2023

Increase the on-hand availability of surgical supplies at the Operating Room department by improving inventory management

MASTER THESIS

ISALA ZWOLLE | UNIVERSITY OF TWENTE

M.B. KLOOSTERMAN

Document

Title: Increase the on-hand availability of supplies at the Operating Room department by improving inventory management
Date: 19-04-2023
City: Zwolle

MSc Industrial Engineering and Management
University of Twente

Host organization
Isala Zwolle (Procurement department)
Dr. van Deenweg 1
8025 BP Zwolle
The Netherlands
T: 038 424 6220
www.isala.nl

External supervisor
J.H. Bosker (Jacqueline)

Educational Institution
University of Twente
Drienerlolaan 5
7522 NB Enschede
The Netherlands
T: 053 489 9111
www.utwente.nl

First supervisor
Dr. ir. A.G. Leeftink (Gréanne)
Centre for Healthcare Operations Improvement & Research (CHOIR)

Second supervisor
Dr. P.B. Rogetzer (Patricia)
Faculty of Behavioural, Management and Social Sciences

Author
Max Kloosterman (s1862308)
MSc Industrial Engineering and Management

Management summary

There are many supplies involved in providing care for patients in healthcare organizations. A hospital's procurement department is responsible for supplying these goods. This is a complicated task, as the inventory concerns a wide variety of items, tight operating schedules, and limited storage capacity. Furthermore, hospitals need to deal with emergency patients and other unpredictable circumstances. Proper inventory management and good agreements with the suppliers are necessary to guarantee the availability of materials. When this is not the case, it leads to disruptions of the main processes. In the best case, this only leads to emergency orders, which take more time and are costly. However, it can also affect the quality of care when materials are not available in time, which should be avoided. On the contrary, overstocking supplies would result in unnecessary costs and waste. Furthermore, it is impossible due to hospital storage capacity limitations.

This research aims to improve local inventory management in a healthcare setting, such that the on-hand availability of surgical supplies increases and the amount of overstocking and number of emergency orders decrease. In this research, we develop a simulation model that provides users with intuitive experiences as they can observe how the inventory behaves over time. We perform a case study on the storerooms of the operations room (OR) department of Isala to determine the applicability in practice. Isala Hospital provides care to the inhabitants living in, or close to, the region Zwolle. Therefore, our research question is:

"How can the on-hand availability of surgical supplies at the Operating Room department of Isala be increased by improving its inventory management?"

The supplies in this research are classified as stock or purchase items. The first step in improving inventory management is by applying clear replenishment policies for these items. We performed a literature study to find alternative replenishment policies.

Stock items have a constant review period and a short lead time. The lead time for stock items is usually less than a day as Isala is supplied multiple times per day from the external warehouse by Isala's logistical partner, Hospital Logistics (HL). The mean usage of these items is higher than the usage of purchase items, and therefore, we use a 2-bin Kanban policy with equal bin sizes, denoted as an (R, r, Q) system, where the reorder point (r) equals the order quantity (Q). This system decreases the workload of employees and the built-in stock rotation reduces the risk of products expiring, leading to better ergonomics.

Purchase items have a constant review period and stochastic lead time. The items can be classified as slow movers due to having a daily usage of fewer than ten units. We use an order-up-to-level replenishment policy with a variable order size, denoted as an (R, s, S) system to use the available capacity efficiently.

In the next step, we developed a simulation model for OR department inventory management. This simulation model uses the developed search heuristics to find the optimal or near-optimal order parameters per item. The optimal order parameters found by the simulation model and its search heuristic significantly outperform the case study's current order parameters in terms of item availability and costs. The average availability increases from 98.92% to at least 99.93% and the costs reduce by at least 8%. The general tendency was that for most items the current order parameters resulted in high inventory levels, and, consequently, high holding costs. On the contrary, other supplies in the current situation have order parameters which are too low. As a result, the shortage costs and number of stockouts increase.

For stock items, the optimal parameters can reduce the average shortage costs by approximately €258 per year, while the yearly average ordering- and holding costs would only slightly increase by €26 and €3, respectively. The optimal parameters reduce the yearly average ordering- and shortage costs of

purchase items by €87 and €28, respectively compared to the initial parameters, while the average holding costs would only slightly increase by €3 per year. When we compare the results of the service-level model to the current parameters for all the items in the dataset (n=1,754), the average availability of items increases from 98.92% to 99.93%, and the total costs of the inventory system reduce by approximately €265,000 a year (from €3,190,000 to €2,925,000). The optimal parameters of the model with a cost objective increase the average availability of items to 100.00% and reduces the total costs of the inventory by approximately €355,000 a year compared to the current order parameters.

By adapting the model's base scenario, several "What-if" scenarios are conducted to observe the impact on the inventory system. The scenario with a varying lead time for supplies that are not delivered by HL shows that there are 50 purchase items with a decrease in their order parameters. These items have the potential to store them at HL, however, storing items at HL can entail additional costs. The other experiment in which we vary the demand variability shows that less variability reduces the required space and average costs for items with daily usage of at least ten units by 15% and 9%, respectively.

We have four recommendations to increase the on-hand availability of surgical supplies at the OR department of Isala. First, we recommend implementing a 2-bin Kanban system as the replenishment policy for stock items and a (*R*, *s*, *S*) replenishment policy for the purchase items. Next, we recommend the OR department of Isala to use the simulation model in their decision-making about setting the order parameters. The order parameters should be regularly revised to cope with non-stationary demand. We recommend doing this at least once every six months, which requires updating the demand distributions and other input data. Next, it is recommended to capture the actual demand for items and the service level of the inventory system. More accurate data further improves the performance of the simulation model. Last, we recommend including the package sizes of items to find more accurate order parameters and possible faster computation times.

This research contributes to the literature on healthcare inventory management. We developed a simulation model similar to Zhang et al. (2014). We expanded their work by including fractional lead times in the model. Furthermore, the search algorithm we use to find near-optimal and optimal order parameters is a simplified version of the algorithm developed by Kapalka et al. (1999), while we extended the algorithm of Esmaili et al. (2019) by also considering costs. In addition, we formulated two variants of the algorithm, as the 2-bin Kanban system and (*R*, *s*, *S*) system have varying characteristics concerning the order parameters. We performed a case study on the storerooms of the OR department of Isala to determine the model's applicability in practice.

Acknowledgements

This thesis is written to complete the Master of Industrial Engineering and Management at the University of Twente. Writing this thesis at Isala Zwolle has been a learning experience, and I look back on it as a great time. Without several people, it was not possible to perform this research. I want to take this opportunity to thank them.

First, I thank my colleagues at Isala. Many of them provided me with useful information and interest in my research. In particular, I acknowledge Jacqueline Bosker for allowing me to conduct my graduation project at Isala. Moreover, Hans Doornink and Jeroen van de Nes provided me with guidance and information about the operating room department during this project.

Furthermore, I thank my supervisors from the University of Twente, Gréanne Leeftink and Patricia Rogetzer. Both of you provided me with great feedback. And I am very grateful that I had both of you as supervisors. Also, I enjoyed the support and conversations of my fellow interns, Josien Mourik, Judith de Boer, and Leon van Beek. The weekly meetings we had with Gréanne were always helpful and joyful.

Lastly, I thank my girlfriend Meike, my parents Wibo and Miranda, and my brother Gijs and his girlfriend Amber for their support during my studies. In addition, I thank the rest of my family and friends for their motivational support.

Contents

1. Introduction

This chapter describes the context of the research conducted in five sections. Section 1.1 provides background information about Isala Zwolle. Section 1.2 outlines the research motivation. Then, Section 1.3 introduces the problem. Based on this analysis, the problem-solving approach and research objective are defined in Section 1.4. The last section describes the scope and deliverables of the research (Section 1.5).

1.1 Background

This section introduces the company, the item types, and the inventory system in the current situation.

1.1.1 Isala Zwolle

The healthcare sector is complex, with a wide variety of services provided to patients by healthcare providers. Isala is a healthcare organisation in the region of Zwolle. They provide hospital care to the 690,000 inhabitants that live in, or close to, this region (Dillmann, 2022). At Isala the patient value comes in the first place. They aim to optimally recover, maintain, and enhance the patient's quality of life. The organisation's core principles are professionalism, heart and soul, and transparency. Isala hospital is one of the twenty-seven top clinical hospitals in the Netherlands, with a total of 1,250 available beds (STZ, 2022). There are 427 medical specialists (394 full-time equivalents (FTE)), and a total of 6,797 employees (5,081 FTE) working at Isala (Isala, 2022). The two hospitals of Isala are in Zwolle and Meppel, these are their main facilities. The hospital in Zwolle has an operating room (OR) capacity of twenty, from which there are fourteen clinical ORs and six ORs available for day treatments. The hospital in Meppel consists of five day-treatment ORs. The other facilities of Isala, which are outpatient clinics, are in Kampen, Heerde, and Steenwijk.

Without supplies, the employees of Isala cannot perform their daily tasks of helping the patients. Figure 1 visualises the supply chain of Isala Zwolle, where the local inventory is stored in the "Department storerooms" and "Point of use" locations (Ahmadi et al., 2018). With the build of their new hospital in 2012, Isala decided to change to a just-in-time (JIT) inventory system. This strategy was first introduced by hospitals in the US and Canada during the 1990s to reduce on-hand inventory and costs (Neve & Schmidt, 2022; Rivard-Royer et al., 2002). With this strategy, Isala receives the goods as closely as possible when needed, and therefore, has more space in the hospital for patient care. However, this means that they had to adjust their inventory system. Instead of a central warehouse where employees could pick up everything when needed, Isala outsourced its warehouse.

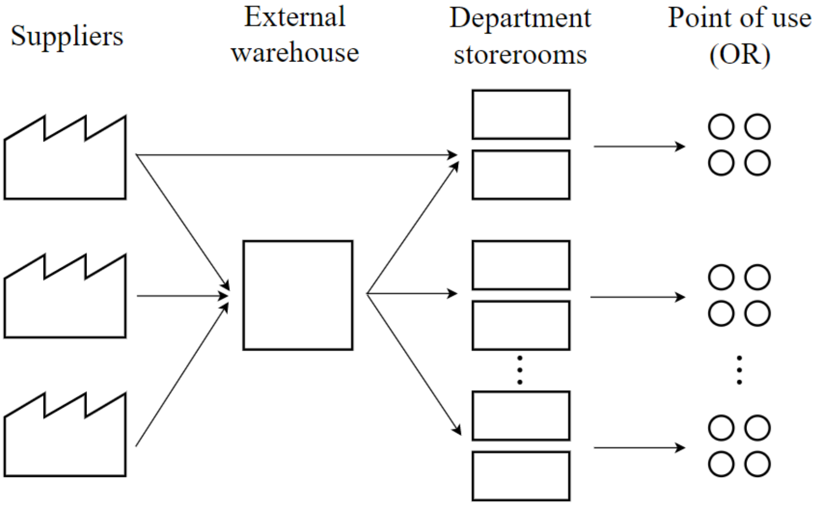

Figure 1: Visualisation of the supply chain of Isala Zwolle.

The external warehouse is managed by its logistics partner, Hospital Logistics (HL). Most of the items the hospital uses are supplied by HL. However, some supplies are supplied directly to Isala. A more detailed description of the different items and suppliers is given in the next section (Section 1.1.2).

When supplies arrive at Isala, the logistics department of Isala receives them. They bring the supplies to the correct department. At that department, there is another logistic employee who stores them in the right storerooms. The logistics department does this for all departments of Isala except for the OR departments. Supplies are brought to the front of the OR department by logistics; however, the OR department has its own logistics department, which takes the supplies and stores them in the right storerooms. Within an OR, no supplies are stored. When items are needed, they are picked from the department storerooms by an employee of OR logistics. Vila-Parrish & Ivy (2013) would describe the supply chain of Isala Zwolle (Figure 1) as a two-echelon supply chain. Echelons are defined as physical locations where supplies are stocked.

1.1.2 Types of items

A hospital uses many materials, from medical supplies and pharmaceuticals to equipment. The OR department storerooms only store surgical supplies. Therefore, only these items are included in the research. We distinguish three types of disposable surgical supplies: (1) stock items, (2) purchase items, and (3) scan-relevant items. These items are all used on a frequent base and ordered from one of the suppliers of Isala. Surgical instruments could also be seen as surgical supplies. However, these are not ordered but instead sterilized after usage in the sterile processing department, after which they can be used again. Therefore, we exclude these from this research.

Stock items are products that are stored in bulk in the external warehouse of the hospital. These items are used daily with a relatively high volume. The items are supplied to Isala by HL multiple times per day. This means that HL should have enough supplies in its warehouse to fulfil the daily orders of Isala. However, the procurement department of Isala has contact with the suppliers about the product details and decides which products HL must store in the warehouse. In this way, when HL wants to replenish its stock levels, it only needs to place an order at the predetermined supplier.

Purchase articles are products that are not stocked in the external warehouse of the hospital and are, therefore, called “non-stock” items (Landry & Beaulieu, 2013). HL may deliver purchase articles to Isala. When this happens, the supplier ships it to the cross-dock of HL. From there, HL sends the products to Isala without stocking them in the external warehouse. This is the case with articles with a low usage rate. However, there are three reasons why another supplier than HL "directly" sends the item to Isala. First, the purchase articles should be stored in a cold environment, which HL cannot provide. Next, the items are too big. HL delivers the products in roll containers. However, when an article does not fit in these containers, HL cannot supply it. Lastly, the product is too expensive. If at least one of the reasons is applicable for the product, it is sent to Isala by another supplier than HL.

Scan-relevant items are also called “non-stock” items. The hospital stores them in the storerooms of a department. The items have a barcode such that they can be scanned when used. When this happens, the system’s inventory level of the item decreases. Most of the time, scan-relevant items are expensive. Therefore, the hospital wants to track the inventory levels of these products. Products that are also considered scan-relevant are implants. These are not always expensive; however, all implants have a unique barcode, which makes them traceable to the patient. Hospitals are obligated to register the implants by government regulations (Rijksoverheid, 2022). In this way, a patient can receive help quickly when there is something wrong with their implant.

1.1.3 Inventory system

As mentioned in Section 1.1.1, Isala has a two-echelon supply chain where supplies are stored at the external warehouse and the department storerooms at Isala. Isala owns the items at both locations; however, it is only responsible for the inventory (policy) at the department storerooms (in the hospital), which are the warehouses at the lowest echelon where the local inventory is stored (Ahmadi et al., 2018). Therefore, Isala has a multi-echelon inventory system environment with an inventory control approach defined as independent single-echelon inventory control policies (Hausman & Erkip, 1994). The number of rooms varies for different departments, which could be divided over multiple levels. This is also the case for the OR department, which has eight storerooms divided over two floors. OR logistics employees prepare the containers with the items for a particular treatment. They retrieve the necessary items from the storerooms at the OR department. A surgery requires different supplies when performed by individual physicians, based on their preference card. The orders from all storerooms are combined into one order and sent to the supplier to replenish the department storerooms. In this system, the department storerooms are responsible for their stocking policies, independent of each other, suppliers, and the external warehouse. Therefore, they can apply replenishment policies that require a multi-item, single-echelon system. Currently, Isala uses different control policies for (1) stock and purchase items and (2) scan-relevant items, see Table 1.

Table 1: Inventory control policy per item type.

Item type	Inventory control policy	Symbol
Stock items	*Periodic review*	(R, s, S)
Purchase items	*Periodic review*	(R, s, S)
Scan-relevant items	*Continuous review*	(r, Q)

The inventory policy for stock and purchase items is described as a *periodic review system* with symbols (R, s, S) (Ahmadi et al., 2018). Every day (R), an OR logistics employee checks the inventory by going by all storeroom shelves. If the inventory level is equal to or less than the "Min" value s, the employee scans the item to replenish the inventory up to the "Max" value S. The "Min" and "Max" parameter values differ per item. Even though the parameters are set, employees frequently overrule them based on personal experience or because physicians ask them to do so. Therefore, items may be ordered before the "Min" level is reached and can be ordered to a level above the "Max" level. The stock items are delivered on the same day unless there is a backorder, whereas purchase items have longer lead times.

Scan-relevant items follow a different inventory policy, which can be described as a *continuous review system* with symbols (r, Q) (Ahmadi et al., 2018). As mentioned in Section 1.1.2, Isala uses a bar-code scanning system in an open-inventory environment. Nurses scan the item when it is used in the OR, which automatically decreases the number of stock items in the system. Hawa (2020) describes the bar-code system as an active tracking system as items need to be scanned to be registered in the system. Whenever an inventory level for a given item reaches or goes below the "Min" value r, an order is placed with a constant size Q.

1.2 Availability of surgical supplies

Equipment and supplies are required to perform treatments. A hospital's procurement department is responsible for supplying these goods. This is a complicated task, as the inventory concerns a wide variety of items, the operating schedules are tight, and there is limited storage capacity. Furthermore, hospitals need to deal with emergency patients and other unpredictable circumstances. Currently, there is a shortage of resources all over the world (RTL Nieuws, 2022). Due to this global shortage, not all surgical supplies are available on time. When the initial product is not available, the procurement departments must look for an alternative product to ensure that treatments can still be performed as scheduled. Proper inventory management and good agreements with the suppliers are necessary to guarantee the availability of materials. When this is not the case, it leads to disruptions of the main processes. In the best case, this only leads to emergency orders, which take more time and are costly. However, it can also affect the quality of care when materials are not available in time, which should be avoided. Therefore, Isala wants to increase the availability of surgical supplies to their desired service level (see Section 2.4).

To have high availability of surgical supplies, the supplies should be on-hand when they are needed. However, due to space limitations, order sizes should be limited. Furthermore, large order sizes can lead to overstocking, which would result in unnecessary costs and waste. Therefore, the right balance of the desired service level and inventory levels is necessary.

