

Devising an inventory control system for an assembly line at Mainfreight

Bachelor Thesis - Industrial Engineering and Management

Casper Slutter

IEM - University of Twente

UNIVERSITY OF TWENTE.



Bachelor Thesis

Title: Devising an inventory control system for an assembly line at Mainfreight

Study: Industrial Engineering and Management

Author

C.P.E. Slutter (Casper)

s1864769

Bachelor IEM

Supervisors University of Twente

1st supervisor: Dr. I. Seyran Topan (Ipek)

2nd supervisor: Dr. E. Topan (Engin)

Supervisor Mainfreight

N. van Benthem (Nick)

University of Twente

Drienerlolaan 5
7522 NB Enschede
The Netherlands

Mainfreight Logistics Services Netherlands

Brede Steeg 1
7041 GV 's-Heerenberg
The Netherlands

Preface

With this thesis, I conclude my Bachelor Industrial Engineering and Management at the University of Twente. I executed my thesis at Mainfreight, where I devised an inventory control system for one of their assembly lines. During the execution, I got support from a number of people and I would like to use this preface to thank them.

First of all, I want to thank Mainfreight for giving me the opportunity to do my bachelor thesis for them. They provided me with a challenging assignment in an interesting production environment. I enjoyed working on it and I gained a lot of practical experience and knowledge. I want to thank all employees, both in the office and on the work floor, who always took the time to answer my questions. Special thanks go to Nick van Benthem, my company supervisor, who showed me around and made me feel at ease at the company. I appreciate his support and the input he gave to me.

I also cannot thank my first supervisor Ipek Seyran Topan enough. She was always there when I needed her help and could always offer good advice. Her involvement goes beyond the thesis itself, which is really outstanding. The meetings we had, the plans we made and the feedback I received were all really valuable towards completing and improving my thesis. Furthermore, I want to thank my second supervisor Engin Topan for providing helpful feedback as well.

Finally, I want to thank my family, especially my mother, for giving me the necessary support during the difficult times that we have been through as a family. It was not always that easy to focus on my thesis, but they tried to make it me as comfortable as possible. Besides, I want to thank my friends Pim and Matthijs for all the helpful conversations and support throughout the execution of my thesis.

Casper Slutter
Uft, August 2020

Management Summary

This management summary shortly discusses the problem, the approach to solve the problem, the solution, the results and recommendations & future research topics.

Introduction

This bachelor thesis is carried out for Mainfreight. Mainfreight is a third-party logistics service provider that operates globally. This assignment is executed for the logistic services business unit in 's Heerenberg. One of the customers of this business unit is an agricultural equipment manufacturer. Mainfreight assembles a large range of their agricultural machines to the specification of the customer in its production area. Here, a lot of different parts of those machines are mounted in different assembly lines. These parts are stored at buffer inventories with limited storage capacity (inventories next to assembly lines), which directly supply the assembly lines. Mainfreight believes that the replenishments to and inventory control of these buffer inventories for different assembly lines are currently not efficient. Because of the time limitations of the bachelor thesis, we focused on only one assembly line (producing two product series), which is considered as all-encompassing. When the problems are solved for this assembly line, it will be easier to solve the problems for other product series at other assembly lines in the future.

We identify the core problem as the lack of a systematic way of inventory control of the buffer inventory. This problem can be solved by devising an appropriate inventory control system. Our aim is therefore to answer the following central research question: *“What is an appropriate inventory control system for the buffer inventory of the A/B series assembly line?”*

Problem solving approach

To answer the main research question, we used the following approach:

1. Analysis of the current situation
To understand the context in which the inventory control system is going to work, we analyse the current production planning process of the assembly line and the current replenishment process of the corresponding buffer inventory. Furthermore, we decide which parts (SKU's) should be included in the inventory control system and determine values for some relevant characteristics of these parts, like storage locations and lead times.
2. Literature study
We conduct a literature study on relevant literature for the inventory control system. This includes study on different demand models (deterministic, stochastic) and different inventory models.
3. Choice of the most appropriate demand model and inventory model
Based on the knowledge gained from the current situation analysis and the literature study, we choose the most appropriate demand model and most appropriate inventory model from the studied literature.
4. Devising the (methodology of the) inventory control system
We devise the methodology of the inventory control system for the buffer inventory. This inventory control system contains the chosen demand and inventory model from literature. In addition, three extensions are made to the inventory control system.
5. Implementation of the devised inventory control system in Excel
The devised inventory control system is implemented in Excel. The models are programmed in VBA and the worksheets are used for the input and output of these models. The company can use this Excel file as a tool to do the replenishments of the parts that are going to be assembled. Furthermore, a manual which explains how to use the Excel tool, is made.

Inventory control system

The devised inventory control system consists of five parts:

- The demand model
Based on a decision rule from literature, demand of 10% of the SKU's is modelled with the Poisson distribution and the demand of the other 90% is modelled with the logarithmic compound Poisson distribution.
- The inventory model
The chosen inventory model is the periodic review stochastic coordinated multi-item inventory model of Fung et al. (2001)¹, which minimises the total expected costs of inventory control per week subject to service level constraints. They used two heuristic algorithms to solve their mathematical model. We used both, but adjusted one of them to find the global minimum expected costs of a specified range at all times. Inputs for the inventory model are mainly: average and variance of the demand, lead times, service levels, major order cost, minor order cost and holding cost. Values for these parameters are not given to us, so we need to determine them. Outputs of the inventory model are: the R_i (the time period between inventory reviews) and S_i (order-up-to-level) parameters for every controlled SKU i . The R_i 's of all SKU's are multiples of a common base period to ensure coordinated multi-item replenishments.
- Extension 1: Calculates the replenishment decisions
This extension automates the derivation of concrete replenishment decisions from the (R_i, S_i) parameters using current inventory levels at a certain replenishment date. With concrete replenishment decisions we mean the required order quantities and an advice on the lots to order to satisfy these required order quantities. This advice consists of the oldest lots, so that a FIFO policy can be adhered to.
- Extension 2: Buffer inventory capacity check
Since the inventory model does not take the capacity of the buffer inventory into account, we include a capacity check on the results of the inventory model. It checks if the current capacity of the buffer inventory is large enough and it indicates which parts of the buffer inventory could have a smaller capacity and which parts should have a larger capacity. This extension is optional and does not influence the results of the inventory model as it is not a constraint in the model but only a check on the results.
- Extension 3: Checks if the replenishment decisions lead to stockouts, based on demand information from outstanding orders
The inventory control system uses stochastic demand to find the replenishment decisions. However, there is short term deterministic demand information available from outstanding orders. This extension uses this information to predict if and where stockouts will occur if the replenishment decisions are followed. This extension is optional and does not influence the replenishment decisions.

Results of the implemented inventory control system

After the devised inventory control system is implemented in the Excel tool, we can find and analyse the results of the proposed inventory control system. It turns out that the least total expected costs of inventory control for a service level of 98% are achieved with a base period of two working weeks. This means that it is optimal to **do a replenishment once every two working weeks (ten working days)**. Some SKU's should be replenished every two working weeks whereas the other SKU's should

¹ Fung, R. Y. K., Ma, X., & Lau, H. C. W. (2001). (T,S) policy for coordinated inventory replenishment systems under compound Poisson demands. *Production Planning & Control*, 12(6), 575-583.

be replenished after a multiple of two working weeks (e.g. every four, six or eight working weeks). The resulting **total expected costs are then equal to €413.89 per week** and the **mean service level will be equal to 98.364%**. However, the capacity check extension (Extension 2) indicates that the capacity of the buffer inventory is not large enough.

Subsequently, we perform sensitivity analyses on different parameters to determine which parameter influences the output of the inventory control system in what degree. Our main findings are that: a change in holding cost has a larger influence on the expected costs than a change in major or minor order cost. Besides, the expected costs increase more quickly with each equal step higher in service level. Furthermore, a change in size of the overflow buffer inventory capacity part has a smaller influence on the net number of pallets short in the buffer inventory than the two other buffer inventory capacity parts, which have approximately the same influence.

Finally, we determine the improvement over the current situation when the proposed inventory control system will be used. First, we measure the improvement in total expected costs (major order cost, minor order cost and holding cost), our KPI. The total expected costs per year of the current situation is determined as the total expected costs over the past year, whereas the new situation is given by the optimal solution of the inventory control system. The results are given in Table MS.1.

