathbb{P}(D_L \leq s). \tag{A.9}$$

If the distribution of lead time demand is known, this probability can directly be computed or approximated using the cumulative distribution function.

Appendix B
Model code

B.1 Cleaning data on batch size

```python
def main():
    for index, row in df.iterrows():
        art = row['Artikelomschrijving']
        if '(' in art and ')' in art:
            value = art[art.find('(') + 1:-1]
            try:
                int(value)
            except:
                df.loc[index, "Aantal ST in eenheid"] = value
            else:
                df.loc[index, "Aantal ST in eenheid"] = int(value)

            if 'X' in value and not 'B' in value:
                new_value = int(value[value.find('X') +1:])
                df.loc[index, 'Aantal ST in eenheid'] = new_value

            if 'x' in value and not 'B' in value:
                new_value = int(value[value.find('x') + 1:])
                df.loc[index, 'Aantal ST in eenheid'] = new_value

            if 'ZAK 8' in value:
                df.loc[index, 'Aantal ST in eenheid'] = 8

            elif 'VER' in value or 'BOTN' in value or 'FT' in value or 'HO' in value:
                df.loc[index, 'Aantal ST in eenheid'] = 1

            if 'RELOADABLE' in value:
                df.loc[index, 'Aantal ST in eenheid'] = 3

        elif '(' in art:
            value = int(art[art.find('(') + 1:])
            df.loc[index, 'Aantal ST in eenheid'] = value

        else:
            df.loc[index, 'Aantal ST in eenheid'] = 1

    df.to_excel('unit-calc.xlsx', sheet_name='Assortiment')
```

B.2 Performing preliminary analysis

Remark. This model can be run on the raw order data provided by the ERP tool of UMC Utrecht. The aim is to gather the product characteristics together with the order data, and compute some basic statistics such as mean and variance of demand, ADI and CV. The correct price of each batch size of a product is also computed.

```python
def main():
    # Import usage data
    df = pd.read_excel('data/data.xlsx')

    product_df = pd.read_excel('data/Analyse producten.xlsx', sheet_name='Product')

    usage_df = df.copy()
    for x in ['Match?', 'Price per batch', 'Aantal ST in eenheid', 'Logical divide',
    ↪ 'ADI', 'CV', 'Category', 'Original Q', 'Original s',]:
        usage_df[x] = None

    final = create_final_df(usage_df, product_df)

def create_final_df(usage_df, product_df):

    for index, product in usage_df.iterrows():
        matching_index = np.where(product_df['Artikel'] == product['Artikel'])[0]
        try:
            matching_index = matching_index[0]

        except:
            product['Match?'] = False
            continue

        else:
            product['Match?'] = True
            product['Aantal ST in eenheid'] = product_df.at[matching_index, 'Adj
            ↪ Aantal ST in eenheid']
            product['Original Q'] = product_df.at[matching_index, 'Original Q']
            product['Original s'] = product_df.at[matching_index, 'Original s']

            if all(product['01.2023':'52.2023'].mod(product['Aantal ST in eenheid'])
            ↪ == 0):
                product['Logical divide'] = True
                product['01.2023':'52.2023'] = product['01.2023':'52.2023'] /
                ↪ product['Aantal ST in eenheid']
            else:
                product['Logical divide'] = False

            if product["Artikelsoort"] == 'ZROH':
                product['Price per batch'] = product_df.at[matching_index, 'Prijs per
                ↪ eenheid']
                D_y = product['Mean'] * 1 / 7 * 365
                correct_price = compute_price(product)
                product['Price per batch'] = correct_price

            product['Mean'] = np.mean(product['01.2023':'52.2023']) * lead_time
            product['Standard deviation'] = np.std(product['01.2023':'52.2023']) *
            ↪ np.sqrt(lead_time)

            demand = product['01.2023':'52.2023'].copy()

            adi = compute_adi(demand)
            cv = compute_cv(demand)
            product['ADI'] = adi
            product['CV'] = cv
            product['Category'] = derive_category(adi, cv)

            usage_df.loc[index] = product

```

```
    usage_df.to_excel('data/Adjusted demand per unit new.xlsx')
    usage_df.drop(columns=usage_df.loc[:, '01.2023':'52.2023'], inplace=True)
    return usage_df

def compute_eoq(D, S, h):
    if h != 0:
        eoq = np.sqrt(2*D*S/h)
        return 1 if round(eoq) == 0 else round(eoq)
    else:
        return 10
def compute_adi(demand):
    n = sum(demand > 0)
    p = sum(demand == 0)
    return p/n

def compute_cv(demand):
    d = np.mean(demand[demand != 0]) * lead_time
    std = np.std(demand[demand != 0]) * np.sqrt(lead_time)
    return (std/d)

def derive_category(adi, cv):
    if adi < 1.32:
        if cv < 0.49:
            return 'smooth'
        else:
            return 'erratic'
    else:
        if cv < 0.49:
            return 'intermittent'
        else:
            return 'lumpy'

def compute_price(product):
    price = product['Price per batch']
    amount = product['Aantal ST in eenheid']

    if price < 2:
        return price * amount
    else:
        return price

if __name__ == '__main__':
    main()
```