1.3 Problem analysis

As described in Section 1.2, Isala's goal is that treatments should not be delayed or cancelled due to the unavailability of surgical supplies. The problem analysis focuses on the inventory for the ORs in the OR department. This inventory contains more items than other hospital inventories as it must also deal with emergency patients. Currently, most of the orders with urgency are from that department. Emergency orders exist to prevent stock-outs; however, such orders cost the hospital additional time and money. In addition, it is also not an option to overstock the inventory, due to space restrictions. Different problems trigger an emergency order or an overstock of the department storerooms. From observing the inventories at the OR department and interviewing the employees, we indicate the causes and effects of those actions and display them in a problem cluster, see Figure 2.

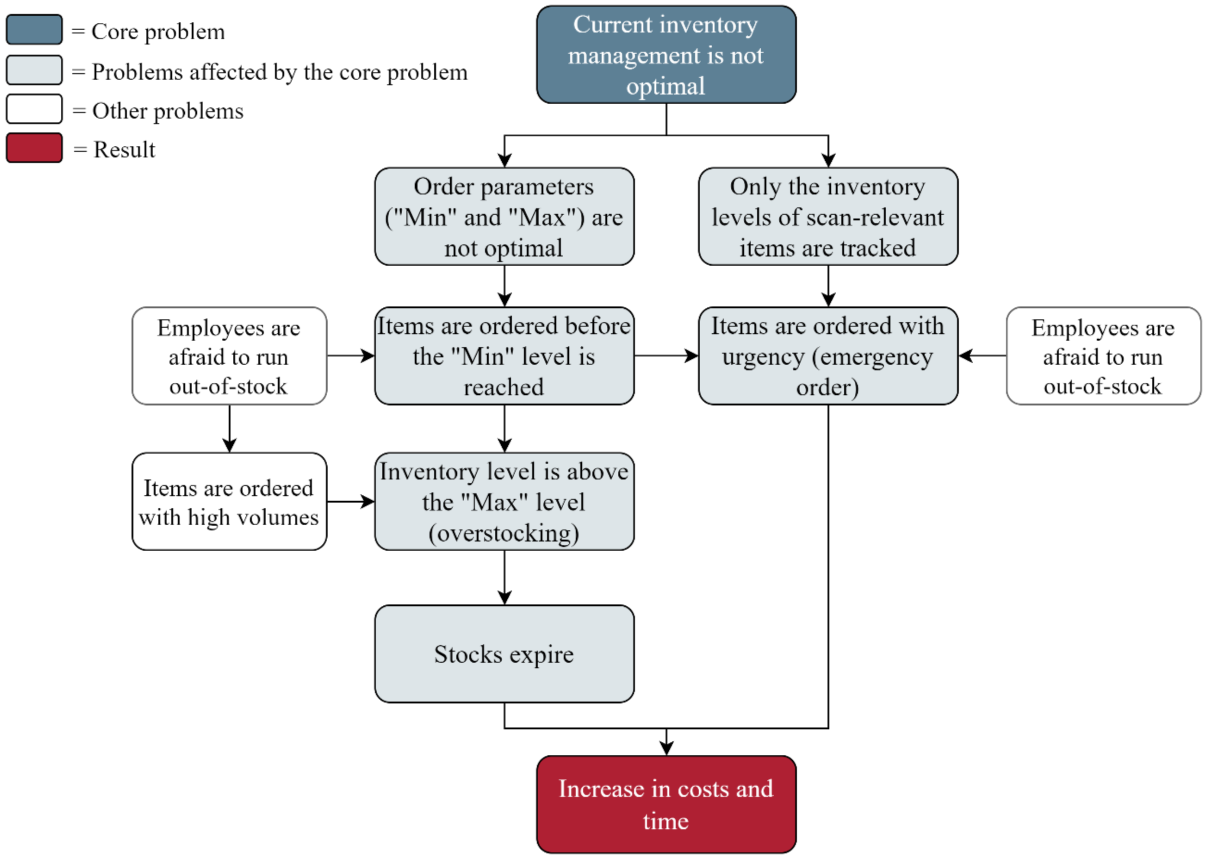

Figure 2: Problem cluster.

The problem cluster visualizes the relations between the problems. From this, we find the core problem solved in this research. Heerkens & van Winden (2017) define the core problem as the problem whose solution will make a real difference.

The core problem is "The current inventory management is not optimal". An effect of this is that the inventory parameters of items are not optimal. For certain products, it means that the "Min" value is too low, which triggers employees to order items in advance. On the other hand, other products have too many items in stock. These stocks are slow movers but take up too much space and can expire when they are not used for a long time. In addition, in the current inventory management system, only scan-relevant items are tracked. The stock levels at the storerooms are unknown for stock and purchase items. In combination with the parameters, which are not optimal, stock levels can become too low. Then employees would order items with urgency to prevent stock-outs. Overstocking and emergency orders lead to increasing costs and time, which optimisation of inventory management could prevent.

Another problem defined is “Employees are afraid to run out-of-stock”. This is not solely the case for Isala, as other hospitals experience the same behaviour (Ahmadi et al., 2018; Veral & Rosen, 2013). Employees can order with higher volumes, exceeding an item’s “Max” stock number. In addition, they often place an order before the item’s “Min” level is reached. Most of the time, employees want a feeling of safety. However, employees also do this because they experience an increase in demand for the item, or when a physician asks them to do so. If so, it also means that inventory management parameters are not optimal.

1.4 Research objective

This research aims to improve local inventory management in a healthcare setting, such that the on-hand availability of surgical supplies increases and the amount of overstocking and number of emergency orders decrease. In this research, we develop a simulation model that provides users with intuitive experiences as they can observe how the inventory behaves over time. The implemented search heuristic in the model delivers optimal or near-optimal order parameter settings of all stored items. This simulation model contributes to the research in healthcare inventory management. We perform a case study on the storerooms of the OR department of Isala to determine the applicability in practice. Therefore, our research question is:

“How can the on-hand availability of surgical supplies at the Operating Room department of Isala be increased by improving its inventory management?”

We divide the main research question into several sub-questions, which outline the rest of this research's structure and approach (see Figure 3). After answering these sub-questions, we can answer the main research question.

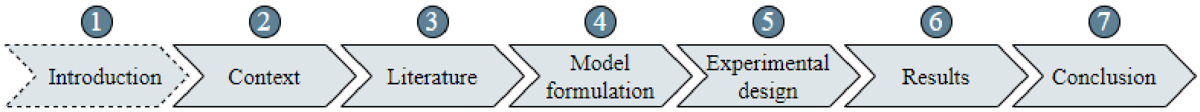

Figure 3: Report structure.

Chapter 2. Context

Sub-question 1: “What are the characteristics of the current inventory policy used in the OR department of Isala?”

- *Which parameters are used in the current inventory policy, such as lead times, usage of items, storage space, (emergency) order frequency, costs, item dimensions, and the “Min” and “Max” values of stock levels?*
- *How do the current storerooms look like and how much space is there for the items?*
- *How is the performance of the inventory policy measured?*
- *How to define the availability of surgical supplies?*
- *What is Isala’s desired availability of supplies?*

In this chapter, we gain insight into the characteristics of the current situation of Isala regarding its inventory management, focusing on the OR department. We gather information about how the order parameters are determined and analyse the (emergency) order data. Furthermore, we determine what situation would be desirable for Isala, in terms of the availability of supplies.

Chapter 3. Literature

Sub-question 2: "What inventory replenishment policy and approach is most applicable to different item types?"

- *What inventory policies exist under the assumption of lost sales?*
- *What other replenishment policies are used in hospitals?*
- *What inventory modelling methods to generate solutions in a healthcare setting are present in the literature?*

This chapter outlines a literature review. The review discusses inventory management replenishment policies and different approaches to define solutions for two item types: stock and purchase items.

Chapter 4. Model formulation

Sub-question 3: "How to develop an optimization approach that models inventory management?"

- *How can the objective be formulated?*
- *What are the model's indices, parameters, and decision variables?*
- *Which constraints should be present in the model?*
- *What stochastic elements are involved in the model?*

We present an inventory management model formulation for Isala in Chapter 4. With this model, the hospital should be able to find the right balance for their inventory levels with an increase in the availability of items and a decrease in costs.

Chapter 5. Experimental design

Sub-question 4: "How should the experimental design look like?"

- *How does the model perform on theoretical experiments?*
- *How does the model perform on the case study (OR department of Isala)?*
- *Which KPIs are selected to compare the scenarios?*

Chapter 5 outlines the different experiments and how they are compared with each other.

Chapter 6. Results

Sub-question 5: "What are the results of the experiments?"

- *What are the results of the case study?*
- *What are the results of the different scenarios?*
- *How can we visualize the results?*

In Chapter 6, we visualize the results from the experimental design.

Chapter 7. Conclusion and recommendations

Sub-question 6: "What are the recommendations to Isala?"

- *Which conclusion can we draw from the results from Chapter 6?*
- *What would be the best policy for Isala to improve its inventory management?*
- *What are suggestions for further research?*

Based on the results from Chapter 6, we present the potential impact of our modelling approach and provide recommendations to Isala on how it can improve the inventory management of its OR department.

1.5 Scope and deliverables

In this section, we outline the scope of this research and the deliverables to Isala. The deliverables are derived from the approach mentioned in the previous section (Section 1.4).

1.5.1 Scope

- Only the stock and purchase items mentioned in Section 1.1.2 are included in this research.
- Only the inventory management and storerooms at the OR department are part of the research.
- The picking process at the storerooms is outside the scope of this research.
- The model we developed in this research only applies to the original items currently in the assortment of Isala. Even though there are currently a lot of items added to the assortment, they may only be temporary items and might have no sufficient data.

1.5.2 Deliverables

1. An analysis of the storerooms at the OR department.
2. An overview of alternative inventory management policies.
3. A model formulation to improve inventory management, which preferably could also be used at the hospital in Meppel.
4. Results of the experimental design.
5. Recommendation regarding inventory management of Isala.

2. Context

This chapter outlines the current situation at Isala. The chapter describes the characteristics of the OR department's inventory management and is divided into five sections. First, we address the current replenishment policy, parameters, and locations of the stock and purchase items (Section 2.1). Second, the (emergency) order and supply processes are outlined (Section 2.2). Then, in Section 2.3, we gather the available order data, which we analyse to further substantiate our research as to why the current replenishment policy is not optimal. Next, we analyse the current service level and describe the desired service level (Section 2.4). We conclude this chapter in Section 2.5.

2.1 Inventory control

This section outlines the order parameters of the current replenishment and storage policy in use at the OR department of Isala. The main part of this research is to improve the inventory management, as we concluded that the current policy is not optimal. Therefore, it is important to identify the current policy and its characteristics.

The current replenishment policy for the stock and purchase items is an (R, s, S) policy. In this policy, Isala uses the terms "Min" and "Max" to define the parameters for their reorder point s and order-up-to-level S, respectively. For almost all the items, the "Max" value is twice the "Min" value with a Min value based on personal experience, which are influenced by several factors. The first factor is the storage location. Various storerooms can store the same product. However, the order parameters for a particular item can vary per location, dependent on the available space. The next factor is the dimension of an item. These are used to eyeball whether the products fit in the available space. However, no data is available about the specific item dimensions, which could be used to calculate the maximum number of items that fit into a bin. Another factor that could influence the order parameters is the package size of a product. Some parameters are rounded to the nearest package size quantity, as some suppliers only deliver full packages. The item demand also influences the order parameters. However, the demand is experienced based as the usage data is not stored. Another factor is the classification of an item. Items are not classified based on a certain method but on two categories. The first one is the item type. As mentioned in Section 1.1, there are three types of items at Isala (stock, purchase, and scan-relevant items), from which we only focus on stock and purchase items. The other category is medical or non-medical items. According to the employees of the OR department, all medical-related items could be necessary for (emergency) operations. The last overlooked factor is the inventory costs because the main priority of the hospital is to provide care to its patients.

When employees order an item, the order quantity is most of the time equal to the "Min" value unless employees expect higher demand in the upcoming days. If the latter is true, the order size is based on the employee's expected demand for the item. The order parameters of an item are noted on the scan card. This card is attached to the bin in which the item is stored. The parameters are not regularly revised because the cards should then be changed as well, which is time-consuming. Therefore, the parameters are often not up to date and do not meet non-stationary demand. Thus, when employees expect more demand, they order an amount of the item based on their experience. These order sizes are mostly larger than necessary due to the uncertainty of the processes in a hospital. In addition, the order sizes can fluctuate as employees react differently when an order quantity is too low. The stock levels are too high for other items, resulting in unnecessary holding costs.

2.2 Order and supply process

This section outlines the order and supply process, and the order data. Subsection 2.2.1 describes the supply process, which provides additional information about the regular and emergency supply process compared to Section 1.1. Furthermore, it outlines the storage of the items in the OR department. After that, we describe the order process in Subsection 2.2.2.

2.2.1 Supply process

All stock items and some purchase articles are delivered to Isala by its logistical partner, Hospital Logistics (HL). They supply Isala at four fixed times a day (7:00, 10:00, 13:00, and 16:15). The purchase articles that are not delivered by HL, are supplied by another supplier. When supplies are ordered with urgency, the logistics department of Isala, which receives the order, is notified in advance. They get an email from the supplier with information about the order. The items that are ordered with urgency are highlighted in this email. Once the supplies arrive at Isala, the logistics department of Isala receives them. They sort the supplies and bring them to the correct department.

In the OR department, the stock and purchase items are stored over twenty-four locations divided over two floors. These items are strictly used at the OR department for surgery, staff, or other activities in this department. Articles are either stored in a bin or on a shelve, depending on the dimensions of the article. A storage bin consists of two space components, with each component containing the "Min" amount of the item. A bin can have three sizes. However, when an item cannot fit into one of the three different bin sizes, it is stored on a shelve.

2.2.2 Order process

Ordering at the OR department happens at fixed review periods in which employees scan the cards of the articles that need to be replenished, which is the same for stock and purchase items. Storage locations that contain medical supplies are reviewed every working day. The storage locations of non-medical items are either reviewed every working day or on Monday, Wednesday, and Friday, which is location dependent. Employees do not review the inventory on Saturdays and Sundays as most surgeries are scheduled during the week, and therefore, the demand for items at the weekend is low.

The order size and frequency are based on the demand and order parameters of the item (see Section 2.1). Even though the daily demand for supplies is not captured and available as data, its structure is displayed in Figure 4. There are several surgeries in a day, spread over multiple ORs. Most of these surgeries are planned except for the surgeries of emergency patients that happen ad hoc. Every surgeon has a way of performing a type of surgery. The preference list contains the items they need, which are then picked at the storage locations for surgery.

The demand for items can fluctuate between the days because of the surgeries that are scheduled on certain days, changes in the schedule, the arrival of emergency patients, changes in the condition of the patient, and more. Fluctuations in demand result in even higher fluctuations in the order size. Another factor that influences the variability in order quantities is the level of outstanding backorders by the suppliers. However, the backorders are outside the scope of this research.

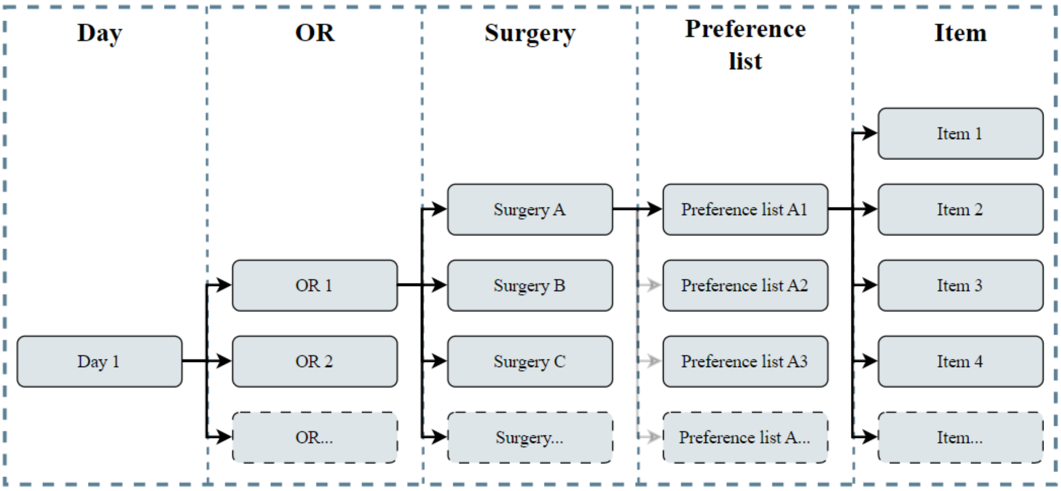

Figure 4: The daily demand structure of items.

When the OR department needs an item with urgency, an employee places a request for order (RFO). An RFO is a document with the exact information about the product, which is sent to the purchasing department. They contact the supplier and logistics department to verify whether the order can be fulfilled. If possible, the emergency order is forwarded to the supplier. Otherwise, an alternative product should be ordered, or the procedure for which the original item was necessary should be postponed.

2.3 Order data

An observation in our data analysis is that the data records concern order quantities and not the actual demand for an item, which is common for hospitals in literature (Bijvank, 2009). As a result, the available data about the usage of an article is limited. To use the data in our research, we define the mean daily usage as the mean order quantity per working day. Other observations from our data analysis concern the “Min” and “Max” levels and the order sizes of items. In the remainder of this subsection, we outline the supplies that stand out based on the available order data of 2022. Table 2 provides a summary of these items.

Table 2: Summary of outstanding items.