Table MS.1 The minor order cost, major order cost, holding cost and total expected costs per year for both the current situation and the new situation, together with the percentage increase or decrease for each cost.

| | Current situation | New situation | Percentage in/decrease |
|--------------------------------------|-------------------|-------------------|------------------------|
| Minor order cost per year | € 4274.00 | € 5621.57 | 31.53% |
| Major order cost per year | € 9080.00 | € 1754.22 | -80.68% |
| Holding cost per year | € 10138.57 | € 14146.57 | 39.53% |
| Total expected costs per year | € 23492.57 | € 21522.36 | -8.39% |

It turns out that there is indeed an improvement in expected costs over the current situation: **the decrease in total expected costs will be 8.39%** according to our calculations, just for one assembly line. This is realised by a **large decrease in number of replenishments per year**, which results in a decrease in total order cost that is larger than the increase in holding cost.

We also expect **other improvements** like setting and achieving service levels, reducing the number of emergency deliveries and standardising the replenishment process, when the proposed inventory control system is used.

Recommendations and future research

Next to recommending to use the Excel tool with the implemented inventory control system, we make other recommendations to the company, such as:

- Making a buffer inventory location for every controlled SKU
- Updating the estimations of demand and cost parameters regularly
- Observing what will happen with the buffer inventory utilization when the inventory control system is used and if needed, enlarge its capacity according to the capacity check results
- Extending the inventory control system with a required trailer loading meters calculation
- Extending the inventory control system to multiple assembly lines

We also propose topics for future research, including:

- Production scheduling of outstanding orders
- Forecasting of demand
- Multi echelon inventory models, if inventory control at external warehouses is also desired

Table of Contents

| | |
|---|----|
| Preface..... | i |
| Management Summary..... | ii |
| 1. Introduction..... | 1 |
| 1.1. Problem description | 1 |
| 1.2. Identification of the core problem | 1 |
| 1.3. Problem solving approach..... | 3 |
| 2. Current situation | 7 |
| 2.1. Production planning process..... | 7 |
| 2.2. Replenishment process | 9 |
| 2.3. Critical parts..... | 12 |
| 2.3.1. Classification of C parts | 12 |
| 2.3.2. ABC inventory classification of L parts | 13 |
| 2.3.3. Conclusion on the selected critical parts..... | 15 |
| 2.4. Characteristics of the critical parts..... | 15 |
| 2.4.1. Part commonality | 16 |
| 2.4.2. Storage locations | 16 |
| 2.4.3. Lead times | 17 |
| 2.4.4. Lot sizes | 17 |
| 2.4.5. Buffer capacity..... | 18 |
| 2.5. Conclusion on current situation | 19 |
| 3. Literature review | 22 |
| 3.1. Theoretical perspective and theoretical framework..... | 22 |
| 3.2. Inventory classification methods | 24 |
| 3.3. Demand models..... | 26 |
| 3.3.1. Deterministic demand models | 26 |
| 3.3.2. Stochastic demand models..... | 27 |
| 3.4. Inventory control system..... | 32 |
| 3.4.1. Single item inventory models..... | 33 |
| 3.4.2. Coordinated multi-item inventory models..... | 39 |
| 3.5. Conclusion on the literature review | 44 |
| 4. Inventory control system methodology | 47 |
| 4.1. Choice of the inventory control system | 47 |
| 4.1.1. Choice of the demand model | 47 |
| 4.1.2. Choice of the inventory model | 51 |
| 4.2. Methodology of the inventory control system | 52 |

| | | |
|--------|--|-----|
| 4.2.1. | Description of the demand model | 53 |
| 4.2.2. | Description of the inventory model | 54 |
| 4.2.3. | Extension which calculates the replenishment decisions | 63 |
| 4.2.4. | Buffer inventory capacity check extension | 66 |
| 4.2.5. | Outstanding orders check extension..... | 70 |
| 4.3. | Implementation of the inventory control system in Excel | 77 |
| 4.3.1. | Relation between the workflows | 79 |
| 4.3.2. | Implementation of workflow 1..... | 80 |
| 4.3.3. | Implementation of workflow 2..... | 92 |
| 4.3.4. | Conclusion on the implementation of the inventory control system | 105 |
| 4.4. | Determination of values for inventory control system parameters | 106 |
| 5. | Results of the implemented inventory control system..... | 111 |
| 5.1. | Raw results | 111 |
| 5.2. | Sensitivity analyses..... | 116 |
| 5.2.1. | Sensitivity analysis with changing cost parameters (K, k_i, h_i) | 117 |
| 5.2.2. | Sensitivity analysis with changing service level (P_i)..... | 120 |
| 5.2.3. | Sensitivity analysis with changing buffer inventory capacity size | 122 |
| 5.3. | Improvement compared to the current situation..... | 123 |
| 5.3.1. | Improvement in expected costs | 123 |
| 5.3.2. | Other improvements | 126 |
| 6. | Conclusions and Recommendations | 128 |
| 6.1. | Conclusions..... | 128 |
| 6.2. | Recommendations and Future research | 132 |
| 6.3. | Contribution to theory and practice..... | 135 |
| | References..... | 137 |
| | Appendix A: Critical parts and values of their relevant characteristics | 138 |
| | Appendix B: Goodness of fit tests | 138 |
| | Appendix C: Values of inventory control system parameters..... | 138 |
| | Appendix D: Manual of the inventory control system | 140 |
| | Appendix E: VBA code of the inventory control system implementation in Excel | 140 |
| | Appendix F: Detailed results of the best replenishment schedule | 140 |
| | Appendix G: Determination of the costs of the current situation | 140 |

1. Introduction

1.1. Problem description

This bachelor thesis is executed for Mainfreight. Mainfreight is a third-party logistics service provider that operates globally. This assignment is executed for the logistic services business unit in 's Heerenberg. One of the customers of this business unit is an agricultural equipment manufacturer. Mainfreight assembles agricultural machines for this agricultural equipment manufacturer. For this assignment, the focus is on one of the assembly lines in Mainfreight's production area, namely the one of the small machines: type A, B small and B large. These machines are assembled to the specification of the customer: all kinds of parts are mounted to the machines. These parts are stored at a buffer inventory (an inventory next to the assembly line), which directly supplies the assembly line. This buffer inventory is replenished by multiple external warehouses of Mainfreight at different locations. Currently, the following agreement is made with the warehouses: "Ordered before 12:00 (noon), delivered at the buffer inventory on the same day". Replenishments are placed from Monday to Friday. Furthermore, assembly takes place in one shift from Monday to Friday from 8:00 till 17:00. The assembly takes place on 15 production spots where, on average, around 40 type A/B machines in total are produced each week. This number should go up to 60 per week within a short time.

This project is started since Mainfreight believes that the replenishments to and inventory control of the buffer inventory are currently not efficient. This management problem was provided to me in the following form: How can Mainfreight control the incoming flow, storage and replenishment between buildings of parts more efficiently to improve productivity? Although this problem exists at multiple assembly lines, it is only researched for one assembly line in order to make the problem narrow enough given the available time. The assembly line of the A/B series is chosen since it is all-encompassing. When the problem is solved for this assembly line, it will be easier to solve the problems at the other assembly lines in the future.

1.2. Identification of the core problem

Problem cluster

The first step in identifying the core problem is to make a problem cluster of the problems related to the management problem. This problem cluster can be found in Figure 1.1.

Mainfreight knows that the current situation is not efficient, since they encounter high costs and frustration among employees. These problems therefore form the start of the problem cluster and are visible at the bottom of the problem cluster. We found that these problems are caused by wrong replenishment of the buffer inventory. This leads to either too many unneeded parts in the buffer inventory (leading to high holding costs and space usage) or too few required parts in the buffer inventory. The latter leads to very costly production stops and extra replenishments in the form of express deliveries, which causes high ordering costs. Too few required parts in the buffer inventory can also be caused by slow delivery of replenishment orders to the buffer inventory and bad communication about the replenishments between the assembly line and external warehouses.

Next, we found a cause for the wrong replenishment of the buffer inventory: the replenishments are done by the employees, who base the replenishment decisions on their feeling and experience only. There is no system or procedure that gives the employees replenishment decisions. Therefore, we conclude that the replenishments based on feeling and experience are caused by the fact that there is no systematic way of inventory control. If there is no system which tells the employees how much of a certain part is needed at a certain moment, the consequence is that they should base the replenishments on their own feeling only.

However, we determined another cause for the replenishments based on the feeling and experience of the employees. Namely, there is no production schedule that they can consult to determine the replenishment decisions. Therefore, the employees cannot base their decisions on anything other than their feeling and experience.