B.3 General model

Remark. This code serves as the main structure of the model. We loop over each product in the data, and compute all relevant statistics using the correct technique based on the ADI-CV categories.

```
import numpy as np
import pandas as pd
import scipy.stats as sct
import matplotlib.pyplot as plt
import math
import time

#constant parameters
```

```
lead_time = 2/7
target_service_level = 0.975
target_fill_rate = 0.99
holding_cost_rate = 0.4
order_costs = 2.5
safety_factor = sct.norm.ppf(target_service_level)
Q_days = {
    'erratic': 14,
    'intermittent': 14,
    'lumpy': 14,
    'general': 14,
    'smooth': 7
}

def main():
    s = time.time()
    df = pd.read_excel('data/Adjusted demand per unit new.xlsx')
    usage_df = df.copy()
    usage_df = usage_df[usage_df['Artikelsoort'] == 'ZROH']
    usage_df = usage_df[usage_df['Match?'] == True]

    for x in ['unit price', 'EOQ', 'Alt Q', 's EOQ', 's alt Q', 'Annual demand', 'Total orders original', 'Total orders EOQ',
              'Total orders alt Q', 'Change EOQ', 'Change alt Q', 'Change s EOQ', 'Change s alt Q', 'Avg inv original', 'Avg inv EOQ',
              'Avg inv alt Q', 'Avg inv EOQ s', 'Avg inv alt Q s', 'Empirical fr orig Q', 'Empirical fr current Q']:
        usage_df[x] = None

    new_df = analyse(usage_df)
    new_df.to_excel('Analysis.xlsx')
    print(f'Runtime: {time.time() - s}')

def analyse(usage_df):
    daily_df = pd.read_excel('data/Dagelijks.xlsx')
    for index, product in usage_df.iterrows():
        demand = product['01.2023':'52.2023'].copy()
        product['Mean'] = np.mean(demand)
        product['unit price'] = product['Price per batch'] / product['Aantal ST in eenheid']
        D_y = product['Mean'] * 1 / 7 * 365

        if math.isnan(product['Price per batch']):
            product['EOQ'] = product['Original Q']

        elif product['Category'] == 'smooth' or product['Category'] == 'erratic':
            holding_costs = product['Price per batch'] * holding_cost_rate
            D_y = product['Mean'] * 1 / 7 * 365
            product['EOQ'] = compute_eoq(D_y, order_costs, holding_costs)

        else:
            product['EOQ'] = math.ceil(D_y / 365 * Q_days['general'])

        product['Alt Q'] = math.ceil(D_y / 365 * Q_days[product['Category']])

        if product['Category'] == 'smooth':
            product['s EOQ'] = math.ceil(compute_reorder_point_smooth(demand, product['EOQ']))
            product['s alt Q'] = math.ceil(compute_reorder_point_smooth(demand, product['Alt Q']))
```

```python

        elif product['Category'] == 'erratic':
            product['s EOQ'] = math.ceil(compute_reorder_point_erratic(demand,
            ↪ product['EOQ']))
            product['s alt Q'] = math.ceil(compute_reorder_point_smooth(demand,
            ↪ product['Alt Q']))

        elif product['Category'] == 'intermittent' or 'lumpy':
            matching_index = np.where(product['Artikel'] == daily_df['Artikel'])[0]
            daily_data = daily_df.iloc[matching_index, 2:].to_numpy()[0] /
            ↪ product['Aantal ST in eenheid']
            product['s EOQ'] = math.ceil(compute_reorder_point_irregular(daily_data,
            ↪ product, product['EOQ'], 'orig'))
            product['s alt Q'] =
            ↪ math.ceil(compute_reorder_point_irregular(daily_data, product,
            ↪ product['Alt Q'], 'alt'))

        product['Annual demand'] = D_y
        product['Total orders original'] = D_y / product['Original Q']
        product['Total orders EOQ'] = D_y / product['EOQ']
        product['Total orders alt Q'] = D_y / product['Alt Q']
        product['Change s EOQ'] = (product['s EOQ'] - product['Original s']) /
        ↪ product['Original s']
        product['Change s alt Q'] = (product['s alt Q'] - product['Original s']) /
        ↪ product['Original s']
        product['Change EOQ'] = (product['EOQ'] - product['Original Q']) /
        ↪ product['Original Q']
        product['Change alt Q'] = (product['Alt Q'] - product['Original Q']) /
        ↪ product['Original Q']
        product['Avg inv original'] = product['Original s'] + product['Original Q']/2
        product['Avg inv EOQ'] = product['Original s'] + product['EOQ']/2
        product['Avg inv alt Q'] = product['Original s'] + product['Alt Q']/2
        product['Avg inv EOQ s'] = product['s EOQ'] + product['EOQ']/2
        product['Avg inv alt Q s'] = product['s alt Q'] + product['Alt Q']/2
        usage_df.loc[index] = product

    return usage_df

if __name__ == '__main__':
    main()
```