Symbol	Items	SKU’s	Description
M	*M1*	2	No consistent use of the policy
	M2	17	
D	*D1*	458	Items that have a “Min” value which is lower than their mean daily usage.
	D2	1585	
	D3	534	
	D4	163	
	D5	340	
	D6	71	
S	*S1*	387	Similar items have the same order parameters.
	S2	388	
	S3	389	
	S4	390	
	S5	391	
C	*C1*	226	Items with a “Max” inventory position of 0, 1, or 2 units. However, they are ordered in a fixed size of ten units.
	C2	616	
	C3	1676	
O	*O1*	400	Other notable items.
	O2	240	

No consistent use of the policy

In the current policy, the “Max” inventory level is twice the “Min” value of an item, except for items *M1* and *M2*. These items have a “Max” level which is ten times the “Min” value. There is no specific reason for these items to have that maximum storage capacity as the daily usage is significantly lower than the “Max” position. Furthermore, the lead time for those items is less than a day, which means that for the “Max” level, twice the “Min” value would satisfy demand as well. The inconsistent usage of the policy can confuse the employees when they need to decide on the order quantity and should therefore be avoided.

High mean daily use

We outline six items that have a “Min” value which is lower than their mean daily usage. As a result, the chance that one of the items runs out of stock is high. All these items are classified as medical items, which makes them crucial for the hospital. Item *D1* has an on-hand minimum of 100 units, a mean daily usage of 113, and is ordered 240 times in 2022. With such high usage numbers, the chance that the item runs out of stock is high. Figure 5 shows the other items with a high probability of running out of stock. Items *D2*, *D3*, and *D4* had an order frequency of more than 180 in 2022. Whereas products *D5* and *D6* had no "Min" and "Max" values assigned to them without a specific reason. In all these scenarios, the order parameters are not optimal.

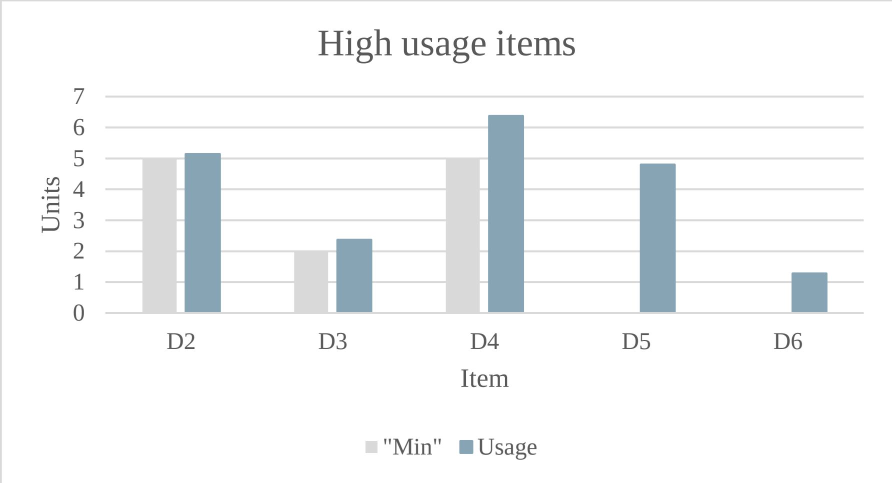

Figure 5: Items with a high mean daily usage compared to their minimum inventory value.

Similar items with the same order parameters

Similar items have the same order parameters for simplicity reasons. The products *S1*, *S2*, *S3*, *S4*, and *S5* are the same products but in different sizes. Therefore, the mean daily usage of these items varies as well. As a result, the order parameters for some of these products are too high, which is especially the case for item *S5* (see Figure 6).

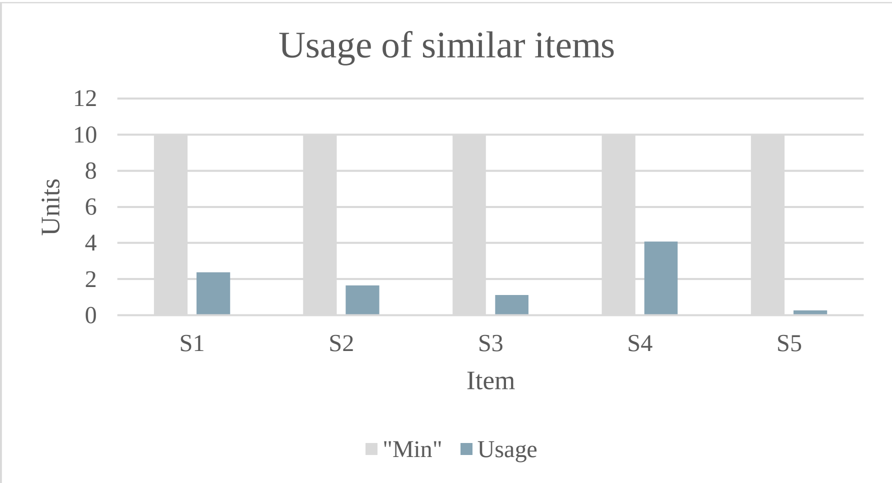

Figure 6: Similar items with the same order parameters but with different mean daily usage.

Neglecting the order parameters

When the order parameters are too low, the employees have to decide how large the order quantity must be. The products *C1*, *C2*, and *C3* have a "Max" inventory position of 0, 1, and 2 units, respectively. However, they are ordered in a fixed size of ten units.

Other items

There are two more items that we want to outline. The first article *O1*, has a “Min” and “Max” value of 100 and 200, respectively. However, the item is ordered 47 times with a mean order size of 248 in 2022. Indicating that the maximum inventory position is exceeded for a large part of the year. The other product *O2*, is only ordered once in 2022 and has a mean daily usage of 0.04 units. However, its order parameters are ten and twenty units for the “Min” and “Max”, respectively.

Concluding, the “Min” and “Max” values of the current inventory policy are not aligned with the demand, and the order sizes do not correspond with the replenishment policy.

2.4 Service level

In this section, we first describe how we can define the availability of items as a service level. Then, we analyse the current situation concerning the service level. Last, we describe the desired service level.

The availability of supplies is the same as the probability of not going out of stock in a period, defined as the service level alpha (Schneider, 1978). It indicates the percentage of fully fulfilled demand in a given time. When there is demand for a product, but it is not on hand, we refer to it as a stockout. Whenever this happens, the article is either derived from another storage location or arrives via an emergency order. However, only the emergency orders are documented via the RFO. Figure 7 shows the number of emergency orders for the OR department from 2017 to 2022 with an upwards trend. In 2022, there were 414 emergency orders out of the 51,848 orders related to the OR department, which means that in 99.2% of the cases, there was no emergency order necessary to fulfil demand. However, we cannot conclude that the OR department has a current service level of 99.2%. In reality, not all emergency orders are correctly registered regarding its department. In 2022, there were 1,027 emergency orders for which the department was unknown and could have been for the OR department. Furthermore, it can happen that the items are supplied to the OR department in Zwolle from the hospital in Meppel by an employee of Isala. However, there is no insight into how often this happens. In addition, some items are simply taken from another storage location, what is organised by employees among themselves. Therefore, the current service level of the OR department of Isala is probably lower than 99.2%.

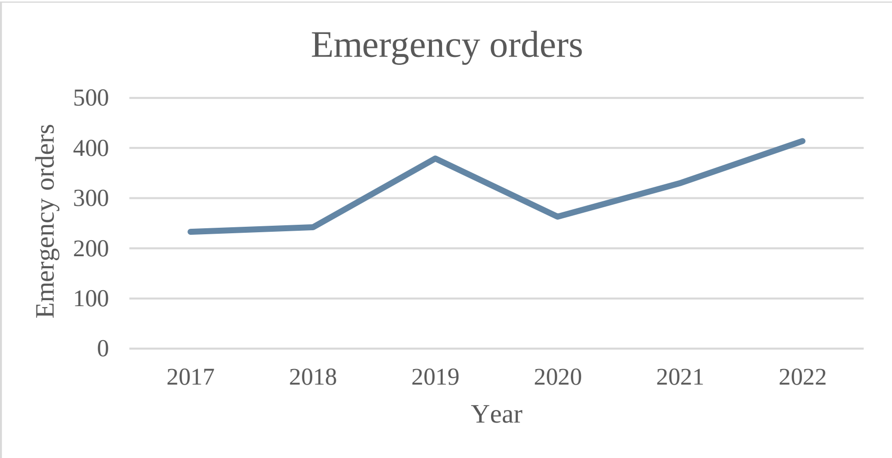

Figure 7: Number of emergency orders for the OR department from 2017 to 2022.

For Isala, it is important to have high availability of supplies such that the daily processes in the hospital are not disrupted (see Section 1.2), which is even more important for medical items. Therefore, Isala has a desired service level of 99.9% and 98.0% for medical and non-medical supplies, respectively.

2.5 Conclusion

The current replenishment policy for the stock and purchase items is an (*R, s, S*) policy. In this policy, Isala uses the terms “Min” and “Max” to define the parameters for their reorder point *s* and order-up-to-level *S*. The values for the order parameters are based on the experience of the employees. When ordering, employees cannot rely on these parameters, which results in high fluctuations in inventory positions and order quantities.

The order data confirmed our observation that the current inventory policy is not optimal. We outlined the items that stand out in our analysis. The policy is not consistently applied to some of these products. Other supplies have order parameters that were too high, just because similar items had those parameters as well. And some products have a reorder point lower than their mean daily use. As a result, these items have a high chance of running out of stock.

The upwards trend in the number of emergency orders is disturbing for Isala. To overcome this, high availability of supplies is important such that the daily processes in the hospital can continue. To reduce the number of emergency orders, Isala wants to increase the availability of supplies, with a desired service level of 99.9% and 98% for medical and non-medical supplies, respectively.

3. Literature

As discussed in Section 1.3, the current inventory management of Isala is not optimal. In this chapter, we answer the third research question: 'What inventory replenishment policy and approach is most applicable to different item types?'. The chapter is divided into three sections. Section 3.1 covers inventory management policies. In the section, we distinguish replenishment policies with a periodic and continuous review period and inventory systems under the assumption of back ordering and lost sales. Section 3.2 outlines background information about different inventory modelling methods to obtain optimal or near-optimal solutions. Also, we discuss the benefits of each method. We conclude this chapter in Section 3.3.

3.1 Inventory replenishment policies

The inventory system of Isala can be considered as a lost-sales inventory system, as when an item is not on-hand but is needed for surgery, an emergency order is triggered to get the item. Therefore, we look for replenishment policies for lost-sales inventory management that improve the current situation of Isala through inventory innovations regarding single-echelon systems in healthcare or other sectors. Literature separates the classic models for inventory control policies into two review periods: (1) Periodic review and (2) continuous review. Based on the review periods, five inventory control policies are distinguished (Ahmadi et al., 2018). Table 3 shows these policies. In the remainder of this section, we first search for periodic review replenishment policies in the literature, as Isala currently uses this review system for their stock and purchase items. Second, we discuss replenishment policies with continuous review. Isala uses this type of policy for its scan-relevant items. Recent technical innovations like radio frequency identification (RFID) make policies with a continuous review more applicable for practical situations, and therefore, offer opportunities for Isala in the future. In addition, we search for alternative replenishment policies without the assumption of lost sales because the assumption of back ordering is more common in literature.

Table 3: Classical inventory control policies, adapted from Ahmadi et al. (2018).

Inventory control policy	Notation	Description
Periodic review	(R, s, S)	Every review period R, if the inventory level is lower or equal to s, an order is placed to refill the inventory position to the "Max" inventory level S.
	(R, r, Q)	Every review period R, if the inventory level is lower or equal to r, an order with a constant amount Q is placed.
	(R, S)	Every review period R, an order is placed with an amount to refill the inventory position to the "Max" inventory level S.
Continuous review	(s, S)	If an inventory level reaches the reorder point s, an order is placed to refill the inventory position to up-to-level S.
	(r, Q)	Whenever the inventory level of an item reaches the reorder point r, an order with constant size Q is placed.

3.1.1 Periodic review

This section outlines some of the available literature on inventory systems with a periodic review period, as summarized in Table 4. In a periodic review, each review period R, the inventory levels are checked, and an order is placed when necessary.

Table 4: An overview of inventory models with a period review replenishment policy.

Policy	Studies	Method	Objective	Constraint	Lost sales
(*R*, *s*, *S*)	(Bijvank, 2009)	Exact and approximate	Service level	Storage capacity	Yes
	(Kapalka et al., 1999)	Heuristics	Costs	Service level	Yes
	(Esmaili et al., 2019)	Heuristics	Service level Costs	-	Yes
	(Bijvank & Vis, 2012b)	Exact and approximate	Costs	Service level	Yes
	(Zhang et al., 2014)	Heuristic and simulation		-	No
(*R*, *r*, *Q*)	(Bijvank & Vis, 2012a)	Exact and approximate	Capacity and service level	Storage capacity and service level	Yes
	(Kapalka et al., 1999)	Heuristics	Costs	Service level	Yes
2-Bin system	(Kanet & Wells, 2019)	Exact	Costs	Service level	No
(*R*, *S*)	(Huh et al., 2008)	Heuristics	Costs	-	Yes
	(Bijvank & Johansen, 2012)	Exact	Costs	Order-size	Yes
(*R*, *s*, *c*, *S*)	(Dellaert & van de Poel, 1996)	Exact and simulation	Costs	Service level	No
Hybrid	(Rosales et al., 2014)	Simulation	Costs	-	No
Linking patient schedule	(Epstein & Dexter, 2000)	Simulation	Costs	-	No

Variable order size (R, s, S) policies

Dependent on the review period, which is a periodic or continuous review, the order-up-to-level models are denoted as (*R*, *s*, *S*) or (*s*, *S*) policy, respectively (Bijvank, 2009). Figure 8 visualizes the inventory levels for the order-up-to-level replenishment policies under a periodic or continuous review. Order-up-to-level replenishment policies have a variable order size and therefore use the available capacity more efficiently than policies with a fixed order size *Q*. Kapalka et al. (1999) consider a periodic review single-echelon inventory system with stochastic demand and deterministic lead times under the assumption of lost sales. They evaluate the long-run average costs and service level for a fixed (*R*, *s*, *S*) policy. Then they would use their formulated Monotone Search Algorithm (MSA) to search for the optimal parameters. They found significant cost savings with their model. Esmaili et al. (2019) use a similar approach as Kapalka et al. (1999). They created a recursive algorithm to determine the optimal parameter values of the (*R*, *s*, *S*) replenishment policy. The system they studied has stochastic demand, lost sales, zero-lead time, and a target service level to be satisfied.

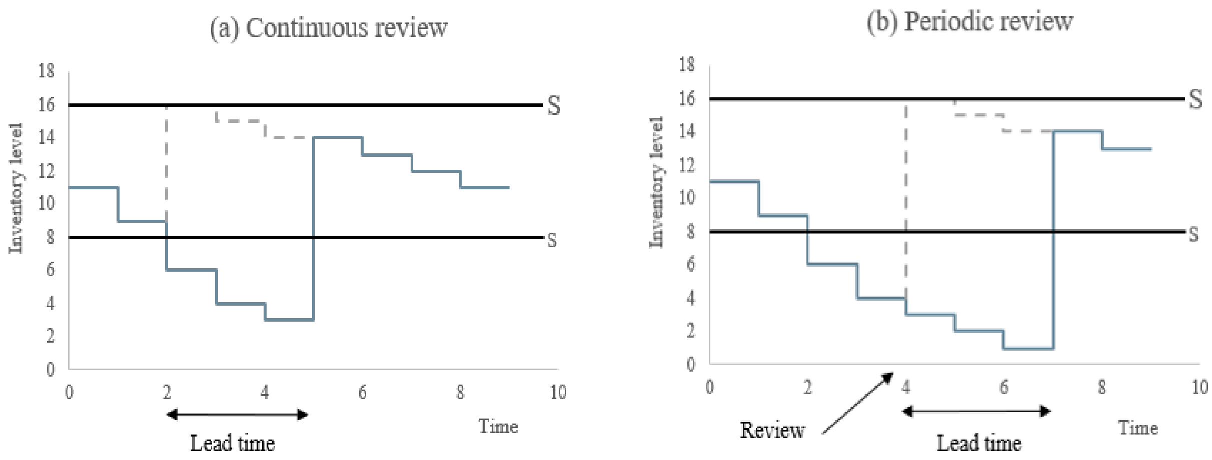

Figure 8: The on-hand inventory level (solid line) and inventory position (dashed line) for order-up-to-level replenishment policies under (a) continuous and (b) periodic review, adapted from Bijvank (2009).

Bijvank (2009) first developed an (*R, r, Q*) policy with reorder point *r* plus order size *Q*, equal to the storage capacity *C*. The order frequency that follows from that policy is used in the (*R, s, S*) policy. They conclude that a variable order size policy can significantly improve the service level. They create a simple spreadsheet approximation to make it more appealing for hospitals to apply the inventory rule. However, in the case study, the order-up-to-level *S* is set to the maximum storage capacity, due to storage space restrictions. Bijvank & Vis (2012b) develop an exact and approximate model to determine the reorder-point *s* and order-up-to-level *S*. This model assumes the lead time to be a multiple of the review period *R* and stochastic demand. Zhang et al. (2014) propose a spreadsheet simulation optimization approach that uses empirical distributions to model demand. The model provides the users with visual intuitive experiences and uses a local search heuristic to find near-optimal results for the (*R, s, S*) policy settings. However, the simulation model cannot cope with factional lead times.