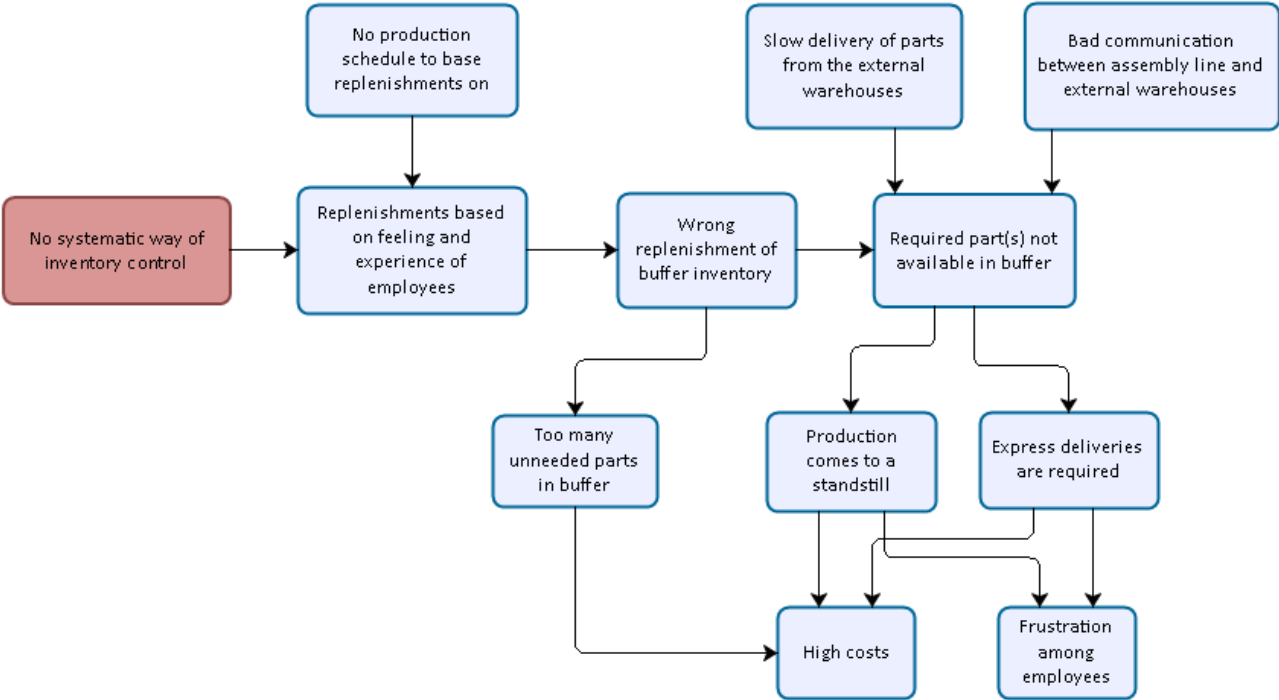


Figure 1.1 The problem cluster with the core problem given in red.

Selection of the core problem

According to Heerkens (2012), for the selection of the core problem, we need to go back in the chain of the problem cluster to find the problems which do not have a cause itself (Heerkens & Van Winden, 2012). We find these problems in the top row and most left column of the problem cluster and they are as follows:

- Bad communication between assembly line and external warehouses
- Slow delivery of parts from the external warehouses
- No production schedule to base replenishments on
- No systematic way of inventory control

We are not sufficiently convinced that the first point really is a (big) problem at the moment, since they use a clear shared Excel file to communicate. In that case, we should not choose it as the core problem according to Heerkens (2012). The second problem seems to be an existing problem at the moment, especially if it concerns the external warehouse which is not located at the same industrial area. However, this is a temporary warehouse, so we should not take this problem into account. Furthermore, we do not have any influence on the location or travel time between the external warehouses and the production location. Since we cannot influence this problem, it cannot be a core problem according to Heerkens (2012).

The third problem, however, can clearly be influenced. There is a list available with all production orders that are in the pipeline, these orders will be called “outstanding orders” in the rest of this report. A weekly production schedule can thus be made based on these outstanding orders.

Nevertheless, the question remains whether solving this problem solves the problem of the wrong replenishments and all problems resulting from them. Namely, the employees should still interpret the production schedule and should come up with replenishment decisions themselves. On the other hand, we can be quite certain that with solving the fourth problem, the problem of the wrong replenishments and all problems resulting from them will be solved. This is the case since a system will provide the employees directly with replenishment decisions. In this way, the interpretation of the employees is not required at all, they should just follow the provided replenishment decisions. An inventory control system will therefore solve the most problems from the problem cluster. Hence, we choose “No systematic way of inventory control” as the core problem.

The action problem can be formulated in the following way: “There is no inventory control system, whereas there should be an appropriate inventory control system to base replenishment decisions on.”

From this, the norm and reality become clear:

Reality: There is no inventory control system. There is only a warehouse management system which indicates what parts are stored at which locations in which quantities.

Norm: There should be an appropriate inventory control system to base replenishment decisions on.

The norm is measured by means of the variable “appropriateness”. This variable is made measurable by measuring whether predetermined criteria for “appropriateness” are met. When all these criteria are met, the norm has been reached and the problem has been solved.

The criteria for “appropriateness” are:

- A high level of parts availability should be attained
- All important parts should be included by the system
- It should give the moments when replenishments are required
- It should give the required parts for each replenishment
- It should give the required quantities of these required parts
- It should take the current stock level into account
- It should take the capacity of the buffer into account
- It should take lot sizes into account
- It should take lead times into account

To find an inventory control system which meets these criteria for appropriateness, the following central research question will be answered in this thesis: *What is an appropriate inventory control system for the buffer inventory of the A/B series assembly line?*

When there is an inventory control system with these criteria, there will be a systematic way of inventory control. In that case, the replenishments can be directly derived from this system and will not be based on the feeling and experience of employees anymore. It is believed that this gives a more optimal buffer inventory in terms of parts availability and thus prevents the resulting problems from the problem cluster from happening.

1.3. Problem solving approach

In this section, the approach to solve the core problem is described. First of all, we determine the scope. Subsequently, we describe the stages of the problem solving approach and the knowledge questions which are answered in each stage. Finally, we discuss the deliverables in this section.

Scope

As mentioned before, this thesis focuses on the replenishment and inventory control processes of

one assembly line only (type A and B machines). Within these processes, we focus on optimizing the buffer inventory. It is treated as a single echelon problem: although there are multiple suppliers (warehouses) for the buffer inventory, we consider them as one supplier with the policy: “ordered before 12:00, delivered on the same day”. So there is only one (buffer) inventory with one flow going inside and one flow going outside to consider. The inventories at other warehouses are not controlled.

Furthermore, the inventory control system is devised only for the most important parts in the first place. This is done to make the problem small enough and to only include parts which are expected to have a substantial influence. We will call these parts “critical parts” in the rest of this report. These critical parts are determined in the current situation analysis, see below. However, we take the future addition of other parts of the A/B series into account.

Stages of the problem solving approach

Stage 1: Current situation analysis (Chapter 2)

Given the scope determined in the previous two paragraphs, the first step to take is to analyse the current situation. This analysis is executed in Chapter 2. It is important to discover how the production planning of the assembly line and the replenishments of the corresponding buffer inventory are currently carried out. Namely, we should understand the context in which the inventory control system is going to work in order to devise an appropriate inventory control system. As a result, the following knowledge question is answered in Chapter 2:

1. What is the current situation regarding the planning and control of the A/B series assembly line and corresponding buffer inventory?

The inventory control system is devised for the critical parts only in the first place and therefore these parts need to be determined, using an inventory classification method. Hence, an appropriate inventory classification method is chosen in Section 3.2. Then, based on the chosen method, the critical parts are determined in Chapter 2. Furthermore, we determine values of characteristics of these parts which are considered relevant for the inventory control system. These characteristics are: part commonality, storage locations, lead times, lot sizes and buffer inventory capacity. This leads to the following knowledge question, which is answered in Chapter 2:

2. Which parts of the A/B series should be included in the inventory control system and what are the values of their relevant characteristics?