B.4 Determining reorder points

Remark. This section includes the code for each technique corresponding to an ADI-CV category. Intermittent and lumpy demand are considered in one function, but we only add the extra safety inventory to lumpy demand through an if-statement. Note that *demand* is an array filled with the historical order data and *EOQ* is the previously computed order quantity Q.

B.4.1 Smooth demand

```python
def compute_reorder_point_smooth(demand, EOQ):
    sigma_L = np.std(demand) * np.sqrt(lead_time)
    mu_L = np.mean(demand) * lead_time

    a, b = (0, 10000)
    epsilon = 0.0001
    target = (1 - target_fill_rate) * EOQ / sigma_L
```

```

    m = (a + b) / 2

    while abs(G(m) - target) > epsilon:
        if G(m) > target:
            a = m
        else:
            b = m
        m = (a + b) / 2

    return mu_L + (m * sigma_L)

def G(z):
    return sct.norm.pdf(z) - z * (1 - sct.norm.cdf(z))
```

B.4.2 Erratic demand

```
def compute_reorder_point_erratic(demand, EOQ):
    sigma_L = np.std(demand) * np.sqrt(lead_time)
    mu_L = np.mean(demand) * lead_time

    alpha = (mu_L**2)/(sigma_L**2)
    beta = mu_L/(sigma_L**2)
    epsilon = 0.0001

    x, y = (0, 100000)
    m = (x+y)/2
    current_fill_rate = fill_rate_gamma(m, alpha, beta, EOQ)

    while abs(current_fill_rate - target_fill_rate) > epsilon:
        if target_fill_rate < current_fill_rate:
            y = m
        else:
            x = m
        m = (x + y)/2
        current_fill_rate = fill_rate_gamma(m, alpha, beta, EOQ)
    return m

def fill_rate_gamma(s, a, b, q):
    esc = (a/b*(1 - sct.gamma.cdf(s, a=a+1, scale= 1/b))) - (s * (1 -
    ↪  sct.gamma.cdf(s, a=a, scale= 1/b)))
    return 1 - (esc/q)
```

B.4.3 Irregular demand

```
def compute_reorder_point_irregular(demand, product, Q, type):
    sigma_L = np.std(demand)
    mu_L = np.mean(demand)

    l = (sigma_L**2 + mu_L**2 - mu_L) / mu_L
    l = l if mu_L < sigma_L**2 else mu_L
    np.seterr(divide='raise')
    try:
        p = (sigma_L**2 - mu_L) / (sigma_L**2 + mu_L**2 - mu_L)
        p = (1 - p) if p < 1 else 0
    except:
        return product['Original s']

```

```
    n = 2
    cum_dist = sct.binom.pmf(0, n, p)

    for i in range(1, n+1):
        cum_dist += sct.binom.pmf(i, n, p) * sct.poisson.pmf(0, i*l)
    j = 1

    while cum_dist <= target_service_level:
        for i in range(1, n+1):
            cum_dist += sct.binom.pmf(i, n, p) * sct.poisson.pmf(j, i*l)
        j += 1
    reorder_point = j

    if type == 'alt':
        Q = product['Original Q']

    ESC, fr, add = compute_add_safety_inv(demand, reorder_point, Q)

    if product['Category'] == 'lumpy':
        reorder_point += add
        ESC, fr, add = compute_add_safety_inv(demand, reorder_point, Q)

    if type == 'orig':
        product['Empirical fr new Q'] = fr
    else:
        product['Empirical fr current Q'] = fr
    return reorder_point

def compute_add_safety_inv(demand, s, Q):
    order = {0:0}
    inv = {0:Q+s}
    ESC, n = 0, 0

    for i in range(len(demand)):
        inv[i+1] = inv[i] + order[i] - demand[i]
        if inv[i+1] <= s:
            order[i+1] = Q
            ESC = n/(n+1)*ESC + abs(min(0, inv[i+1])) / (n+1)
            n += 1
        else:
            order[i+1] = 0

    fr = 1 - ESC/Q

    if fr < target_fill_rate:
        return ESC, fr, ESC
    else:
        return ESC, fr, 0
```