Fixed order size (R, r, Q) policies

Bijvank & Vis (2012a) discuss a fixed order size (*R, r, Q*) policy for the point-of-use locations in a hospital setting. They mention that the order size in such settings usually contains fixed quantities as this system is transparent and easy to understand for hospital staff. In the inventory system, demand is stochastic and excess demand is lost. Furthermore, the lead time is assumed to be a fraction of the review period. They formulate a capacity model with a service level constraint and a service model with a capacity constraint. Kapalka et al. (1999) propose a search procedure for an (*R, s, S*) policy. However, for the (*R, r, Q*) policy, the order-up-to-level is equal to the fixed order size and safety stock (i.e., $S = Q + r$).

2-Bin system

The literature describes a "visual" version of a periodic review (*R, r, Q*) system with two bins as the 2-bin system (Kanet & Wells, 2019; Landry & Beaulieu, 2013; Landry & Beaulieu, 2010). In the system, the first bin contains the working item capacity with capacity *Q*, and the second bin contains the reserve items with capacity $r = DL + ss$, where *DL* is the expected demand during the lead-time period, and *ss* is the desired safety stock. Staff consumes from the working bin until its depleted. This triggers replenishment in the following review period, and the order is expected to arrive *L* periods later. This system has been used in a hospital setting since the late 1980s and has developed since then. An adaption to the system is to have two equal bins to increase simplicity, described as a "2-equal-bin" system or the "2-bin Kanban", where $r = Q$. According to Landry & Beaulieu (2013), the system is, in most situations, superior to other inventory management systems with a periodic review. A 2-bin Kanban system reduces the time taken for the ordering process, as counting stock levels is no longer necessary.

Furthermore, the built-in stock rotation reduces the risk of products expiring, leading to better ergonomics in high-density storage systems (Landry & Beaulieu, 2013). A recent literature review by Lanza-Léon et al. (2021) reflects on the benefits and barriers of applying a Kanban system in healthcare settings. The case study of Persona et al. (2008) shows improvement from an economic and management point of view. However, in the case of "non-stock" products, the system is inefficient as the required quantities exceed the actual need because of the risk of a stockout. In addition, Papalexi et al. (2016) also experience that the effective implementation of a 2-bin Kanban system depends on the products used in the system. The researchers use an ABC classification method in combination with the demand profile to identify suitable products. Kanet & Wells (2019) extend the research on the 2-bin Kanban system by increasing the number of bins. They make use of the Economic Bin Quantity and the Economic Order Quantity in their study to determine the bin quantities and number of bins for a given situation. However, a limitation is that the calculations are based on a continuous review, while many instances without technology improvements follow a periodic review.

Base-stock (R, S) policies

Base-stock policies are a form of order-up-to-level policies. In the policy, the satisfied demand is immediately ordered in the upcoming review period. For such policies, the reorder level $s = S - 1$. Inventory control policies with a base stock are denoted as (R, S) or $(S - 1, S)$ for periodic and continuous review, respectively (Bijvank, 2009). Huh et al. (2008) proposes a heuristic order-up-to level S based on two newsvendor expressions. They assume positive lead times (i.e., 1 to 4) and consider several stochastic demand distributions in their model. Bijvank & Johansen (2012) found that a pure base-stock policy is not optimal for lost-sales inventory systems, and they propose a restricted base-stock policy (RBSP) instead. This policy limits the order size to a maximum amount.

(R, s, c, S) policy

Dellaert & van der Poel (1996) extend the EOQ model to their so-called (R, s, c, S) inventory control policy with Poisson demands. In the model, in every review period R, if the inventory level of an item is lower or equal to s, all other products with the same supplier and an inventory level below the can-order level c are also ordered to refill all the inventory positions to the "Max" inventory level S. They propose a simple inventory rule that calculates the s, c, and S values, with a given R. The model aims to minimize costs, considering the ordering and holding costs. In addition, they study the service rate of the model. However, any storage capacity restrictions are disregarded.

Hybrid policy

Rosales et al. (2014) describe a hybrid policy which combines a low-cost periodic (R, s, S) replenishment policy with a high-cost continuous (r, Q) replenishment option to avoid stock-outs. Every day, at the beginning of the shift, inventory of the point-of-use location is replenished up to level S when the inventory level is lower or equal to s. However, whenever the inventory level reaches a threshold r during the shift, a replenishment with constant size Q is triggered to prevent a stock out. In the system, the value of r is lower than s. They developed a parameter search engine using simulation to optimize the long-run average costs. Their results show that the costs, inventory levels, and the number of replenishments reduce when using the hybrid policy instead of solely a periodic or continuous review policy.

Linking patient schedule

Epstein & Dexter (2000) propose a simulation approach to assess the integration of the supply and surgery schedules. They use simulation to analyse the system in which items are ordered and delivered just-in-time based on planned surgeries. They conclude that it only saves costs for expensive items as the strategy would result in more orders compared with a sophisticated, stand-alone material management inventory control. However, further analyses of this strategy should be conducted to reveal other possible benefits.

3.1.2 Continuous review

This section outlines the available literature on inventory systems with a continuous review period, as summarized in Table 5. The inventory levels are continuously checked, and an order is placed when necessary.

Table 5: An overview of inventory models with a continuous review replenishment policy.

Policy	Studies	Method	Objective	Constraint	Lost sales
(*s*, *S*)	(Kelle et al., 2012)	Approximate	Costs	Storage space and service level	No
	(Archibald, 1981)	Exact and approximate	Costs	-	Yes
(*r*, *Q*)	(Hadley & Whitin, 1963)	Exact and approximate	Costs and service level	-	Yes
	(Johansen & Thorstenson, 1993)	Exact	Costs	-	Yes
2-Bin system with RFID	(Rosales et al., 2014)	Exact	Costs	Inventory balance	No
(*S - 1, S*)	(Smith, 1977)	Approximate	Costs	-	Yes

Variable order size (s, S) policies

Kelle et al. (2012) provide simplified models to set the reorder point and order up to level for an (*s, S*) replenishment system with continuous review. These parameters are based on a near-optimal allocation policy of stock under storage capacity and service level constraints. They performed a case study on a hospital pharmacy's automated ordering system. The research supports the hospital in achieving its main goals: (1) to reduce the number of emergency and daily refills, and (2) to reduce holding costs. However, no assumption of lost sales is included in their models. Archibald (1981) developed an order-up-to-level policy for a continuous review inventory system with the demand that follows a compound Poisson process, assuming a constant lead time.

Fixed order size (r, Q) policies

Bijvank & Vis (2011) performed a literature review on lost-sales inventory models with an (*r, Q*) replenishment policy. They mention that the first work dates to the classical textbook of Hadley & Whitin (1963), in which models with an exact and approximate treatment are developed for a case with lost sales. In the models, at most one order is outstanding, demand is assumed to follow a Poisson distribution, and lead times are discrete. This work is extended by including gamma-distributed lead times by Johansen & Thorstenson (1993). They formulated a semi-Markov decision model to obtain exact solutions for the reorder point r and the order quantity Q.

2-Bin system with RFID

The 2-bin Kanban system is further improved with the introduction of RFID technology (Çakici et al., 2011; Landry & Beaulieu, 2013). This technology enables the policy to change from a periodic review to a continuous review inventory control policy (*r, Q)*. The 2-bin system with RFID is either on an item or bin level. When a company uses this technology on an item level, products have a radio transponder called "tags" attached to them. A tag that is within range of a reader communicates with the systems such that the system knows how many items are stored at every location. However, it can be prohibitively expensive for low-cost items (Johnson, 2002). On a bin level, the tag is attached to the label of a bin. When a bin empties, staff must place the label on a replenishment board connected to the system. Hospital management knows the status of the stocks when the labels are removed from the bin. The increased visibility alerts them to potential stock-outs (Landry & Beaulieu, 2010). To improve inventory management for a continuous-review 2-bin system, Rosales et al. (2014) present a semi-Markov decision model, which determines the optimal number of empty bins for replenishment.

Base-stock (S - 1, S) policies

In this policy, every demand is immediately reordered to the base-stock level *S*. The study of Smith (1977) presents an approximation for finding the optimal order-up-to levels *S* for an inventory system with Poisson demands, arbitrary replacement time distribution, and emergency handling costs for lost sales.

3.2 Modelling methods

The previous section outlines some of the available literature on different replenishment policies. The various replenishment policies can be formulated as optimization problems with one or multiple objectives. These objectives need to be minimized or maximized, with usually some constraints. Solution methods obtain the optimal or near-optimal solution to solve the optimization problem. In this section, we outline the modelling methods mentioned in Section 3.1 with their corresponding benefits and classify them into three categories. These categories are (1) exact methods, (2) non-exact methods, and (3) simulation.

3.2.1 Exact methods

Exact methods use standard processes that use mathematical principles to achieve optimal solutions (Saha & Ray, 2019). Many modelling approaches have been developed by researchers, which may be classified as deterministic, stochastic, and distribution-free approaches. Inventory problems under certainty are modelled as deterministic models. An EOQ-based method and mathematical programs are used to formulate these models. A mathematical program consists of a mathematical structure in which the variables represent the problem choices. These decision variables are set to minimize or maximize the objective function and could be restricted by constraints. A mathematical program seeks to find the optimal solution given a defined domain or input, also known as the global optimum (Grond, 2016). A stochastic approach is applied to formulate a model that includes uncertainty with complete knowledge of probability distribution. It is possible to use a newsvendor model in addition to mathematical programs and a hybrid-EOQ-based model to formulate these models. Models without complete knowledge of probability are formulated using distribution-free analysis approaches (Saha & Ray, 2019). A disadvantage of exact methods is that they do not scale well to large problems, taking a considerable amount of processing time to find the optimal solution.

3.2.2 Non-exact methods

Non-exact methods are formulated to solve optimization problems. When exact methods have long processing times, which are not desirable in practice, they may require simplifications and approximation procedures to determine the solution more efficiently (Bijvank & Vis, 2012b). Two forms of non-exact methods are approximation algorithms and heuristics. Approximation algorithms are simplifications of an exact method, which guarantee to find a near-optimal solution. Heuristic algorithms are designed to find, generate, or select a search algorithm that may provide a sufficiently good solution within a 'reasonable' amount of time (Grond, 2016). They typically handle problems with incomplete or imperfect information or limited computational time better than exact algorithms or mathematical programs. However, the main disadvantage of heuristics is that they cannot guarantee the solution's optimality (Grond, 2016).

3.2.3 Simulation

Robinson (2014) defines simulation as "experimentation with a simplified imitation (on a computer) of an operations system as it progresses through time, for better understanding and/or improving that system". Usually, simulation involves the form of a set of assumptions about the system's operation, expressed as mathematical relations between the objects of interest in the system. Simulation is used when mathematical optimization is too complex or when real-world experiments are too expensive or time-consuming. Furthermore, it provides the users with visual intuitive experiences to overcome resistance from managers who want to deploy solutions that their teams can easily learn, accept, and use (Zhang et al., 2014). A disadvantage of a simulation model is that it is not an optimization method (Winston & Goldberg, 2004). However, it is possible to implement heuristics into it to find solutions, and simulation makes it possible to evaluate different replenishment policies and scenarios.

3.3 Conclusion

Different inventory replenishment policies exist in literature to improve the service level, decrease the costs, or both. Literature separates these policies into policies with a periodic or continuous review. Currently, Isala uses a periodic review system to control the stock and purchase items. We consider a different replenishment approach for stock and purchase items as these item types vary in their mean usage and lead time.

Stock items have a constant review period and a short lead time. The lead time for stock items is, most of the time, less than a day as they are supplied multiple times per day from the external warehouse by Isala's logistical partner. The mean usage of these items is higher than the usage of purchase items, and therefore, we use a 2-bin Kanban policy with equal bin sizes, denoted as an (R, r, Q) system, where the reorder point (r) equals the order quantity (Q). This system decreases the workload of employees and the built-in stock rotation reduces the risk of products expiring, leading to better ergonomics.

Purchase items have a constant review period and stochastic lead time. The items can be classified as slow movers due to their relatively low daily usage. To use the available capacity efficiently, we want to formulate an order-up-to-level replenishment policy with a variable order size, denoted as an (R, s, S) system.

Several modelling methods can be used to formulate the optimization problem. However, determining optimal parameter values for a 2-bin Kanban and (R, s, S) replenishment policy is complex and time-consuming, making it less applicable for practice. Furthermore, the model should be appealing to use and easy to understand, such that managers at Isala can implement the model in practice. Therefore, we want to create a simulation model with a search heuristic that finds the near-optimal solution for all items. The simulation model will be similar to the model of Zhang et al. (2014), however, we would need to adjust it such that it can cope with lead times that are an integral multiple of the review period and factional. The heuristic will be similar to the algorithm of Kapalka et al. (1999) and Esmaili et al. (2019), which they also developed for an (R, s, S) policy. In addition, the heuristic we develop is also suitable for the 2-bin Kanban model. With the model, we can evaluate and optimize both replenishment policies with a variety of parameter options.

4. Model formulation

In Chapter 3, we decided to use a 2-bin Kanban system for stock items and an order-up-to-level replenishment policy for purchase items. Furthermore, we concluded that a simulation model with a search procedure is the best approach for solving the optimization problem. This chapter outlines the mathematical formulation of these systems and the solution approach, divided into four sections. First, we outline the solution design in Section 4.1. Moreover, it provides an overview of the steps of the new replenishment method and the model assumptions. Then, we describe the objective and constraints of the search heuristic in Section 4.2. In Section 4.3, we outline the local search algorithm and the reasoning behind its use. Section 4.4 presents the outputs from the model and the key performance indicators to analyse these outputs. In Section 4.5, we outline the assumptions and simplifications of the model. Section 4.6 describes the verification and validation of the model. We conclude this chapter in Section 4.7.

4.1 Solution design

The problem involves a single-item inventory control under periodic review R, which is the same for stock and purchase items. We use simulation in combination with a search heuristic to find the values of the reorder parameter r (s) and the order quantity parameter Q (order-up-to-level parameter S) that minimize the costs of stock (purchase) items for a period J. This section starts with outlining the five steps of the simulation model. Figure 9 shows a visualisation of these steps.

Step 1: Initialize simulation model

In the first step, we initialize the parameters of the simulation model and the search heuristic. Section 4.3 describes the initialization of the parameters starting reorder point and order quantity (or order-up-to-level), the minimum reorder point and maximum order quantity (or order-up-to-level), and the package size, i.e., the algorithm's decreasing/increasing step size. Section 5.3 describes how we use a warm-up period to set the initial inventory position. Furthermore, we set the service level per item for medical items to the desired service level of 0.999 and non-medical items to 0.98 (see Section 2.4).

Step 2: Determine demand

In the second step, the simulation loops over the days (j = 1 to J) to simulate how the inventory behaves for a given policy setting. Figure 10 shows an overview of this simulation process. The remainder of this step description explains the process in more detail. Every day starts with determining the expected demand, which takes the order data of the item as input (see Chapter 5). After that, we calculate the start on-hand inventory and the inventory position. The on-hand inventory position is equal to the products which are physically in stock. The inventory position is the on-hand inventory position plus the outstanding orders. Then, we determine whether we should replenish the item. An item should be replenished when the current day is a review period and the inventory position is equal to or below the reorder point. If we do not need to order the item, we determine if there was a stockout and calculate the shortage, i.e., the number of products that were not on stock. Otherwise, we determine the order quantity and the demand during lead time first. Next, we calculate the end on-hand inventory, which is always equal to or greater than zero, as we assume that the demand during a stockout is lost, or equivalently, satisfied from an alternative source. The day ends with calculating the day's realized demand, which is the expected demand minus the shortage.

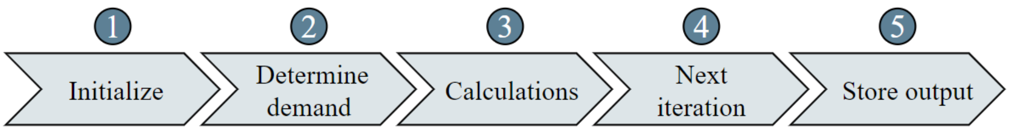

Figure 9: Overview of simulation model steps.

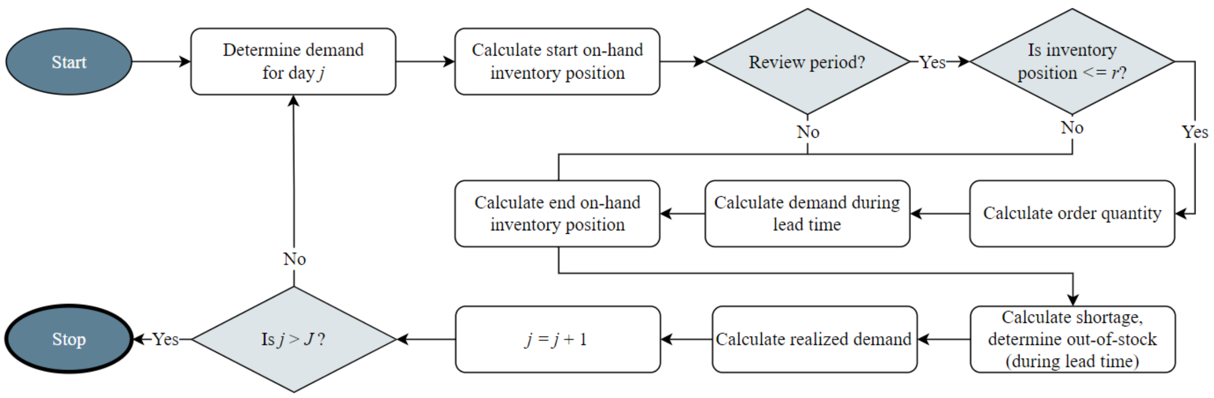

Figure 10: Inventory system simulation process.