Stage 2: Literature study on inventory control systems (Chapter 3)

After analysing the current situation, the next step towards solving the core problem is conducting study on relevant theory from literature. This literature study is performed in Chapter 3. To determine which literature is relevant for inventory control systems, we first make a theoretical framework in this chapter. From this framework, it appears that a demand model is needed in our inventory control system. There are a lot of demand models and each one models demand in a different way. Therefore, it is important to have knowledge regarding the different demand models. Hence, we conduct a literature study on demand models by answering the following knowledge question:

3. What demand models are proposed in inventory management literature?

Furthermore, the inventory control system uses an inventory model to come up with the replenishment decisions. A large variety of inventory models can be used in an inventory control system. It is important to gain knowledge about these different inventory models. Therefore, we conduct a literature study on inventory models by answering the following knowledge question:

4. What inventory models are used in inventory control systems according to literature?

Stage 3: Choose the most appropriate inventory control system (Section 4.1)

After analysing the current situation and conducting literature study, we choose the most appropriate inventory control system based on the gained knowledge. This choice is made in Section 4.1 of Chapter 4. We make the choice by choosing the most appropriate demand model and the most appropriate inventory model from literature. These two models form the foundation of the inventory control system: the demand model is an input to the inventory model and the inventory model provides the replenishment decisions.

Stage 4: Devising the inventory control system (Section 4.2)

In the previous stage, the foundation for the most appropriate inventory control system is chosen. In this stage, the inventory control system is devised based on this foundation. This is performed in Section 4.2 of Chapter 4, where we describe the methodology behind the devised inventory control system. The devised inventory control system should fulfil the criteria for “appropriateness”. However, it cannot fulfil all these criteria only by the demand and inventory model (foundation) as found in literature. Therefore, the inventory control system is extended with tailor-made parts so that all criteria are fulfilled. We call these tailor-made parts “extensions” and their methodology is also described in Section 4.2. The complete devised inventory control system therefore consists of the demand model, the inventory model and extensions.

Stage 5: Implementation of the inventory control system (Section 4.3)

After devising the inventory control system, it is implemented so that the company can actually make use of it. The implementation is described in Section 4.3 of Chapter 4. The implementation is in the form of a tool which the company can use to do replenishments with. This tool is made in an Excel file. In this file, the models behind the inventory control system are programmed in VBA. The inventory control system can be controlled by user forms in the worksheets. The worksheets act as a user interface where the employees give input and receive output of the inventory control system. The output is the replenishment decisions: when to order each critical part and how much to order at these instances. The employees can directly use this information to do the replenishments. In order to guarantee the usefulness of the tool, immediately and in the future, a manual is made. This manual describes how the tool is set up and how it should be used. Besides, it describes how the input can be changed, in case that is necessary in the future.

Stage 6: Finding and analysing the results of the inventory control system (Chapter 5)

Now that the devised inventory control system is implemented in an Excel tool, we can use this tool to determine the results of the inventory control system in Chapter 5. To this end, we first determine values for the required input to the inventory control system in Section 4.4. We use these values in the tool with the implemented inventory control system of the previous stage, to find the results (output) of the inventory control system. Furthermore we analyse the results by doing sensitivity analyses on different parameters that are inputs to the inventory control system. Finally, we determine the improvement over the current situation by comparing the results of the current situation with the results of the situation in which the inventory control system is used.

Stage 7: Drawing conclusions and making recommendations (Chapter 6)

After all stages have been carried out, we can draw conclusions based on the knowledge that we gained and the solution of the core problem that we made. We draw these conclusions in Chapter 6. Besides, we make recommendations to the company and propose topics for future research in this chapter. We conclude with a description of the contributions that we made to theory and practice.

Deliverables

From the stages of the problem solving approach, several deliverables can be derived. They are summarised below:

- Analysis of the current situation: current production planning and replenishment process, analysis of the critical parts and values of their relevant characteristics
- Overview and classification of different demand and inventory models
- Models of all the parts of the devised inventory control system
- Excel tool in which the inventory control system for the critical parts is implemented
- Manual of the Excel tool, which describes how the tool is set up, how it should be used and how input can be changed in the future.

2. Current situation

In this chapter, the current situation regarding the planning and control of the A/B series assembly and the required parts is analysed. It is important to understand the current situation for a number of reasons. First of all, we should become familiar with the processes to know the context in which the inventory control system is going to work. Secondly, we can determine the scope of the inventory control system further, i.e., which parts should be included and which not. Thirdly, we can derive properties which the inventory control system should take into account, e.g., lot sizes and lead times. The current situation is analysed by answering two knowledge questions.

The first knowledge question is *“What is the current situation regarding the planning and control of the A/B series assembly line and corresponding buffer inventory?”*. This knowledge question is subdivided into two sub-questions, which are answered in the separate Subsections 2.1 and 2.2:

- Subsection 2.1: *“What does the production planning process of the A/B series assembly line currently look like?”*
- Subsection 2.2: *“What does the replenishment process of the A/B series buffer inventory currently look like?”*

The second knowledge question is *“Which parts of the A/B series should be included in the inventory control system and what are the values of their relevant characteristics?”*. This knowledge question is subdivided into two sub-questions. The first sub-question is answered in Subsection 2.3 and is concerned with determining the parts which should be included in the inventory control system. Then, values of several characteristics of the selected parts are determined. This is executed in Subsection 2.4 by means of answering the second sub-question:

- Subsection 2.3: *“What are the critical parts of the A/B series?”*
- Subsection 2.4: *“What are the values of some relevant characteristics of these critical parts?”*

At the end of the chapter, a conclusion on the current situation is given in which the two knowledge questions will be answered. This conclusion can be found in Subsection 2.5.

2.1. Production planning process

To get a good understanding of the current situation, it is important to discover how the production planning regarding the A/B assembly line is currently performed, by answering the sub-question: *“What does the production planning process of the A/B series assembly line currently look like?”*.

At the assembly line of the A/B series, production takes place according to a week planning. At the end of the week, the production planning of the next week is made. This planning is based on a list with orders to be built (outstanding orders). This list is updated with new orders from time to time by the administration department. When the planning is made, at least all the orders that should be finished in the next week are certain. Generally, already more orders are included in the list.

The most relevant columns from this list are the build week, requested ship date, the type of the machine and the expected build time. The build week is the week in which the order should ultimately be built, in order to meet the requested ship date. This requested ship date is the due date on which the order should be ready. The build week is therefore the week preceding the week in which the requested ship date falls. The type of machine is one of the 3 types that is assembled at the assembly line: A, B small (BS) or B large (BL). The expected build time is the time that it should take the mechanic to build the order. This expected build time is based on time measurements and includes every activity related to building that specific order. This expected build time is therefore related to the size of the order. A large order has a large amount of options and/or complex

combinations of options to be built on the machine and has therefore a higher expected build time. The expected build times can strongly fluctuate based on the type of machine and the size of the order. Although there is no clear definition of the size of a small order or a large order, we can give some examples of extreme values, see Table 2.1.

Table 2.1 Examples of extreme expected build times

| Type of machine | Build time small order (minutes) | Build time large order (minutes) |
|-----------------|----------------------------------|----------------------------------|
| A | 130 | 536 |
| BS | 492 | 1246 |
| BL | 18 | 1018 |

Given these strong fluctuations in expected build times, the weekly production output also fluctuates. There can be weeks with a lot of large orders and weeks with a lot of small orders, where the weekly production output is lower than average and higher than average respectively.

Based on the list with orders to be built, the production planning is made. The following three steps are taken to come up with the weekly production planning:

- Selecting the orders to produce next week

The main criteria on which they decide on which orders to produce is the build week. In any case, the orders with the build week corresponding to the current week should be built in that week. When they are behind schedule, orders from past build weeks should be treated first. When they are ahead of schedule, orders from future build weeks are built. However, this is on the condition that the orders can be built. If they already know that a certain type of machine cannot be built, it will not be planned for assembly. This happens when there is a production stop of that type because there are no basic machines or other important parts for that type on hand.
- Selecting the number of mechanics for next week

The number of mechanics is chosen dependent on the production planning. There is a maximum capacity of 15 mechanics, since there are 15 production spots and one mechanic occupies one spot. Every mechanic works 480 minutes per day and 5 days per week. This means that every mechanic is available for 2400 minutes per week to build orders. As long as there are enough outstanding orders for the current and future build weeks, the maximum capacity of 15 mechanics is used. However, when they are so far ahead of schedule that there are not enough outstanding orders for 15 mechanics, then less mechanics are scheduled. This also shows strong fluctuation since the machines are subject to seasonal demand: one period there could be relatively many outstanding orders, while the other period could have relatively few outstanding orders.
- Allocating the mechanics to the types of machines

The scheduled mechanics are then allocated to one of the three types. However, not all mechanics have the same skill level. Therefore, it needs to be taken into account that some mechanics do not have the skill to assemble certain types of machines. Furthermore, some mechanics are more skilled and thus faster than others, so these mechanics are assigned to the type of machine that needs the most attention. However, Mainfreight is working towards a situation where every mechanic has the skill to assemble every order of every type of machine. In that case, the production planning becomes more flexible and thus easier. This is the case, since all combinations of mechanics and orders are possible then. However, in the meantime a learning curve will likely be visible: when mechanics are learning to make new

types of orders, they will initially not make it within the expected build time. However, when they get more and more experienced, they will likely be able to assemble the orders within the expected build time.