Step 3: Calculate the service level and costs

After the simulation model has run over the days, it calculates the service level and the costs for a given reorder point r (s) and order quantity Q (order-up-to-level S) setting. Section 4.2 describes how these are calculated. Based on these calculations, the simulation models continues to the next iteration.

Step 4: Next iteration

There are three possibilities on how the simulation model continues when it is in the next iteration step, based on the decision of the search algorithm. These options are:

1. For a particular item i, the parameter setting changes. Depending on the current parameter value and the corresponding service level and costs, both the reorder point r and order quantity Q either increase or decrease in the case of the 2-bin Kanban System. Based on the (R, s, S) policy, the reorder point s and the order-up-to-level S can also increase or decrease separately. Return to step 2.
2. The simulation model continues with the next item $i + 1$. This happens when, for item i, the search heuristic goes beyond the lower or upper bound of the search area. In addition, for particular scenarios, the search algorithm can stop and continue to the next item when the costs do not improve or the service level is below the threshold (see Section 4.3). Return to step 2.
3. Go to step 5: 'Store output'. When the search algorithm has found the best settings for every item in the list, it terminates, and the simulation model continues to the next step.

Section 4.3 provides a more comprehensive description about the reasoning and flow of the search algorithm.

Step 5: Store output

After the simulation has looped over all items on the list, the output data of all these items are stored. The output data consist of:

- The reorder point r (s) and order quantity Q (order-up-to-level S) setting for every stock (purchase) item.
- The costs and service level of the corresponding setting.
- The number of stockouts, number of orders, and average inventory level of every item.
- The shortage and fill rate of all items.

Once we have the output data of every item, we create an overview of these data in Excel.

4.2 Objectives and constraints

This section outlines the objective function of both proposed replenishment policies. First, we formulate the objective function and constraints of the 2-bin Kanban system. After that, we describe the objective function and constraints of the (*R, s, S*) policy, which has only minor modifications to the 2-bin Kanban model.

4.2.1 2-Bin Kanban

In the 2-bin Kanban system, an order is placed with a fixed size Q (equal to *r*) when the inventory level in a period falls below the reorder point *r*. In this system, the costs are a combination of ordering, holding, and shortage costs per product *i*, where *i* = 1 to *I*.

Ordering costs

The ordering costs consist of a fixed order cost, *A*, and the cost per unit that is ordered, *c*. The ordering cost only incur when the current day is a review period and the on-hand inventory position, *X*, is less than or equal to the reorder point *r*. We define the ordering costs as:

$$O\{X_{ij} \leq r_i\}(A_i + c_i Q_i), \tag{1}$$

where,

$O\{\cdot\}$ = indicator function of set {·},

X_{ij} = inventory position of product *i* in period *j*,

r_i = reorder point of product *i*,

A_i = fixed order costs of product *i*,

c_i = per-unit ordering cost of product *i*,

Q_i = order quantity of product *i*.

Holding costs

The holding costs consist of the holding cost rate, *h*, times the cost per unit, *c* (same as for the ordering costs Equation 1), multiplied by the on-hand inventory at the start of period *j*. We define the holdings costs as:

$$h_i c_i X_{ij}^{+}, \tag{2}$$

where,

h_i = holding cost rate of product *i*,

c_i = per-unit ordering cost of product *i*,

X_{ij}^{+} = max(*x*, 0), on-hand inventory of product *i* at the start of period *j*.

Shortage costs

The shortage costs include the fixed emergency order costs, *K*, and the costs per lost demand. The cost per lost demand is equal to the shortage of the item, X^{-}, multiplied by the cost per unit *c* (same as for the ordering costs Equation 1). This only occurs when there is a stockout in period *j*. We define the shortage costs as:

$P\{X_{ij}^{-} < 0\}(K_i + |c_i X_{ij}^{-}|)$, (3)

where,

$P\{\cdot\}$ = the indicator function of set $\{\cdot\}$,

X_{ij}^{-} = min(x, 0), shortage of product i at the end of period j,

K_i = the emergency order costs for product i,

c_i = the per-unit ordering cost of product i.

Objective function

Using the cost mentioned above and the decision variables, r and Q, we formulate the objective function (Equation 4), which aims at minimizing the costs for the 2-bin Kanban systems:

$$min \sum_{j=1}^{J} [O\{X_{ij} \leq r_i\}(A_i + c_i Q_i) + \ h_i c_i X_{ij}^{+} + \ P\{X_{ij}^{-} < 0\}(K_i + |c_i X_{ij}^{-}|)] \quad (4)$$

This objective is subject to the following constraints:

$\alpha_i \geq \ \bar{\alpha}_i$, (5)

$r_i = \ Q_i$, (6)

$r_i, Q_i \geq 0$ integer. (7)

Even though shortage costs are involved in the objective function, the solution is constrained to meet at least the desired service level $\bar{\alpha}$, which is defined as the probability of not going out of stock in a review period (Schneider, 1978). Whenever there is a stockout, the articles are either derived or taken from another storage location or arrive via an emergency order. In both cases, it disrupts the daily process in the hospital, which should be avoided. Therefore, we use service level α, which is calculated as follows:

$$\alpha_i = 1 - \left[\frac{1}{J}\sum_{j=1}^{J} Stockout_i\right] \quad (8)$$

4.2.2 (*R*, *s*, *S*) policy

For the system with an (*R*, *s*, *S*) policy, an order with a variable size (equal to *S* - *X*) is placed when the inventory level in a period falls below the reorder level s. Also, in this system, the costs are again a combination of ordering, holding, and shortage costs. For product *i*, where *i* = 1 to *I*, the holding and shortage costs are equal to these costs for the 2-bin Kanban system with Equations 2 and 3, respectively. The ordering costs are defined as:

Ordering costs

The ordering costs are again composed of a fixed and per unit costs (see Equation 1). However, the order size, *S* – *X*, varies per period *j*. Thus, we define the ordering costs as:

$O\{X_{ij} \leq s_i\}(A_i + c_i(S_i - X_{ij}))$, (9)

where,

$O\{\cdot\}$ = indicator function of set $\{\cdot\}$,

X_{ij} = inventory level of product i in period j,

s_i = reorder level of product i,

A_i = fixed order costs of product i,

c_i = per-unit ordering cost of product i,

S_i = order-up-to level of product i.

Objective function

Using these costs and the decision variables, s and S, we formulate the objective function (Equation 10), which minimizes the costs for the (R, s, S) policy:

$$min \sum_{j=1}^{J}[O\{X_{ij} \leq s_i\}(A_i + c_i(S_i - X_{ij})) + \ h_i c_i X_{ij}^{+} + \ P\{X_{ij}^{-} < 0\}(K_i + |c_i X_{ij}^{-}|)] \tag{10}$$

This objective is subject to the following constraints:

$$\alpha_i \geq \ \bar{\alpha}_i, \tag{11}$$

$$s_i \leq \ S_i - 1, \tag{12}$$

$$s_i, S_i \geq 0 \text{ integer.} \tag{13}$$

We use Equation 8 to calculate the service level α for product i, which is the same as for the 2-bin Kanban system (see Section 4.2.1).

4.3 Search Algorithm

This section outlines the search heuristic of our simulation optimization approach. In this study, we formulate a search algorithm based on the monotone search algorithm (MSA) proposed by Kapalka et al. (1999) and the recursive algorithm of Esmaili et al. (2019). We simplified the algorithm of Kapalka et al. (1999) by reducing the number of steps, while we extended the algorithm of Esmaili et al. (2019) by also considering the cost. The algorithm is described as a local search heuristic that iteratively determines the best reorder point and order-up-to-level (or order quantity) for a given article. The benefit of using such algorithms is that it assumes monotonicity to reduce computation time. More specifically, for a fixed S, "*if the service level of (s, S) is lower than the required service level, there is no need to check policies (s-1, S), (s-2, S), …, because they will fail to provide the required level of service*" (Kapalka et al., 1999). They recognise the same pattern for the costs with a fixed S, which always increases when s increases. However, for our search algorithm, we can only assume monotonicity for the service level since the shortage costs disrupt this for the costs. We formulate two variants of the algorithm, as the 2-bin Kanban system and (R, s, S) system have varying characteristics concerning the order parameters. The upcoming sections describe the algorithms for each of the systems in more detail.

4.3.1 2-Bin Kanban system

The algorithm search procedure for the 2-bin Kanban system (stock items) is visualised in Figure 11. The procedure starts by initializing the parameters. The upper bound (*Qmax)* of the search area is the maximum inventory position given by $UB = m * (\text{mean order size} + 2)$, where m is a number chosen such that $UB > \text{max order size}$. Furthermore, we multiply it with a value which is two units more than the mean order size to account for low mean order sizes. Kapalka et al. (1999) propose to use a multiple of the EOQ + Safety Stock for the upper bound. However, we always increase or decrease both r and Q. As a results, the algorithm reaches the upper bound faster than in the case of Kapalka et al. (1999). Therefore, we propose to use the mean order size instead. We found that $m = 6$ is more than adequate in a wide range of test situations based on the order data. The lower bound is the minimum reorder point (*rmin*), which is 0. When the parameters are set, the tool starts searching from the starting point for the first setting that satisfies the service level constraint. The algorithm either reaches the upper bound and stops, or it finds a setting and continues the search to find a solution with lower costs. When the algorithm continues its search, it decreases the reorder point and order quantity by one unit until the solution no longer satisfies the service level constraint. Some products are only allowed to be ordered in fixed-packed sizes (batch sizes) of, for instance 10 units. In this case, the reorder point and order quantity would both decrease by 10 units. During this process, the setting with the least amount of costs is stored. Since we assumed monotonicity there is no need to check policies ($r-1, Q-1$), ($r-2, Q-2$), …, once the service level of (r, Q) is lower than the required service level because they will fail to satisfy the desired service level as well. To reduce the execution time, we want to start at a point which is near the optimal solution. Therefore, a good starting point for the 2-bin Kanban search algorithm is r and Q equal to the Economic Bin Quantity (EBQ), which is the optimal bin quantity for a 2-bin Kanban system unless its value is lower than the summation of the demand during lead time and the safety stock (Kanet & Wells, 2019). The EBQ is equal to $\frac{EOQ}{\sqrt{3}}$.

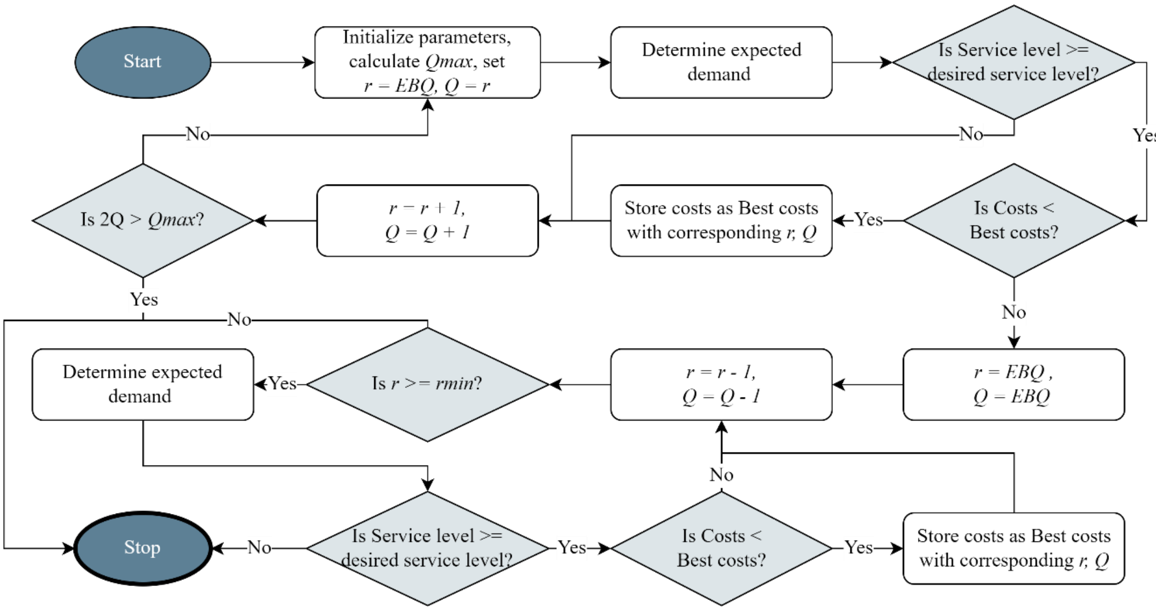

Figure 11: The flowchart of the 2-bin Kanban search algorithm.

4.3.2 (*R, s, S*) policy

Figure 12 shows the algorithm search procedure for the purchase articles with an (R, s, S) policy. For this policy, the algorithm first initializes its parameters. Even though the search procedure is different for the 2-bin Kanban system, we found that the same procedures apply to the lower bound (*smin*) and upper bound (*Smax*) of the (R, s, S) system, based on test situations with the data set. In these situations, the algorithm found a solution for all the items. Then, the tool starts its search at the minimum reorder value $s = 0$ and an order-up-to-level $S = s + 1$. Next, it iterates over four steps to find the near-optimal setting: In the first step, both s and S increase by one unit each with every iteration until the tool either reaches the upper bound and stops or finds an (s, S) setting that satisfies the service level constraint and has higher costs than the previous setting with the lowest costs. The current order parameters setting is stored, and we return to the setting with the best costs ($s = s-1, S = S-1$). In Step 2, the algorithm decreases the reorder point s while S remains unchanged to find the lowest reorder point that satisfies the service level constraint for that specific order-up-to-level. When the reorder point is lower than the minimum reorder point or the service level constraint cannot be satisfied, the algorithm returns to the stored setting (Step 3). From there, the tool does the procedure of Step 2 until it does not satisfy one of the constraints anymore. In the last step, it increases the order-up-to-level S while s stays unchanged until it reaches the maximum order-up-to-level value or the order-up-to-level S is larger than the summation of the reorder point and the EOQ.

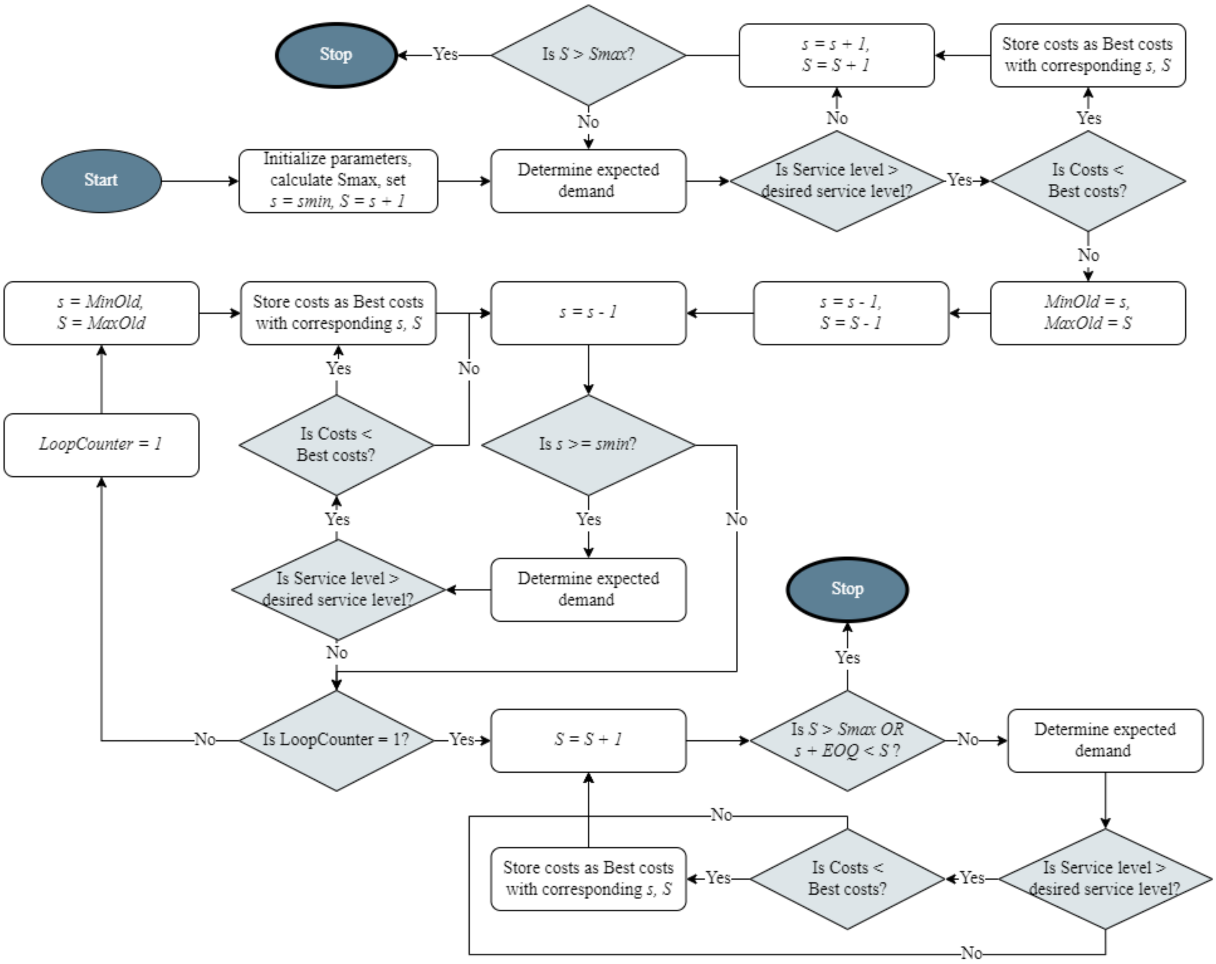

Figure 12: The flowchart of the (R, s, S) search algorithm.

4.4 Model outputs

We use two main KPIs to measure the performance of the various interventions. These KPIs are:

Costs

The costs are a combination of the holding, ordering, and shortage costs. A more detailed explanation of the costs calculations is provided in Section 4.2. The costs KPI does not apply to the model objective in which we only consider the service level constraint.

Service level

The service level is defined as the availability of an item, which is the probability of not going out of stock in a review period.

Other outputs from the model are the number of orders, average inventory level, fill rate, and the number of stockouts. These outputs have less impact on decision-making. However, they can still support certain decision-making as it provides the user with more information about the scenarios.

4.5 Assumptions and simplifications

Various assumptions and simplifications are made to model the inventory of the OR department. This section provides an overview of them.

1. The lead time of stock items that are delivered by HL is assumed to be constant based on the agreed order and delivery times between Isala and HL.

2. The lead time of stock items that are supplied by other suppliers and purchase articles is assumed to be between three to five days, which is justified by subject-matter-experts (SMEs) from Isala.
3. Lead time demand is the fraction of the day demand during lead time. This is only calculated for stock items, as purchase items have a lead time of at least one day. Therefore, the lead time demand for purchase articles would be equal to the daily demand times of the lead time days.
4. The end inventory of a day is always equal to or greater than zero. We assume that demand that is not realized via the original location, is either fulfilled via an emergency order or taken from another storage location.
5. We assume that the inventory position at a review day is reviewed at the beginning of the day.
6. Only workdays with a total of ten hours are considered in the simulation. Since, most of the operations are scheduled on those days. Also, it is not possible to place regular orders outside office hours.
7. There are no backorders at suppliers. Therefore, it is always possible to order the necessary items.
8. We assume that the supplier always delivers the correct items.

4.6 Verification and validation

According to Robinson (2014), it is not possible to prove that a model is valid. Therefore, model verification and validation are concerned with creating enough confidence in the model so it can be used for decision-making. We define verification as correctly translating the conceptual model into a computer model. Whereas validation concerns whether the simulation model is an accurate representation of reality for the particular objectives of the study (Law, 2015).

To verify the model, we applied seven out of the eight verification techniques of Law (2015). These applied techniques are:

1. Write and debug the computer program in modules or subprograms. To maintain an overview of the various processes, the model consists of multiple modules with multiple subprograms and functions.
2. More than one person should review the model. The model has been verified by multiple experts of Isala and by an independent researcher.
3. Simulate a variety of settings. We run the model for 1,754 items that all have different characteristics. Furthermore, we vary parameters in the experiments (see Section 5.2).
4. Use a “trace” to debug the program. We traced the steps of the simulation model by visualizing the inventory process, demand, and results on a sheet (see Appendix A). Furthermore, we used the debugger during the construction of the model.
5. Run the model under simplifying assumptions. At the start, we mainly ran the model for a few items and used constant parameters. Later we added more variability into the model.
6. Observe animations of the simulation output. By observing the visualization of the inventory position and process of items, possible errors were detected and redressed.
7. Compute the sample mean and sample for each simulation input probability distribution, and compare them with the historical mean and variance. We did this partly as only the historical mean of the demand is available for comparison.

The technique that we did not apply in the verification of the model is using a commercial simulation package to reduce the amount of programming. However, we decided to use Excel which is not a simulation package. However, the program is broadly used at Isala, which increases useability. Even though, we did not use all verification techniques, we believe that the conceptual model is correctly translated into the computer program.

To validate the model, we cannot compare the model outputs with the current situation of the real system as the data records do not concern actual demand for items, and Isala has no information on the current service levels (see Section 2.3 and Section 2.4). Instead, we did two other steps to validate the model.

First, we showed the simulation model, its process, and the results to several SMEs at Isala for validation. Once they believed that the simulation model represent reality sufficiently well, we continued with running experiments (Kleijnen, 1998). The aim of running various experiments is to detect the effect of changing a particular parameter (see Section 5.1.2). With a sensitivity analysis we want to derive conclusions about the importance of the parameters.

4.7 Conclusion

Our simulation model consists of five steps to simulate the inventory system of Isala and to find near optimal or optimal order parameters for all items. In this simulation, the order parameters are the decision variables. First, we initialize the simulation model and set the item's order parameter setting. Then, we determine the demand for a certain item i. Next, we calculate the costs and service level for the item. Based on these calculations, we return to step 2 with another order parameter setting for the same product or with the order parameters of the next item $i + 1$. We store the outputs of the model when there are no items left to simulate.

We measure the performance of the model by two KPIs, which are the costs and availability of an item for a certain order parameter setting. In addition, we store the number of orders, number of stockouts, average inventory, and fill rate to support certain decision-making as it provides the user with more information about the scenarios. To construct the model, we made several assumptions and simplifications. We verified these assumptions, and the overall model by using seven out of the eight verification techniques of Law (2015). We validated the model by first showing it to experts of Isala and then ran various experiments with it.

5. Experimental design

This chapter outlines the experimental design of the simulation model, which we described in Chapter 4. This chapter is divided into four sections. Section 5.1 describes the input demand distributions of the model. These include the demand input and various parameter options for the holding cost rate, ordering costs, and demand variability. Section 5.2 outlines the various experimental parameters and the experimental scenarios. Then, we present the experimental setup in Section 5.3. We conclude this chapter in Section 5.4.

5.1 Input demand distributions

We have to specify the demand probability of the items to carry out the simulation using random inputs. There are a variety of distributions that can be used in the simulation model. These can be split into three main types, presented in Table 6.

Table 6: Types of distributions, adapted from Robinson (2014).

Distribution	Description
Continuous distributions	For sampling data that can take any value across a range
Discrete distributions	For sampling data that can take only specific values across a range
Approximate distributions	Used in the absence of data

It is preferred to use a continuous distribution over a discrete distribution, whereas it is preferred to use a discrete distribution over an approximate distribution (Law, 2015). However, it is not possible to fit a continuous distribution in every situation. We need to identify the category of the collected data to select which distribution we use in our model.

- *Category A data* is available. This category of data is collected previously or is known. For this data it is important that it is accurate and in the right format for the simulation model.
- *Category B data* is not available but collectable. To collect the data, it might require conducting interviews with experts or observing the real-world system.
- *Category C data* is not available and not collectable. It can be too time-consuming to observe a real-world system or the system is not operational.

An observation in our data analysis is that the data records concern order quantities and not the actual demand for an item, which is common for hospitals in literature (Bijvank, 2009). As a result, the available data about the usage of an article is limited. Furthermore, we cannot collect the usage data as it would take too much time to count the usage of all the items for a given period. Therefore, we classify it as category C data.

There are two main ways of dealing with category C data. We can either estimate the data or treat the data as an experimental factor rather than a fixed parameter (Robinson, 2014). The unobtainable data at Isala is the exact usage of an item, further referred to as the demand. However, we can estimate the average demand of an item by taking its mean order quantity per working day. Based on the average demand, we can split the items into two categories: slow-moving and high demand items.

Slow-moving products have an average daily demand lower than 10 units. For these items, it is important to be able to deal with discrete units. A Poisson distribution is a discrete distribution which is appropriate to use for slow-moving items when the observed standard deviation is within 10% of the square root of the average daily demand (Silver et al., 2017). We cannot confirm this statement as the observed standard deviation is part of the unobtainable data. However, literature support the assumption that demand of items in a hospital environment follows a Poisson distribution (Epstein & Dexter, 2000; Bijvank & Vis, 2012a). Therefore, we make this assumption for slow-moving items.

High demand products have an average daily demand of at least 10 units. For these items, Silver et al. (2017) suggest using a normal distribution to simulate the demand, which Dellaert & van de Poel (1996) also assume in their research. A drawback of this continuous distribution is the possibility of generating a random variable with a negative value when the variance is relatively large. However, we cannot collect information about the variance of the items demand. Instead of using a fixed variance, we use experimental factors for the variance to mean (VTM) ratio, similar to Bijvank & Vis (2012b). We determine the variance by multiplying the average demand with the VTM ratio. It is important that the VTM value does not result in a variance which is to large. Otherwise, we should truncate the normal distribution such that the generated value is always non-negative. Another approach to simulate the daily demand of high-demand products to overcome the absence of data is to use approximate distributions. The simplest form of approximate distribution is the uniform distribution. We can use a similar approach for this distribution as for the normal distribution, where we use a multiple of the mean to set the *a* and *b* boundaries, which are the minimum and maximum values. The normal distribution and uniform distribution are both used in separate scenarios to simulate the demand for high-demand items.

5.2 Experiments

This section first outlines the various parameters we use in the experiments. Next, we provide an overview of the scenarios and their corresponding parameter values.

5.2.1 Experimental parameters

There are two main reasons for designing various experiments for the simulation model. First, we want to acquire insight into the performance of the search heuristic with its corresponding replenishment policies in the base situation and compare the performance to the current order parameters. We run this experiment for a model with the cost objective that is explained in Section 4.2 and with a service-level model. The service-level model determines the order parameter settings that take up the least amount of storage capacity while satisfying the service-level constraint. The information about the model performances helps Isala by determining the order parameter settings for their items. Table 7 shows the parameters of the base situation. Second, we perform a "what-if" analysis on two parameters, lead time and demand variability, to observe their impact on the order parameters. In an experiment, we vary a certain parameter while the other parameters remain fixed (ceteris paribus) and compare the cost objective outcomes. In each scenario, we consider 1,754 items, and we use the normal or uniform distribution and Poisson distribution for high-demand items and slow-moving items, respectively. The parameters that vary are the holding cost rate, *h*, fixed ordering costs, *A*, lead time of suppliers other than HL, and the demand variability of high-demand items. The fixed ordering costs for products are different for the supplier HL, *A (HL)*, and other suppliers, *A (Other)*. By varying the holding cost rate and the fixed ordering costs we want to see if the simulation model can cope with various parameters. We vary the lead time to gain insight into supplying an item via HL instead of another supplier. Demand variability is a topic that gets more attention from the experts from Isala. However, the impact is not clear. We want to get more insight into this by varying the variability of high-demand items. These parameters vary as follows:

h	{0.05,	0.30,	0.60,	1}
A (HL)	{0,	1,	5,	10}
A (Other)	{0,	5,	22.50,	60}
Lead time (Other)	{0.5 or 0.8,	Uniform (3, 5)}		
VTM (Normal)	{0.5,	1,	2}	
a, b (Uniform)	{(0.7, 1.3),	(0.5, 1.5),	(0.2, 1.8)}	

Table 7: Base situation parameters.

Category	Parameter value	
Holding cost rate	*h* 0.30	
Fixed ordering costs	*A (HL)* 1	*A (Other suppliers)* 22.50
Stockout costs	*K* 50	
Demand variability	*VTM (Normal) high-demand items* 2	*VTM (Poisson) slow-moving products* 1
Lead time	*HL* 0.5 or 0.8	*Other suppliers* U (3, 5)

Holding cost rate and fixed ordering costs (HL)

We vary the holding cost rate, *h*, and fixed ordering costs of HL, *A (HL)*, parameters similar to the factors that Neve & Schmidt (2022) used in their experiments. In the base scenario, we use a holding cost rate of 30% (per year). This number is similar to the holding cost rate Dellaert & van de Poel (1996) used in their research for storage in the sterile department (29.5%). Furthermore, in the base scenario, we assume that the fixed ordering costs of HL are €1.

Fixed ordering costs (Other)

The factors for the fixed ordering costs of other suppliers, *A (Other)*, are based on available data, with zero costs being the minimum, an average of €22.50, and a maximum of €60. Furthermore, we added a fixed ordering cost of €5 for convenience regarding the calculations. With the experiments, we want to get insight if these costs have a noticeable impact on the order parameters of purchase items.

Stockout costs

The stockout costs, *K*, remain constant in the experiments. An expert of Isala approximates that the cost of an emergency order would be around €225. However, not every stock out triggers an emergency order, as items can also be taken from another storage department. Nevertheless, employees of Isala always spend time handling a stockout. Therefore, we set the stockout costs, *K*, to €50.

Lead time

Based on expert suggestions, the lead time (*Other*) is between three and five working days. As there is no clear indication of the probability per lead time, we assume that it is uniformly distributed on the interval (3, 5). In consultation with HL is it also possible to store purchase items in their warehouse. The items would then be classified as stock items instead of purchase items with the main advantage that they would have a shorter lead time. Items that are delivered by HL are scanned at a fixed interval time window, which is location dependent. The items are also delivered at a fixed time. Therefore, we can assume that all the items that are delivered by HL have a constant lead time, which is either 0.5 or 0.8 working days, depending on the location the item is stored in. Isala is curious about the effect of supplying the purchase items via HL on the order parameters.

Demand variability

We set the parameters for the VTM of high-demand items similarly to the parameters Bijvank & Vis (2012b) used in their experiments. In the current situation, the variability of demand is relatively high. Isala is curious about the effect of less variability on demand on the order parameters. Furthermore, we want to gain insight into the effect when demand is uniformly distributed. We also use boundaries that result in even higher variability of demand. These boundaries are set in collaboration with an expert from Isala.

The demand for slow-moving products follows a Poisson distribution. The VTM ratio is always 1 for a Poisson demand distribution.

5.2.2 Overview experiments

We categorize the experiments into four categories: model verification, base scenario, lead time, and demand variability. The experimental parameters in the model verification category are the holding cost rate and the fixed ordering costs of HL and other suppliers. We used this category mainly to verify our model and did not include them in Chapter 6. However, we can conclude that our model can handle various parameters and that the fixed ordering costs of other suppliers do not have a noticeable impact on the order parameters of purchase items. We do outline the outputs of the other three categories in Chapter 6. Table 8 shows a summary of all experiments, which includes the varying parameter, the number of experiments, and a description. We perform a total of eleven experiments.

Table 8: Summary of the experiments.

Category	Parameter	Experiments	Values	Description
1. Base scenario	-	3	From Table 7	We compare the outputs of the current order parameters and the parameters chosen by the search heuristic for a cost and service-level objective.
2. Lead time (purchase articles)	*Lead time (Other)*	2	{Lead time HL, U(3, 5)}	Run the experiments on the model with cost objective and gain insight into the items with lowered order parameters.
3. Demand variability	*VTM* *U(a, b)*	6	{0.5, 1, 2} {(0.7, 1.3), (0.5, 1.5), (0.2, 1.8)}	Run the experiments on the model with cost objective and gain insight into the change in the order parameters.

5.3 Experimental setup

In this section, we determine the initial conditions, the run length, and the number of replications. We consider two model outputs, costs and average inventory, to determine those numbers for the base scenario. To save time, we take the same experimental setup for all experiments, instead of determining them separately.

5.3.1 Initial conditions

We can set the initial conditions in two ways. We can run a warm-up period for the model or take initial conditions from the real system (Winston & Goldberg, 2004). As there is no accurate data available about the usage and inventory levels of items, we determine the initial conditions by running a warm-up period. With a warm-up period, we delete the observations from the beginning of the simulation. These observations depend on initial conditions and therefore do not representative of steady-state behaviour.

We use Welch's graphical method to determine the length of the warm-up period. With this method, we estimate when the system reaches a steady state performance by visually inspecting the time-series data (Law, 2015). We make five independent replications with a run length of 520 days. For each i^{th} observation, we calculate the mean. Then we use a moving average of 5 (w=5), for which the plot is reasonably smooth. Figure 13 shows the plots of the moving averages of the costs and the average inventory. From this figure, we estimate that the output seems stable after 50 days. Therefore, we decide to use a warm-up period of 50 days.

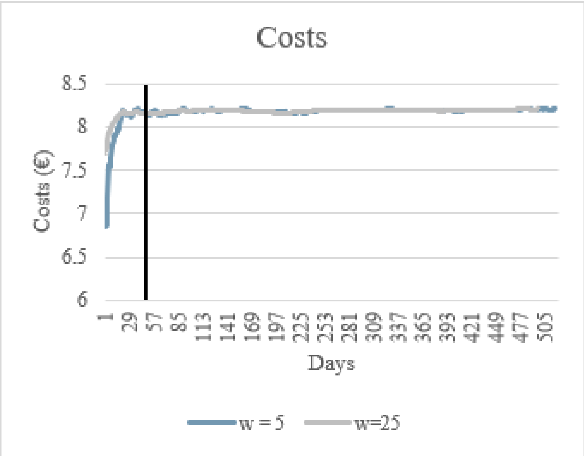

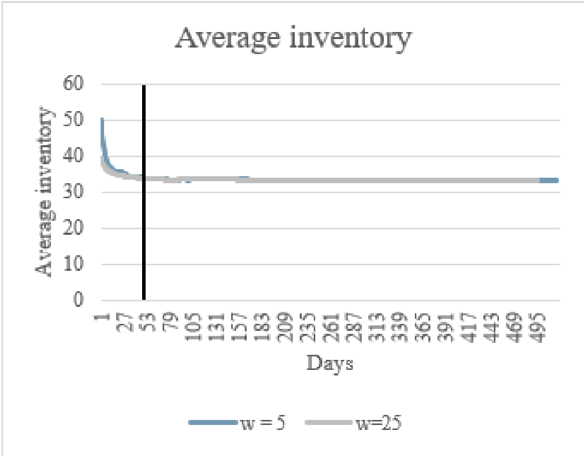

Figure 13: Determining the warm-up period for two model outputs, using Welch's graphical method.

5.3.2 Run length and number of replications

The model should run much longer than the length of the warm-up period to obtain sufficient data (Law, 2015). We decided to run the model for 260 days, which is one year in working days. The total run length of one replication of the simulation would then be 310 days including the warm-up period.