With the current production planning approach, there is only a production planning for the week. This means that it is known which orders are going to be built next week, but it is not known in which sequence they are going to be built and by whom exactly. In other words, they know the production schedule on a weekly basis but not on a daily basis. On a daily basis, orders are assigned to mechanics when they are (almost) finished with their previous order. Generally, the list of orders, which is in ascending order based on build week, is worked down from top to bottom. The next order in line is put on a written list. When a mechanic is finished with his order, he puts his name on this written list behind this next order in line and he will assemble that order. However, a skilled mechanic can be asked to do a more difficult order instead of the order next in line.

Therefore, a clear, daily production schedule is not determined in advance. The scheduled mechanics are given the set of scheduled orders and they can determine themselves how they are going to complete the orders. The only condition is that the scheduled orders should all be finished at the end of the week. This indefinite way of scheduling orders has as result that it is not exactly known in advance when each order is started. Consequently, the part requirements are not known on a day level.

2.2. Replenishment process

Next to the production scheduling, the replenishment process regarding the A/B assembly line needs to be analysed to get a good understanding of the current situation. Therefore, the following sub-question is answered in this section: *“What does the replenishment process of the A/B series buffer inventory currently look like?”*.

Mainfreight does the assembly of machines according to the specification of the customer. For the assembly, it needs both basic machines and parts which should go onto these machines, as required by the customers. Therefore, it needs to keep inventory of the basic machines and all possible parts which go onto these machines. These machines and parts are supplied to Mainfreight, mostly by containers via the ocean, but also by trucks via the road. The inbound department of Mainfreight decides where the contents of the containers/trucks are stored. It is not possible to store everything at the production location and therefore external warehouses exist.

The arrangement of production and inventory of the A/B series is as follows. Basically, there are a few locations where Mainfreight is occupied with operations for this customer. There are inventory-only locations where the parts are stored. These are the external warehouses, which have different distances to the production location. This production location is one large location, which consists of multiple halls where both production takes place and inventory is held. The relevant halls for the A/B series include hall 5, 6, 7, 8, 21 and 22. The actual production of the A/B series takes place in hall 21. Inventory of the A/B parts is held in all the aforementioned halls. This inventory can be divided into buffer inventory and bulk inventory.

Buffer inventory is the inventory of the A/B series parts which is directly available for production. This means that it can be quickly picked for assembly by the employees. Almost all buffer inventory is located in hall 21, since this is very close to the assembly line. All (buffer) inventory in this hall is completely dedicated to the A/B series. Next to that, hall 6 and 7 also contain a very small buffer inventory for certain A/B parts. In hall 6, only two critical parts are stored as buffer inventory. These parts are stored there because they are more often used by another series, which is assembled in

this hall. In hall 7, part of the inventory of the MD parts is stored due to a lack of space in hall 21. However, the largest and most important buffer inventory is in hall 21.

The rest of the inventory of A/B parts at the production location is stored in the bulk inventory. Bulk inventory is the inventory which is stored far away and often high on the shelves in hall 5, 6, 7, 8 and 22. This inventory cannot be considered as buffer inventory as parts cannot easily be picked for assembly from here. They should first be moved to the buffer inventory instead. The bulk inventory is very large and is used by parts from all machine series. It does not contain fixed places for the A/B series.

In conclusion, there are three inventory types: inventory at external warehouses, bulk inventory at the production location and buffer inventory at the production location. This is visualised in Figure 2.1.

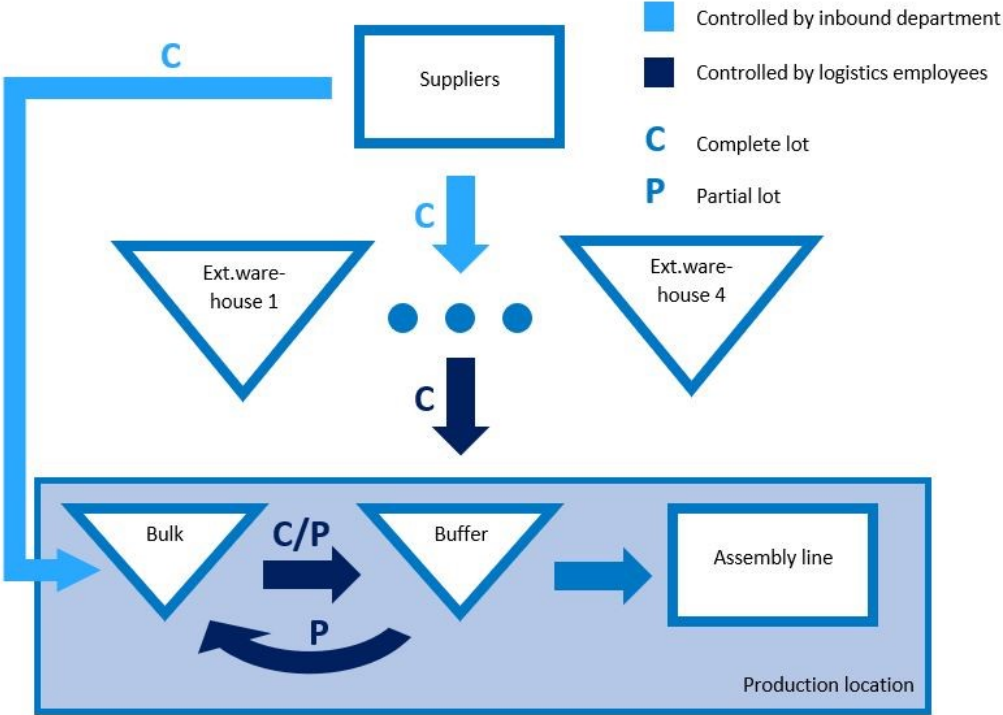


Figure 2.1 The replenishment process

The arrows in Figure 2.1 indicate the control of the inventory. The control of the inventory at the external warehouses and the bulk inventory is done by the inbound department. Since the parts arrive in large quantities via containers to Mainfreight, this department is concerned with distributing these parts among the external warehouses and the bulk inventory. On the other hand, the control of the buffer inventory is done by the logistics employees of the A/B series. This means that they decide what parts should be held at the buffer inventory in which quantities. They also decide on the replenishments from the bulk inventory and the external warehouses to the buffer inventory. This project only concerns these replenishments and therefore not the replenishments to the bulk inventory and the external warehouses themselves. How the replenishment process from the external warehouses and the bulk inventory to the buffer inventory currently occurs, is explained now.

The buffer inventory consists of multiple shelves which accommodate pallets. These pallets contain the parts required for the assembly. The pallet size and the quantity that the pallet holds is

dependent on the part and can greatly vary. For example, a pallet can only hold one CB part, while the same pallet could also hold hundred boxes of a small part. The buffer inventory is designed in such a way by the employees that it contains (almost) all the required parts for assembly. They allocated a certain number of pallet locations to a certain part. Generally, multiple pallet locations are assigned to parts of which only one or a few fit on a pallet and only one pallet location is assigned to parts of which a large quantity fit on a pallet.

To determine whether parts need replenishment, employees observe the buffer inventory. For the parts that are stored on multiple pallet locations, the employees look at the number of empty pallet locations. When they see that a large number of pallet locations allocated to a part, is empty, they decide to replenish that part so that its empty pallet locations become full again. For the parts that are stored on one pallet location, the employees look at the amount that is left on this pallet. When they see that only a small amount is present on the pallet, they decide to replenish that part. However, a "small amount left" is rather subjective. To determine what "small" is, the employees use their feeling and experience. They roughly know how fast-moving a part is and therefore know till when they can replenish safely, that is, when there is still enough left in stock to fulfil the demand during the replenishment lead time. However, they do not know this for sure, as they do not use the actual information about the demand during the replenishment lead time, even while this information is available through a file with the outstanding orders.

If they have decided that a part needs replenishment, they proceed as follows. They search for the part in the Warehouse Management System (WMS). Here, they see all the available lots of this part, including their quantities and locations at the external warehouses and/or bulk inventory. According to the company's procedure, they should order the lot with the oldest lot date. In fact, this is a FIFO policy, to make sure that the parts which have first arrived at Mainfreight also get used the first. However, this procedure is not strictly followed as sometimes picking a newer lot is more convenient than picking the oldest lot. This is not much of a problem when the difference in lot dates is small, but it is not desirable when the difference is large. Again, whether a difference is small or large is rather subjective.