With the sequential procedure, we decide how many replications provide us with an acceptable estimate of the simulation models mean performance (Law, 2015). With a relative error $\gamma = 0.05$, and a significance level $\alpha = 0.05$, the maximum number of replication necessary for our simulation model is three. Therefore, we use three replications in our simulation. Appendix B displays the sequential procedure for the cost and average inventory outputs.

5.4 Conclusion

We designed various experiments to determine the performance of the simulation model. The parameters that vary in those experiments are the holding cost rate, h, fixed ordering costs, A, lead time of suppliers other than HL, and the demand variability of high-demand items.

6. Results

This chapter outlines the results of the simulation model for the various scenarios and parameter settings of Chapter 5. Section 6.1 compares the results of the base scenario with the current order parameter settings and for a cost objective and service level model. Section 6.2 shows the performance of each experimental scenario. We conclude this chapter in Section 6.3.

6.1 Base scenario results

This section compares the performances of the costs and service-level model to the order parameters of the current replenishment policy in the case study. In the simulation, all item types are replenished according to the policy of the item type, i.e., stock items have a constant order size, and purchase articles have a variable order size. In reality, it could also happen that stock items have variable order sizes. However, we assume that stock items have a fixed order quantity, as this assumption holds for most historical orders, and the replenishment policy was intended to work that way.

Table 9 shows a summary of the performance of the models (with n=1,754), in which the cost and service-level models outperform the current order parameters in terms of service level and costs. The average order parameters of the service-level model are lower than the current ones. Nevertheless, on average, it results in a higher service level, fewer costs, and fewer stock-outs. With the decrease in the order parameters, less inventory space is necessary for the items. For the model with a cost objective, the required storage capacity increases by 16% (comparing the optimal order-up-to-level S (or $r + Q$, depending on the item) with the current one). On the contrary, the service-level model reduces the required space by 29%. Furthermore, the workload of employees reduces because they need to spend less time on emergency orders. And the average costs are significantly lower for the cost- and service-level model compared to the current order parameters, as the yearly average costs reduce by €204 and €152, respectively.

Table 9: Summary of the base scenario results. In the table, the Min is equal to the reorder point s (r) and the Max is equal to the order-up-to-level S ($r + Q$).

Model	Average Min	Average Max	Service level (%)		Costs per year (€)		Number of stock-outs per year	
			Average	SD	Average	SD	Average	SD
Current	7.5	15.8	98.92%	0.14%	€1,819	€148	2.8	0.4
Service-level	5.4	11.3	99.93%	0.05%	€1,667	€136	0.2	0.1
Cost	8.9	18.3	100.00%	0.00%	€1,615	€134	0.0	0.0

Figure 14b shows how the average costs are divided into the ordering, holding, and shortage costs. The ordering costs are only slightly lower for the current order parameters at the cost of significantly higher holding- and shortage costs. For some items, the models decrease the order parameters while still maintaining the desired service level criterion. This results in a lower average holding cost. Whereas, the order parameters increase for items that fail to satisfy the service level criterion, which results in lower average shortage costs.

Most of the items in this experiment are stock items (n=1,383). The other items are classified as purchase items (n=371). We used a different replenishment policy for them. Table 10 shows a summary of the item types in the experiment. For stock items, the optimal parameters can reduce the average shortage costs by approximately €258 per year, while the yearly average ordering- and holding costs would only slightly increase by €26 and €3, respectively. The optimal parameters reduce the yearly average ordering- and shortage costs of purchase items by €87 and €28, respectively compared to the initial parameters, while the average holding costs would only slightly increase by €3 per year. Appendix C contains a more comprehensive summary of the stock and purchase items.

Table 10: Summary of the base scenario results per item type. In the table, the Min is equal to the reorder point s (r) and the Max is equal to the order-up-to-level S ($r + Q$).

Model	Item type	Average Min	Average Max	Average service level (%)	Average costs (€)
Current	Stock	9.3	19.5	98.65%	€2,073
	Purchase	0.9	2.1	99.94%	€872
Service-level	Stock	6.9	14.0	99.93%	€1,879
	Purchase	0.2	1.2	99.95%	€874
Cost	Stock	11.2	22.6	100.00%	€1,844
	Purchase	0.2	2.2	99.99%	€760

Overall, the models significantly outperform the current parameters in the base scenario. When we compare the results of the service-level model to the current parameters for all the items in the dataset (n=1,754), the average availability of items increases from 98.92% to 99.93% (Figure 14a), and the total costs of the inventory system reduce by approximately €265,000 a year (from €3,190,000 to €2,925,000). The optimal parameters of the model with a cost objective increase the average availability of items to 100.00% and reduces the total costs of the inventory by approximately €355,000 a year, compared to the current order parameters.

Other results of the models are found in Appendix C, which contains all the stored output of the base scenario.

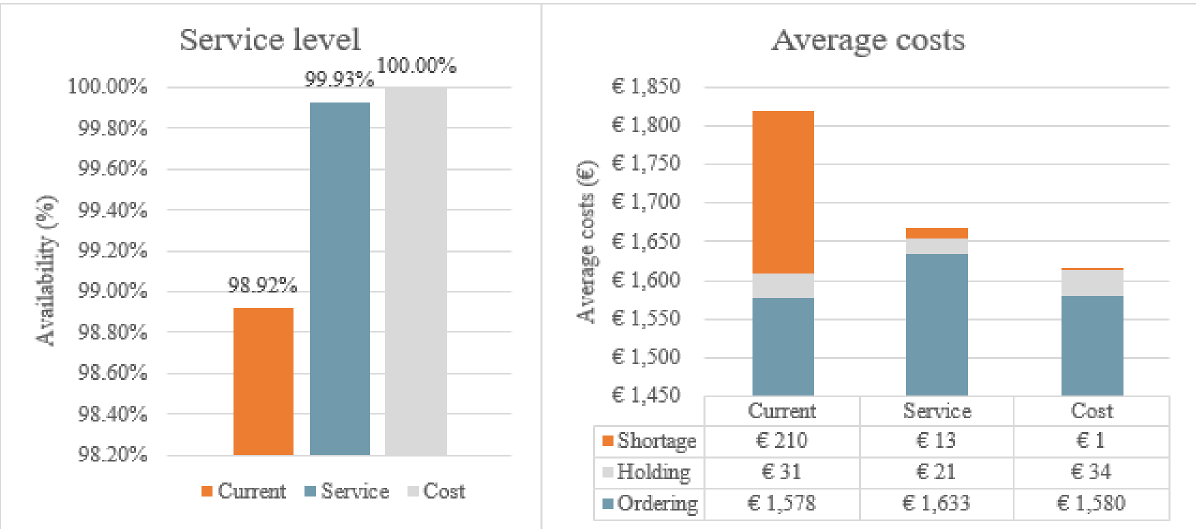

Figure 14: The (a) service level and (b) average costs per year for the current order parameters, service-level model order parameters, and cost model order parameters.

6.2 Experimental scenario results

In this section, we outline the results of the experimental scenarios. Subsection 6.2.1 contains the results of the model with a varying lead time of other suppliers, which only includes the purchase items. In Subsection 6.2.2, we outline the results of various demand variabilities for high-demand items.

6.2.1 Varying lead time (Others) results

We test this experiment with the cost model and a data set of n=371 purchase items. Using the lead time of HL would reduce the average Min level and Max level by 58% and 14%, respectively, compared to the lead times of the base scenario. However, the analysis shows that the order parameters only change for 50 items, which is 13%. For these items, the total Max level reduces by 65 units. The order parameters remain the same for the other supplies (n=321).

Appendix D shows the items that can reduce their order parameters by delivering them via HL. However, it does not automatically mean that Isala should do this. Items that are delivered by HL are stored in their warehouse, which can also be seen as Isala's external warehouse. These are the property of Isala, as HL stores them specifically for Isala. Isala receives an invoice from HL if they deliver the items or if the items can no longer be used because they have expired. As a result, storing items at HL can entail additional costs that fall outside the scope of this research.

6.2.2 Varying demand variability results

We test this experiment with a data set of n=89 high-demand items for six scenarios. Table 11 shows the performance of the scenarios with the cost model. The demand variability of Scenario 1 is equal to the base scenario. The results of the scenario are quite close to the results of Scenario 5. The demand variability is higher for Scenario 6. These results are, therefore, worse with an increase in the average Min, Max, and costs compared to the base scenario. A lower demand variability significantly reduces the average Min, Max, and costs of the inventory system. Scenarios 2, 3, and 4 have average costs per year of €7,211, €6,901, and €6,868, respectively. A VTM ratio of 0.5 (Scenario 3) reduces the required space and average costs for high-demand items by 15% and 9%, respectively compared to the base scenario.

Table 11: The performance of the demand variability scenarios with the cost model order parameters. In the table, the Min is equal to the reorder point s (r) and the Max is equal to the order-up-to-level S ($r + Q$).

Scenario	Parameter value	Average Min	Average Max	Average costs per year (€)
1. (Base)	VTM 2	59.1	118.2	€7,611
2.	VTM 1	54.0	108.0	€7,211
3.	VTM 0.5	50.4	100.9	€6,901
4.	U(0.7, 1.3)	52.4	104.9	€6,868
5.	U(0.5, 1.5)	59.2	118.4	€7,322
6.	U(0.2, 1.8)	69.4	138.7	€8,021

6.3 Conclusion

The optimal order parameters found by the simulation model and its search heuristic significantly outperform the case study's current order parameters. While the performance of the service-level model and the cost model are quite similar. The general tendency was that for most items the current order parameters resulted in high inventory levels, and, consequently, high holding costs. In the contrary, other supplies in the current situation have order parameters which are too low. As a result, the shortage costs and number of stock outs increase. When we compare the results of the service-level model to the current parameters for all the items in the dataset (n=1,754), the average availability of items increases from 98.92% to 99.93%, and the total costs of the inventory system reduce by approximately €265,000 a year (from €3,190,000 to €2,925,000). The optimal parameters of the model with a cost objective increase the average availability of items to 100.00% and reduces the total costs of the inventory by approximately €355,000 a year, compared to the current order parameters.

The scenario with a varying lead time for suppliers other than HL shows that there are 50 purchase items with a decrease in its order parameters. These items have the potential to store them at HL, however, storing items at HL can entail additional costs that fall outside the scope of this research.

The other experiment in which we vary the demand variability shows that less variability can reduce the required space and average costs for high-demand items by 15% and 9%, respectively.

7. Conclusions and Recommendations

This chapter is divided into three sections. First, we answer the main research question in Section 7.1. Then in Section 7.2, we outline the recommendations for Isala. And Section 7.3 describes the discussion of this research.

7.1 Conclusion

In this section, we answer the main research question. This question is formulated as follows:

"How can the on-hand availability of surgical supplies at the Operating Room department of Isala be increased by improving its inventory management?"

Since no information on current service levels is available, we set the desired service levels for medical and non-medical items at 99.9% and 98.0%, respectively. The inventory management of stock and purchase items must improve to achieve these service levels.

Stock items have a constant review period and a short lead time. The lead time for stock items is, most of the time, less than a day as they are supplied multiple times per day from the external warehouse by Isala's logistical partner. The mean usage of these items is higher than the usage of purchase items, and therefore, we use a 2-bin Kanban policy with equal bin sizes, denoted as an (R, r, Q) system, where the reorder point (r) equals the order quantity (Q). This system decreases the workload of employees and the built-in stock rotation reduces the risk of products expiring, leading to better ergonomics.

Purchase items have a constant review period and stochastic lead time. The items can be classified as slow movers due to having a daily usage of fewer than ten units. To use the available capacity efficiently, we use an order-up-to-level replenishment policy with a variable order size, denoted as an (R, s, S) system.

We have developed a simulation-optimization system for OR department inventory management. This simulation-optimization approach uses one of the two developed search heuristics to find the optimal or near-optimal order parameters, depending on the item type. The optimal order parameters found by the simulation model and its search heuristic significantly outperform the case study's current order parameters. The general tendency was that for most items the current order parameters resulted in high inventory levels, and, consequently, high holding costs. On the contrary, other supplies in the current situation have order parameters which are too low. As a result, the shortage costs and number of stockouts increase.

For stock items, the optimal parameters can reduce the average shortage costs by approximately €258 per year, while the yearly average ordering- and holding costs would only slightly increase by €26 and €3, respectively. The optimal parameters reduce the yearly average ordering- and shortage costs of purchase items by €87 and €28, respectively compared to the initial parameters, while the average holding costs would only slightly increase by €3 per year. When we compare the results of the service-level model to the current parameters for all the items in the dataset (n=1,754), the average availability of items increases from 98.92% to 99.93%, and the total costs of the inventory system reduce by approximately €265,000 a year (from €3,190,000 to €2,925,000). The optimal parameters of the model with a cost objective increase the average availability of items to 100.00% and reduces the total costs of the inventory by approximately €355,000 a year, compared to the current order parameters.

By adapting the model's base scenario, several "What-if" scenarios are conducted to observe the impact on the inventory system. The scenario with a varying lead time (Others) shows that there are 50 items with a decrease in its order parameters. These items have the potential to store them at HL, however, storing items at HL can entail additional costs. The other experiment in which we vary the demand variability shows that less variability can reduce the required space and average costs for items with a daily usage of at least ten units by 15% and 9%, respectively.

7.2 Recommendations for practice

This section outlines three recommendations for Isala based on this research.

Change the replenishment policies of stock and purchase items

The first step in improving inventory management is by applying clear replenishment policies. Currently, stock and purchase items are replenished according to a (R, s, S) policy. The order sizes fluctuate a lot as they are either equal to the item's "Min" value or based on personal experience. Changing the replenishment policy of stock items to a 2-bin Kanban system with equal bin sizes reduces this fluctuation and decreases the workload of employees. Furthermore, the built-in stock rotation reduces the risk of products expiring, leading to better ergonomics. We recommend keeping a (R, s, S) replenishment policy for purchase articles. These items are ordered less frequently than stock items, which reduces the impact of fluctuating order sizes on the system. A benefit of using a (R, s, S) policy is that it uses the available storage capacity more efficiently than policies with constant order sizes.

Implement and use the simulation model

We recommend the OR department of Isala to use the simulation model in their decision-making about setting the order parameters. Implementing the model gives more insight into the inventory management process. The visual intuitive experience that it provides to its users could overcome possible resistance during the implementation of the model. The order parameters should be regularly revised to cope with non-stationary demand. We recommend doing this at least once every six months, which requires updating the demand distributions and other input data. Once the system is implemented in the OR department and provides positive results. It should be expanded to other departments in the hospital, which, to the best of our knowledge, requires little change to the model when these other departments use similar data records as the OR department.

Capture the actual demand for items and the service level of the inventory system

It is recommended to capture the actual demand for items and the service level of the inventory system. More accurate data further improves the performance of the simulation model. With the right information, it is easier to see improvement in the system. And items with a lower service level than desired are detected faster.

Expand the use of package sizes in the model

The data set used in the experiments only concerns items with a package size of one unit. However, the model should handle other sizes as well. We recommend including the package sizes of items to find more accurate order parameters and possible faster computation times.

7.3 Discussion

This section is divided into three subsections. Subsection 7.3.1 describes the contribution of this research to theory. Subsection 7.3.2 explains this research's contribution to practice. In Subsection 7.3.3, we describe the limitations of this research and the suggestions for further research.

7.3.1 Contribution to theory

This research contributes to the literature on healthcare inventory management. We developed a simulation model similar to Zhang et al. (2014). We expanded their work by including fractional lead times in the model. Furthermore, the search algorithm we use to find near-optimal and optimal order parameters is a simplified version of the algorithm developed by Kapalka et al. (1999), while we extended the algorithm of Esmaili et al. (2019) by also considering costs. In addition, we formulated two variants of the algorithm, as the 2-bin Kanban system and (R, s, S) system have varying characteristics concerning the order parameters. We performed a case study on the storerooms of the OR department of Isala to determine the model's applicability in practice.

7.3.2 Contribution to practice

This research supports Isala in its decision-making on its inventory management. We provide insight into the current situation and an overview of alternative inventory management policies. The simulation model helps by setting the order parameters of items and a better understanding of the inventory process. The near-optimal or optimal order parameters found by the simulation model significantly outperform the current parameters in terms of service level and costs. Increasing the service level from 98.92% to 100.00% and reducing the total costs from €3,190,000 to €2,835,000 per year. Furthermore, slight adjustments to the model could make it possible to further expands its use in the rest of the hospital.

7.3.3 Limitations and further research

Some of the limitations of this research are:

- Since the available data about the lead time of other suppliers includes backorders and emergency orders it does not represent the actual lead time. Therefore, we assumed that the lead time of other suppliers is uniformly distributed between three to five working days based on information of the experts at Isala.
- Excel is used to simulate the inventory of Isala. Using Excel results in high calculation times and a lot of VBA coding as the program is not optimal for simulations. However, Excel is widely known and available in the healthcare sector. Making it more likely to be implemented into practice. Therefore, it was a suitable option for this research.
- There are no data records available of the actual demand for a certain item. Instead, we used the order data of 2022.

Furthermore, we give some suggestions for further research:

- Research the effect of the 2-bin Kanban system with RFID technology on an item and bin level.
- Perform further research on the decision of which items to store at HL.
- Incorporate the capacity per storage location when determining the optimal order parameters, which changes it from a single-item inventory system to a multi-item one.
- Expand to a multi-echelon supply chain in which also the external warehouse of Isala, which is the storage at HL, is considered.

References

Ahmadi, E., Masel, D. T., Metcalf, A. Y., & Schuller, K. (2018). Inventory management of surgical supplies and sterile instruments in hospitals: a literature review. *Health Systems*, 134-151.

Archibald, B. C. (1981). Continuous review (s,S) policies with lost sales. *Management Science*, 1171-1177.