When they have chosen the lot, two scenarios can occur. When the lot is located in the bulk inventory, the employees can pick this lot themselves and use it to replenish the buffer inventory. When the lot is located in an external warehouse, the employees should fill in a so called "move file". This move file consists of multiple lists, each list covers one specific replenishment from an external warehouse to the production location. The employees should specify the exact location and the quantity of the part to be moved in the correct list from this file. There exists a move file for every day of the week. These move files are filled in by all the logistics employees of each assembly line independently. The deadline for filling in the move file is currently set at 12:00. Everything that is filled in before 12:00 in the move file of the current day will be delivered at the production location on the same day. Replenishments after 12:00 should be filled in in the move file of the next day and will be delivered the next day. Obviously, it is also possible to fill in the move file of a few days later, if the replenishment should be delivered only then.

In principle, only complete lots are moved from the external warehouses to the production location. This means that only complete pallets are moved, since a lot fits on a pallet. This is not a problem for parts which are stored at multiple pallet locations in the buffer inventory. The employees can then decide to replenish the same number of pallets as the number of empty pallet locations for this part. In this way, all empty locations for this part in the buffer get filled up again. It is more difficult for the parts of which only one pallet is kept in the buffer inventory. If they wait with replenishing until this pallet is empty, then they are probably already too late, since there is no part available to cover the

demand during the lead time. That is why they already replenish when there are some parts left on this pallet. They then order a complete lot from the external warehouse. However, this will generally not fit completely on the pallet location in the buffer inventory, as there are still parts left on this pallet. For example, when a pallet location of a part consists of a pallet with 100 boxes and when this part is replenished if there are 20 boxes left, then 20 of the 100 boxes from the new lot cannot be stored in the buffer inventory. Then, this remainder is moved to a new location in the bulk inventory. If the part needs replenishment again, then this remainder is used first. Of course, when the part is replenished from the bulk inventory in the first place, this problem of ordering complete lots does not apply. In that case, the lot is picked from the bulk inventory, used to fill up the buffer inventory and placed back at its place in the bulk inventory if it is not empty yet.

The action of moving a portion of a lot instead of the complete lot is called a partial move. As described in the previous paragraph, this only happens between the buffer inventory and the bulk inventory (both at the production location) and not between an external warehouse and the production location. However, it is possible to partially move from an external warehouse to the production location. It is generally not done because it saves time at the external warehouse since employees are not concerned with counting a certain number of parts and moving only that amount. Also, it is believed that it is inefficient to move one half of a lot in one week and the other half of the lot in the next week as the increase of transport and handling costs will likely be higher than the decrease in holding costs. Therefore, partial moving from external warehouses is not done, even though it is possible. A complete overview of the current movements and whether they are complete lots (C), partial lots (P) or both (C/P) can be found in Figure 2.1.

2.3. Critical parts

In this section, an analysis is made of the parts which should be included in the inventory control system. These parts will be called the critical parts. To this end, the following sub-question is answered in this section: *“What are the critical parts of the A/B series?”*.

To find the critical parts, it is important to understand the different kind of parts and their relation to each other. A basic distinction can be made between C and L parts. C parts can be regarded as complete subassemblies whereas L parts are components which go into a subassembly. The relevant C part families for this project are CM, CL and CD. In Subsection 2.3.1, these C part families will be described one by one in more detail and the corresponding critical C parts will be given. The L parts will be covered when discussing the CD family. Subsequently, an ABC inventory classification will be performed on these L parts to determine the critical L parts in Subsection 2.3.2. This section is concluded with a short conclusion on the selected critical parts in Subsection 2.3.3.

2.3.1. Classification of C parts

CM part family

The CM part family concerns the TR parts which are mounted to the basic machines. Each of the 3 types of machines have their own types of TR parts. Therefore, there are multiple Stock Keeping Units (SKU's) for each machine type. These SKU's, together with a description of each SKU and the type of machine that each SKU belongs to, are given in Table A.1 in Appendix A: Critical parts and values of their relevant characteristics.

TR parts are packaged on pallets in certain quantities. They are replenished from external warehouses or the bulk inventory and are stored in the buffer inventory, from where they are picked for assembly. The only exception are the TR parts for the BL type, which have no fixed position in the buffer inventory, these are stored somewhere in the bulk inventory. At the assembly line, the TR parts are mounted to the basic machines. Some basic machines come in crates, so without TR parts, whereas other basic machines already come on TR parts. In the case of machines which already come

on TR parts, these TR parts might need to be changed when the customer requires other TR parts than the ones which are already mounted. These TR parts are changed with TR parts from the buffer inventory and the TR parts which come from the machine are stored in the buffer inventory again. Since the TR parts are obviously important parts and part of them are replenished from external warehouses, all TR parts should be included in the inventory control system. Therefore, all 18 SKU's from Table A.1 are regarded as critical parts.

CL part family

The CL part family concerns the kits and MD parts. These are options which the customer can order on his machine. A kit is, generally, a carton box with a collection of parts which belong together. Certain combinations of kits form an option. A MD part is an option in itself. The CL parts come on pallets and are replenished from external warehouses or the bulk inventory. They are stored in the buffer inventory and picked from there for assembly. There are a large variety of CL parts which could be ordered by the customer. Therefore, all these CL parts should be available in the buffer inventory. That is why it is important that all these CL parts should be controlled by the inventory control system, i.e. all CL parts are critical parts.

Given the large variety of CL parts, there should be a good overview of all these parts. To this end, the CL part SKU's of the A/B series were filtered from the delivered orders of the past year. This list with delivered orders was retrieved from the WMS and concerned the periods May 2018 up to and including April 2019. It is believed that this covers all the CL parts for the A/B series. It appeared that all these 57 CL parts were currently included in the buffer inventory as well, so this is a double-check that confirms that this are the relevant CL parts. All these 57 CL SKU's are given in Table A.2 in Appendix A: Critical parts and values of their relevant characteristics, with their descriptions and the quantities in which they were used last year.

CD part family

The CD part family indicates how the basic machine should be build up. There are 6 CD numbers, 2 numbers for each basic machine type. These 2 numbers differentiate between a machine which should be build up with a CB part or with a RB part. The 6 CD numbers and their descriptions can be found in Table A.3 in Appendix A: Critical parts and values of their relevant characteristics. A CD number is not a SKU itself, it is a top number which is linked to a specific Bill Of Materials (BOM). There are thus 6 BOM's, one for every specific type. These BOM's include all the parts which are in any case needed, regardless of the TR parts and options ordered by the customer. This are the L parts, which show great differences. On the one hand, some of these L parts are used for multiple types and are thus in multiple BOM's. On the other hand, some of the L parts are BOM-specific, this means that these parts are only compatible with CB parts or RB parts.

You could imagine that not all L parts are equally important for the company, so it does not make sense to include all the L parts in the inventory control system. For example, a CB part which is a complex part that takes a lot of space is more important than a bolt which is small, simple and could be replenished easily in large quantities. To get a view of the L parts which are worthwhile to include in the inventory control system (critical L parts), an inventory classification method is used. A discussion on different inventory classification methods and the best one to use in this case, can be found in Section 3.2. It turns out that the ABC analysis is the most appropriate method to classify the L parts. This analysis is conducted in the next subsection.

2.3.2. ABC inventory classification of L parts

The ABC inventory classification discriminates between parts based on their usage value. The usage value is defined as the usage rate multiplied by the individual value. To do the classification, this

usage value should be calculated for every L SKU from every BOM. The individual value of these SKU's can easily be found in these BOM's. The calculation of the usage rate requires a bit more work. The usage rate of a SKU will be defined as the annual demand of that SKU. To calculate the annual demand of a SKU, the number of delivered orders for a one year period of a CD machine is checked first. Then, the quantity required of this SKU per machine is checked in the BOM of this CD machine. The two numbers are multiplied to find the annual demand of that SKU for that specific CD machine. This calculation is repeated for each CD machine which has this SKU in its BOM. Finally, the total annual demand of the SKU is calculated by adding up these individual annual demands. Then, the usage value is calculated by multiplying the annual demand and the individual value. This approach was used to calculate the usage value for every L SKU. Subsequently, the total usage value was calculated and the individual usage values were expressed as percentages of the total usage value. Also, the annual demand of each SKU was expressed as percentage of total units. These percentages will now be used to determine whether a part is classified as a class A, B or C part.

It appears that the three CB part types account for a very large percentage of the total usage value, while only accounting for a very small percentage of total units. In fact, the three CB part types together had an annual demand of only 840 units, which is about 0.48% of total units. The usage value that this 0.48% of total units represented is about 78.42% of total usage value. Since the three CB part types account for a large percentage of usage value, while accounting for only a small percentage of total units, these three parts will be classified as class A parts. No other parts show this extreme behaviour and for this reason, only the 3 CB part types are classified as class A parts. They are visible in Table A.4 in Appendix A: Critical parts and values of their relevant characteristics. The extreme behaviour of the CB parts is due to the fact that these parts have a much higher individual value than the other L parts, so their usage values become very high, even when their annual demands are less than a lot of the other L parts.

Next, some parts were classified as class B parts. These parts contribute substantially to the total usage value, but do not contribute that much as the class A parts do. It was chosen to classify parts as class B parts when they account for at least 0.20% of total usage value. Furthermore, only L parts which are stored at external warehouses are considered, because only these L parts need replenishment. This means that parts which account for more than 0.20% of total usage value but are not stored at external warehouses, will not be classified as class B parts. To this end, all parts which account for more than 0.20% of total usage value were checked in the WMS on their presence at external warehouses. The parts which were present at external warehouses at the time of checking, were classified as class B parts. These 12 parts are visible in Table A.5 in Appendix A: Critical parts and values of their relevant characteristics. It appears that these are mainly large parts like RB parts and SE parts. These class B parts account for 15.82% of total usage value, while only accounting for 3.98% of total units.

The remaining L parts are classified as class C parts. These parts either account for less than 0.20% of total usage value or are not present at external warehouses or both. This concerns 139 parts, which account for only 5.74% of total usage value while accounting for 95.55% of total units. These parts will not be considered in the rest of the research and will therefore not be listed. A large number of these parts are nuts and bolts, which are needed in large quantities, but are very low in value. These parts are directly supplied to the assembly line and a Two Bin system is already developed for their replenishment. Therefore, it is not a problem that these parts are not included in the research.

The ABC classification is visualized in the graph in Figure 2.2. Here, the percentage of total usage value is plotted against the percentage of total units. The parts of the graph corresponding to the different classes are pointed out. You can see again that the class A parts (the three CB part types)

account for a very large percentage of total usage value, but only a very small percentage of total units. In the graph, this is visible as a steep increasing line. The part of the line concerning the class B items is diminishing in steepness. This is due to the fact that there are already items included which deliver only a small contribution to total usage value. The part of the line concerning the class C items is strongly decreasing in steepness and becomes a flat line. This can be explained from the fact that the class C parts increasingly less contribute to the total usage value. In fact, a large amount of these class C parts have a negligible impact on the total usage, which causes the (almost) flat line.

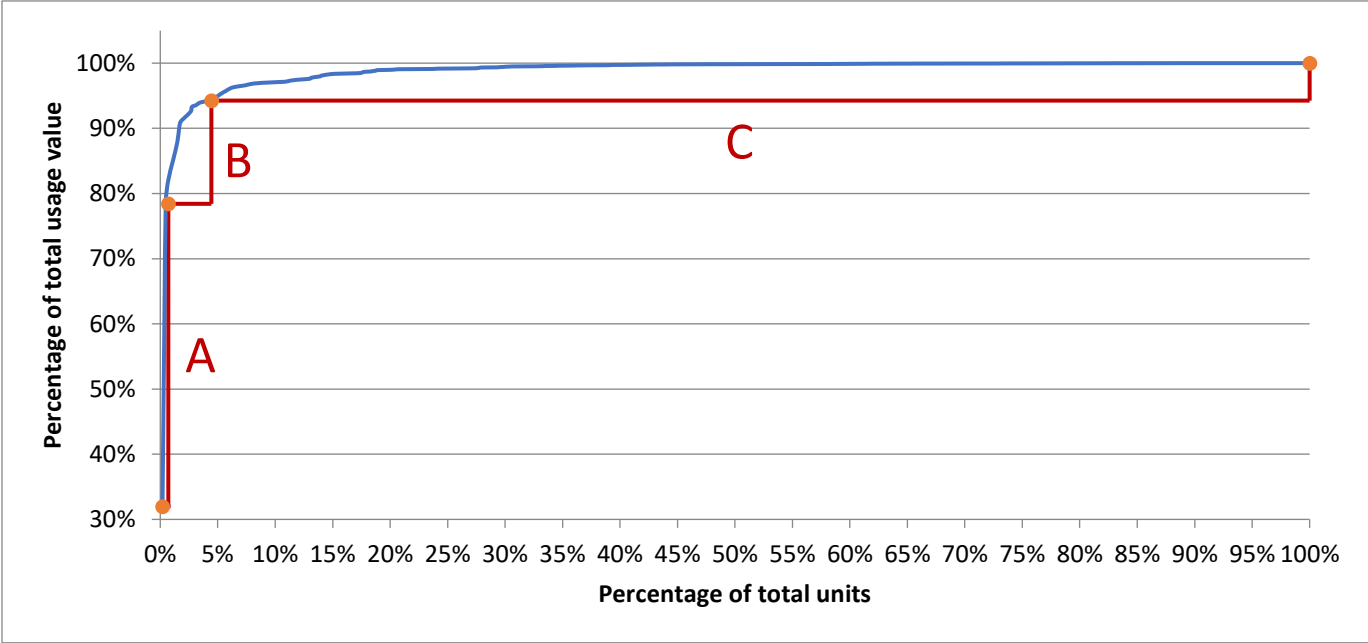


Figure 2.2 ABC classification of the L parts

The conclusion of this ABC classification of L parts is that only class A and class B parts will be taken into account in the inventory control system. In that way, the inventory control system becomes efficient, since 94.26% of total usage value is covered by only dealing with 4.45% of total units. This means that a total of 15 L parts will be taken into account, namely 3 class A parts and 12 class B parts.

2.3.3. Conclusion on the selected critical parts

In conclusion, all C parts and L parts tabulated in Appendix A: Critical parts and values of their relevant characteristics, are going to be included in the inventory control system and will from now on be referred to as the critical parts. This boils down to 90 critical parts. Of these 90 parts, 75 parts are C parts: all the parts from the two product families CM (TR parts) and CL (options). The remaining 15 parts are L parts: 3 class A parts and 12 class B parts from the BOM’s underlying all the CD top numbers. The CD numbers are therefore not critical parts, they are merely a top number indicating which BOM should be used.

2.4. Characteristics of the critical parts

Now that the critical parts have been selected, values of some characteristics of these parts which are relevant for the inventory control system are determined. Therefore, the following sub-question is answered: “What are the values of some relevant characteristics of these critical parts?”. Characteristics considered relevant for the inventory control system are: part commonality, storage locations, lead times, lot sizes and buffer capacity. These characteristics will now be treated in this order in Subsections 2.4.1 through 2.4.5. In these subsections, there will be multiple times referred

to Table A.6 in Appendix A: Critical parts and values of their relevant characteristics. This table contains the values for some of these characteristics for all the critical parts. However, these values provide only a snapshot of the situation from the 29th of May 2019. Some values are not constants and may therefore change over time.

2.4.1. Part commonality

Some critical parts might be used in more than one machine type. To check whether this is the case, the BOM's of all machine types were checked in case of the L parts and lists with C parts per machine type were checked in case of the C parts. The results are visible in the column "Used in types" in Table A.6.

72 of the 90 critical parts are used exclusively in one of the machine types under study (A/BS/BL). Next to that, there are 9 critical parts which are used for both the A and BS. Besides, there are 3 critical parts which are used for both the BS and BL. There are also 3 critical parts which are used in all three series (A, BS and BL). So far, this does not give any problems when devising the inventory control system. This is the case since all these critical parts are only used by the machine types under study. The orders for these three types will be taken into account and therefore all demand for these critical parts will be considered. However, this does not apply for the 3 remaining critical parts. These parts are namely used by both the BL and the D series. Since the D series machine is not included in the scope of the research, part of the demand will not be considered. If the demand for the D series is high and the inventory control system does not take this into account, then stockouts may still occur. It is therefore important to also consider the orders for the D series for these 3 critical parts.

Furthermore, the critical parts can also be ordered by dealers who will do the assembly of machines themselves. However, these orders are not handled at the production location but are fulfilled from one of the external warehouses directly. As a consequence, no replenishments to the production location are required for these orders and they will therefore not be taken into account in the inventory control system.