Bijvank, M. (2009). *Service Inventory Management: solution techniques for inventory systems without backorders.*

Bijvank, M., & Johansen, S. (2012). Periodic review lost-sales inventory models with compound Poisson demand and constant lead times of any length. *European Journal of Operational Research*, 106-114.

Bijvank, M., & Vis, I. F. (2011). Lost-sales inventory theory: A review. *European Journal of Operational Research*, 1-13.

Bijvank, M., & Vis, I. F. (2012a). Inventory Control for Point-of-Use Locations in Hospitals. *The Journal of the Operational Research Society*, 497-510.

Bijvank, M., & Vis, I. F. (2012b). Lost-sales inventory systems with a service level criterion. *European Journal of Operational Research*, 610-618.

Çakici, Ö. E., Groenevelt, H., & Seidmann, A. (2011). Using RFID for the management of pharmaceutical inventory — system optimization and shrinkage control. *Decision Support Systems*, 842-852.

Dellaert, N., & van de Poel, E. (1996). Global inventory control in an academic hospital. *International Journal of Production Economics*, 277-284.

Dillmann, R. (2022, September 12). *Together, we make Isala!* Retrieved from www.iris.isala.nl.

Epstein, R. H., & Dexter, F. (2000). Economic Analysis of Linking Operating Room Scheduling and Hospital Material Management Information Systems for Just-In-Time Inventory Control. *Economics and Health Systems Research*, 337-343.

Esmaili, N., Norman, B. A., & Rajgopal, J. (2019). Exact analysis of (R, s, S) inventory control systems with lost sales and zere lead time. *Naval Research Logistics*, 123-132.

Grond, M. O. (2016). *Computational capacity planning in medium voltage distribution networks.* [Phd Thesis 1 (Research TU/e / Graduation TU/e), Electrical Engineering]: Technische Universiteit Eindhoven.

Hadley, G., & Whitin, T. (1963). *Analysis of Inventory Systems.* Englewood Cliffs: Prentice-Hall.

Hausman, W. H., & Erkip, N. K. (1994). Multi-Echelon vs. Single-Echelon Inventory Control Policies for Low-Demand Items. *Management Science*, 597-602.

Hawa, T. (2020). *Reimagining the Operating Room Inventory as a Center of Savings.* Master of Applied Science, University of Toronto.

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

Huh, W. T., Janakiraman, G., Muckstadt, J. A., & Rusmevichientohg, P. (2008). Asymptotic Optimality of OrderUp-To Policies in Lost Sales Inventory Systems. *Management Science*, 404-420.

Isala. (2022, October 5). *Feiten en cijfers*. Retrieved from www.isala.nl: https://www.isala.nl/over-isala/bestuur-en-organisatie/feiten-en-cijfers/

Johansen, S., & Thorstenson, A. (1993). Optimal and approximate (Q,r) inventory policies with lost sales and gamma-distributed lead time. *International Journal of Production Economics*, 179-194.

Johnson, E. (2002). *Quad wants to be a Savi player in agribusiness.* Case study: 6-0015: Tuck School of Business Dartmouth, Glassmeyer/McNamee Center for Digital Strategies.

Kanet, J., & Wells, C. (2019). Setting bin quantities for 2-Bin Kanban systems (version 3). *Omega*, 142-149.

Kapalka, B. A., Katircioglu, K., & Puterman, M. L. (1999). Retail inventory control with lost sales, service constraints, and fractional lead times. *Production and Operations Management*, 393-408.

Kelle, P., Woosley, J., & Schneider, H. (2012). Pharmaceutical supply chain specifics and inventory solutions for a hospital case. *Operations Research for Health Care*, 54-63.

Kleijnen, J. (1998). Validation of simulation, with and without real data.

Landry, S., & Beaulieu, M. (2010). Achieving lean healthcare by combining the two-bin kanban replenishment system with RFID technology. *International Journal of Health Management and Information*, 85-98.

Landry, S., & Beaulieu, M. (2013). The Challenges of Hospital Supply Chain Management, from Central Stores to Nursing Units. In B. Denton, *Handbook of Healthcare Operations Management* (pp. 465-482). Springer New York, NY.

Lanza-Léon, P., Sanchez-Ruiz, L., & Cantarero-Prieto, D. (2021). Kanban system applications in healthcare services: A literature review. *The International Journal of Health Planning and Management*, 2062-2078.

Law, A. M. (2015). *Simulation Modeling and Analysis.* New York: McGraw-Hill Education.

Neve, B. V., & Schmidt, C. P. (2022). Point-of-use hospital inventory management with inaccurate usage capture. *Health Care Management Science*, 126-145.

Papalexi, M., Bamford, D., & Dehe, B. (2016). A case study of kanban implementation within the pharmaceutical supply chain. *International Journal of Logistics Research and Applications*, 239-255.

Persona, A., Battini, D., & Rafele, C. (2008). Hospital efficiency management: The just-in-time and Kanban technique. *International Journal of Healthcare Technology and Management*.

Rijksoverheid. (2022, October 5). *Registratie van implantaten*. Retrieved from www.rijksoverheid.nl: https://www.rijksoverheid.nl/onderwerpen/medische-hulpmiddelen/registratie-van-implantaten

Rivard-Royer, H., Landry, S., & Beaulieu, M. (2002). Hybrid stockless: A case study. Lessons for health-care supply chain integration. *International Journal of Operations & Production Management*, 412-424.

Robinson, S. (2014). *Simulation: The Practice of Model Development and Use.* Basingstoke: Palgrave Macmillan.

Rosales, C. R., Magazine, M., & Rao, U. (2014). Point-of-Use Hybrid Inventory Policy for Hospitals. *Decision Sciences*, 913-937.

RTL Nieuws. (2022, July 20). *RTL Nieuws*. Retrieved from www.rtlnieuws.nl: https://www.rtlnieuws.nl/onderzoek/artikel/5322022/ziekenhuizen-tekort-medische-materialen-operatie-uitgesteld

Saha, E., & Ray, P. K. (2019). Modelling and analysis of inventory management systems in healthcare: A review and reflections. *Computers & Industrial Engineering*.

Schneider, H. (1978). Methods for Determining the Re-order Point of an (s, S) Ordering Policy when a Service Level is Specified. *The Journal of the Operational Research Society*, 1181-1193.

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2017). *Inventory and Production Management in Supply Chains.* Boca Raton: Taylor & Francis Group.

Smith, S. A. (1977). Optimal inventories for an (S-1,S) system with no backorders. *Management Science*, 522-528.

STZ. (2022, October 19). *STZ ziekenhuizen*. Retrieved from www.stz.nl: https://www.stz.nl/over-ons/stz-ziekenhuizen/

Veral, E., & Rosen, H. (2013). Can a focus on costs increase costs? *Hospital Material Management Quarterly*, 28-35.

Vila-Parrish, A. R., & Ivy, J. S. (2013). Managing Supply Critical to Patient Care: An Introduction to Hospital Inventory Management for Pharmaceuticals. In B. Denton, *Handbook of Healthcare Operations Management* (pp. 447-463). New York, NY: Springer.

Winston, W. L., & Goldberg, J. B. (2004). *Operations Research: Applications and Algorithms.* Belmont: Thomson/Brooks/Cole.

Zhang, X., Meiser, D., Liu, Y., Bonner, B., & Lin, L. (2014). Kroger Uses Simulation-Optimization to Improve Pharmacy Inventory Management. *Interfaces*, 70-84.

Appendix A: Visualization of the simulation model

Product information			
SKU	16	Policy	2-Bin Kanban
Article	[Confidential]	Mean	8.75
Location	[Confidential]	Price	€ 0.42
Type	Grijp	Lead time	0.8
Medical	No	Review	2
HL	Yes	Package size	1

Parameters	
Variability	3
Order costs	1
Holding rate	0.3
Stockout costs	50

Decision variables	
Min	Max
30	**60**
Original	
Min	Max
25	50

Results		
Costs	Availability	Stock outs
€ 1,109.91	99.62%	1
Orders	Average Inventory	Fill-rate
77	32.6	99.78%
Ordering costs	Holdings costs	Stock out costs
€ 1,053.67	€ 4.13	€ 52.11

Simulate Inventory

Dag	Demand	Dag vd week	Dag nummer	Review	Start Inventory	Inventory position	Order	Size	Lead demand	Order ontvangen	End Inventory	OutofStock	Realized demand	Total Realized	Total Demand	Fill rate	Inventory Som	Total Ordered	Missed demand
1	8	Maandag	1	1	36	36					28	0	8	8	8	1	36		0
2	6	Dinsdag	2	0	28	28					22	0	6	14	14	1	64		0
3	7	Woensdag	3	1	22	22	1	30	6	30	45	0	7	21	21	1	86	30	0
4	7	Donderdag	4	0	45	45					38	0	7	28	28	1	131		0
5	8	Vrijdag	0	1	38	38					30	0	8	36	36	1	169		0
6	7	Maandag	1	1	30	30	1	30	6	30	53	0	7	43	43	1	199	60	0
7	10	Dinsdag	2	0	53	53					43	0	10	53	53	1	252		0
8	9	Woensdag	3	1	43	43					34	0	9	62	62	1	295		0
9	6	Donderdag	4	0	34	34					28	0	6	68	68	1	329		0
10	13	Vrijdag	0	1	28	28	1	30	10	30	45	0	13	81	81	1	357	90	0
11	8	Maandag	1	1	45	45					37	0	8	89	89	1	402		0
12	6	Dinsdag	2	0	37	37					31	0	6	95	95	1	439		0
13	5	Woensdag	3	1	31	31					26	0	5	100	100	1	470		0
14	9	Donderdag	4	0	26	26					17	0	9	109	109	1	496		0
15	6	Vrijdag	0	1	17	17	1	30	5	30	41	0	6	115	115	1	513	120	0
16	9	Maandag	1	1	41	41					32	0	9	124	124	1	554		0
17	11	Dinsdag	2	0	32	32					21	0	11	135	135	1	586		0
18	8	Woensdag	3	1	21	21	1	30	6	30	43	0	8	143	143	1	607	150	0
19	8	Donderdag	4	0	43	43					35	0	8	151	151	1	650		0
20	8	Vrijdag	0	1	35	35					27	0	8	159	159	1	685		0
21	2	Maandag	1	1	27	27	1	30	2	30	55	0	2	161	161	1	712	180	0
22	11	Dinsdag	2	0	55	55					44	0	11	172	172	1	767		0
23	14	Woensdag	3	1	44	44					30	0	14	186	186	1	811		0

Figure 15: Visualization of the simulation model.

Appendix B: Sequential procedure

Costs		Mean	Variance	Tvalue	CIHW	RelError	Check
1	2268,74196						
2	2184,05700	2226,39948	3585,77123	12,70620	538,01222	0,24165	Not OK
3	2226,81492	2226,53796	1792,94314	4,30265	105,18624	0,04724	OK
4	2227,25244	2226,71658	1195,42305	3,18245	55,01637	0,02471	OK
5	2187,04416	2218,78210	1211,34747	2,77645	43,21539	0,01948	OK
6	2214,25776	2218,02804	972,48958	2,57058	32,72640	0,01475	OK
7	2207,01504	2216,45475	827,73458	2,44691	26,60816	0,01200	OK
8	2222,06484	2217,15602	713,42091	2,36462	22,33007	0,01007	OK
9	2206,18416	2215,93692	637,61903	2,30600	19,40973	0,00876	OK
10	2172,37488	2211,58072	756,53761	2,26216	19,67606	0,00890	OK
11	2212,98936	2211,70877	681,06423	2,22814	17,53234	0,00793	OK
12	2248,63140	2214,78566	732,75599	2,20099	17,19912	0,00777	OK
13	2256,00096	2217,95607	802,36230	2,17881	17,11723	0,00772	OK
14	2225,17884	2218,47198	744,36844	2,16037	15,75280	0,00710	OK
15	2229,60540	2219,21421	699,46280	2,14479	14,64606	0,00660	OK
16	2255,67744	2221,49316	735,92990	2,13145	14,45551	0,00651	OK
17	2236,07424	2222,35087	702,44063	2,11991	13,62689	0,00613	OK
18	2195,73912	2220,87244	700,46422	2,10982	13,16137	0,00593	OK
19	2216,65488	2220,65046	662,48574	2,10092	12,40570	0,00559	OK
20	2226,94176	2220,96503	629,59709	2,09302	11,74331	0,00529	OK

Figure 16: Sequential procedure with the costs.

Average Inventory		Mean	Variance	Tvalue	CIHW	RelError	Check
1	37,07308						
2	35,82692	36,45000	0,77645	12,70620	7,91694	0,21720	Not OK
3	36,24231	36,38077	0,40260	4,30265	1,57621	0,04333	OK
4	36,69615	36,45962	0,29327	3,18245	0,86172	0,02363	OK
5	35,47692	36,26308	0,41309	2,77645	0,79804	0,02201	OK
6	35,86154	36,19615	0,35734	2,57058	0,62733	0,01733	OK
7	36,52308	36,24286	0,31305	2,44691	0,51746	0,01428	OK
8	35,65769	36,16971	0,31113	2,36462	0,46633	0,01289	OK
9	36,93846	36,25513	0,33791	2,30600	0,44682	0,01232	OK
10	36,35385	36,26500	0,30134	2,26216	0,39269	0,01083	OK
11	35,82308	36,22483	0,28896	2,22814	0,36113	0,00997	OK
12	36,90385	36,28141	0,30111	2,20099	0,34865	0,00961	OK
13	36,24615	36,27870	0,27611	2,17881	0,31754	0,00875	OK
14	35,31154	36,20962	0,32169	2,16037	0,32748	0,00904	OK
15	36,32692	36,21744	0,29963	2,14479	0,30313	0,00837	OK
16	36,67692	36,24615	0,29285	2,13145	0,28836	0,00796	OK
17	36,52308	36,26244	0,27906	2,11991	0,27160	0,00749	OK
18	35,30000	36,20897	0,31410	2,10982	0,27870	0,00770	OK
19	36,81538	36,24089	0,31601	2,10092	0,27094	0,00748	OK
20	36,24615	36,24115	0,29938	2,09302	0,25608	0,00707	OK

Figure 17: Sequential procedure with the average inventory position.

Appendix C: Simulation model output

Table 12: Summary of the average output.

Avg. Output	Initial	Service	Costs
Availability	98,92%	99,93%	100,00%
Costs	€ 1.819	€ 1.667	€ 1.615
Out of Stock	2,81	0,18	0,01
Fill-rate	99,06%	99,66%	99,98%
Number of orders	32,1	36,1	21,7
Inventory level	10,1	6,6	11,9
Ordering	€ 1.578	€ 1.633	€ 1.580
Holding	€ 31	€ 21	€ 34
Shortage	€ 210	€ 13	€ 1

Table 13: Summary of the average output per item type.

Current

	Items	Availability	Costs	Out of Stock	Fill-rate
Stock	1383	98,65%	€ 2.073	3,5	98,92%
Purchase	371	99,94%	€ 872	0,2	99,60%

	Ordering	Holding	Shortage	Number of orders	Inventory level
Stock	€ 1.795	€ 20	€ 259	39,5	12,3
Purchase	€ 767	€ 74	€ 31	4,3	2,0

Service-level model

	Items	Availability	Costs	Out of Stock	Fill-rate
Stock	1383	99,93%	€ 1.879	0,2	99,72%
Purchase	371	99,95%	€ 874	0,1	99,42%

	Ordering	Holding	Shortage	Number of orders	Inventory level
Stock	€ 1.855	€ 13	€ 11	44,4	8,1
Purchase	€ 806	€ 49	€ 18	5,1	1,1

Cost model

	Items	Availability	Costs	Out of Stock	Fill-rate
Stock	1383	100,00%	€ 1.844	0,0	100,00%
Purchase	371	99,99%	€ 760	0,0	99,91%

	Ordering	Holding	Shortage	Number of orders	Inventory level
Stock	€ 1.821	€ 22	€ 1	26,8	14,6
Purchase	€ 680	€ 77	€ 4	2,6	1,7

Appendix D: List of SKUs

Table 14: List of the SKUs with a decrease in the order parameters when delivered by Hospital Logistics.

SKU	Min	Max	Verschil Min	Verschil Max
33	2	3	0	-2
173	0	1	0	-1
185	0	1	-1	-2
283	0	1	-1	-1
305	0	1	-1	-1
339	0	1	0	-1
354	1	2	-1	-1
356	0	2	-1	0
357	0	2	-1	0
358	0	2	-1	0
359	0	2	-1	0
362	1	2	0	-1
367	1	2	-1	-1
368	1	2	-1	-1
598	0	1	-1	-1
628	0	1	-1	-1
803	1	2	-1	-1
873	0	1	0	-1
960	1	2	-1	-1
961	1	2	-1	-1
962	1	2	-1	-1
982	1	2	0	-1
985	1	2	0	-1
995	0	1	0	-1
1008	0	1	-1	-1
1028	1	2	-1	-1
1066	1	2	-1	-1
1081	1	2	-1	-1
1084	2	3	-2	-3
1085	1	2	-1	-3
1177	0	1	-1	-3
1186	1	2	-2	-2
1190	1	2	0	-1
1208	1	2	-1	-1
1209	1	2	-2	-2
1212	0	2	0	-1
1221	1	2	0	-1
1254	1	2	0	-1
1311	0	2	0	-5
1316	0	1	0	-1
1377	2	3	-3	-3
1378	0	1	-1	-1
1448	1	2	-1	-3
1658	0	1	-1	-1
1660	0	1	-1	-1
1666	0	1	-1	-1
1686	0	1	-1	-1
1711	1	2	-1	-1
1712	1	2	-2	-2
1713	1	2	-1	-1