2.4.2. Storage locations

The next characteristic which will be treated is the location where the critical parts are stored. This data is retrieved from the WMS and is visible in the column "Locations" in Table A.6. In this column, all locations where a specific critical part is stored, have been written down. These locations were determined on a hall level. A location consists of three figures. When it consists of three numbers, the first number indicates the warehouse: numbers 3,4,7 and 8 all indicate a different warehouse. Locations starting with a 7 indicate the production location, while the other numbers indicate external warehouses. The second and third number indicate the hall number, ranging from 01 for hall 1 to 28 for hall 28. The only exception is 7Bx, which indicates a field location outside location 7 (where the movable machines are stored). Furthermore, there are locations starting with a letter instead of a number: G,I and Z. These locations are all at separate external warehouses.

Given this structure of location names, it is easy to make the division of the three inventory types. Buffer inventory is all the inventory located at location 721, since hall 21 is completely dedicated to buffer inventory for the A/B series. Next to that, there is some buffer inventory at locations 706 and 707 for a few SKU's. Bulk inventory is all the inventory at the production location apart from the buffer inventory. This means that every location starting with a 7 and not ending with 21 (or 06/07 for the few exceptions) is considered as bulk inventory. Finally, there is the inventory at external warehouses. This is all the inventory located at locations not starting with a 7. The presence of the critical parts in these three inventory types is made insightful in the columns: "Buffer", "Bulk" and

“Ext.wh.” in Table A.6. A green cell indicates that a critical part is present in the buffer inventory, the bulk inventory and/or an external warehouse respectively.

However, this is only a snapshot of the locations where the critical parts were stored on the 29th of May 2019. This is not a given, the locations change when new parts are supplied by suppliers or when parts are moved from one warehouse to another for example. Nevertheless, it gives a good insight in the arrangement of the critical parts among the different inventory types.

2.4.3. Lead times

The lead time will be defined as the time period between the initiation of a replenishment order and the delivery of that order. The lead time of a part is dependent on the location from where it is replenished. The parts are generally stored at multiple locations, but not always at the same locations. Therefore, we cannot determine one specific lead time for a certain part. However, we can determine lead times from the different locations.

Replenishments from the bulk inventory to the buffer inventory are carried out instantaneously since the bulk inventory and buffer inventory are both stored at the production location. The only actions that are required are driving with a reachtruck to the right location in the bulk inventory, picking the lot, bringing it to the right location in the buffer inventory and filling up the inventory there. These actions will not take long and it is therefore reasonable to assume a zero lead time.

Replenishments from the external warehouses to the production location require more time, naturally. This requires a lot more actions: filling in a move file, arranging a trailer, loading the trailer at the external warehouse, driving to the production location, unload the trailer there and moving the lots to the correct location. Shipments from different warehouses are combined, which means that if half a truckload is required from external warehouse A and half a truckload from external warehouse B, the truck will first visit both locations before stopping at the production location. Therefore, the lead times from all the external warehouses can be considered the same. Currently, the procedure is that when the move file is filled in before 12:00, the parts will arrive at the same day, i.e., before 17:00. The truck will be unloaded at the same day, which means that the parts are available at least at the end of the day. This means that the lead time is a few hours when the order is placed before 12:00. When the order is placed after 12:00, the order is delivered in the afternoon of the next day. The lead time is in that case approximately one day. The lead time from external warehouses is thus dependent on the time that the order is placed. For simplicity, we will always assume a lead time of one day from external warehouses. In that way, the inventory control system will always be on the safe side.

2.4.4. Lot sizes

The next characteristic that will be treated is the lot size of a critical part. The lot size will be defined as the number of parts that a lot of a critical part contains. This information has been retrieved from the WMS and is presented in the column “Lot sizes” in Table A.6. This column contains all the existing lot sizes on all the locations from a certain critical part. However, the buffer inventory has not been taken into account as parts are continuously picked from these lots and they are therefore incomplete.

It appears that some SKU’s have only one specific lot size whereas other SKU’s have a lot of different lot sizes. These different lot sizes are due to the suppliers, who might ship in different boxes/quantities or due to logistics employees, who placed incomplete lots back at the bulk inventory or external warehouses. For these SKU’s, no fixed lot size can be determined.

The lot size is currently relevant for the replenishments from external warehouses to the production location. As can be seen in Figure 2.1, only complete lots are replenished from the external warehouses and therefore the complete lot sizes from there are ordered. However, as mentioned before, it is possible to replenish partial lots from external warehouses. If the optimal replenishment policy requires other lot sizes than complete lots, then it would be feasible. However, we should still take certain box quantities into account for certain SKU's. For example, there are four SE parts in one box, so it would be inconvenient to open the box and move only one SE part from the box. SE parts should therefore be replenished in multiples of four. Nevertheless, if a lot consists of two boxes with four SE parts, then it is not necessary to replenish both boxes.

2.4.5. Buffer capacity

The following characteristic that will be determined is the buffer inventory capacity of the critical parts. As described before, the inventory of parts at the production location consists of both bulk inventory and buffer inventory. The bulk inventory is shared with all parts from all series. Besides, the inbound department determines for a large part which parts are stored in which quantities in this bulk inventory. Therefore, there are no dedicated places or fixed quantities for A/B series parts in this bulk inventory. For this reason, it is not possible to determine the capacity of the bulk inventory regarding the A/B series. However, the capacity of the buffer inventory regarding the A/B series can be determined. There are namely 3 halls (hall 21 completely and a small part of hall 6 and 7) where shelves are dedicated to storing A/B parts as buffer inventory and the capacity of these shelves can be determined. The shelves are on different height levels on top of each other. Only on the lowest level, parts can be picked manually. However, there are separate locations on pick level where the parts from higher levels are stored. These pick locations are supplied by these higher buffer inventory shelf levels. They are also considered as buffer inventory and therefore also taken into account in the buffer capacity. Consequently, all the parts from the buffer inventory are available to pick manually.

All the critical parts are stored in this buffer inventory, with a few exceptions: all TR parts of the BL series (part of CM SKU's) are not stored in this buffer inventory. These parts also do not have a fixed location in the bulk inventory, so their capacities cannot be determined. The rest of the critical parts all have fixed places in the buffer inventory and this buffer inventory is not used by other assembly lines (except for the 3 SKU's shared with the D series). Therefore, the buffer capacities of these critical parts can easily be determined.

The buffer capacity of a SKU will be defined as the maximum quantity of the SKU that can be stored in the buffer inventory. However, it is challenging to determine the maximum buffer capacity. This is the case since there are a lot of different pallet types, each with different dimensions. Not all pallet types go together well on a shelf. They either do not fit together on a shelf or leave large unused spaces. Also, some pallet types can only be stored on certain locations. Therefore, optimizing the buffer inventory utilization will be a challenge in itself, which is not in the scope of this research. Besides, the employees think that, through experience, they have managed to create a fairly optimal design of the buffer inventory. The buffer design of the employees will therefore be respected. This means that there is only limited flexibility in changing the capacity of the SKU's in the buffer inventory.

The current buffer design has been observed and the buffer capacity of a part according to this design has been determined. This has been done for every critical part which is stored in the buffer inventory. This boils down to all critical parts, except for the BL TR parts. The buffer capacity of the critical parts is visible in Table A.7 in Appendix A: Critical parts and values of their relevant characteristics.

Each critical part has one or more pallet locations for its own. These locations are visible under “Individual capacity” in the table. Besides, there are some shared pallets locations: 2 RB part SKU’s (L parts) share pallet locations, 2 CB part SKU’s (L parts) share pallet locations, 4 SE part SKU’s (L parts) share pallet locations, all the A/BS TR part SKU’s (CM parts) share pallet locations and 4 MD part SKU’s (CL parts) share pallet locations. All these shared pallet locations are visible under “Shared capacity” in the table. Furthermore, there are overflow locations where all the SE part SKU’s, all the TR part SKU’s and almost all the CL SKU’s can be stored, in case there are more pallets on hand than the individual and shared capacity can hold. For example, a certain SE part SKU has 2 pallet locations for its own and 30 pallet locations shared with the other SE parts. When all these locations are full and there are still pallets of this certain SE part SKU left, then these pallets are stored at the overflow locations.

At the buffer inventory, multiple types of pallets are stored. The distinction is made between small pallets (SP), of which 3 fit on a shelf, big pallets (BP), of which 2 fit on a shelf and large pallets (LP), of which 1 fits on a shelf. The pallet type is dependent on the SKU. The assumption is made that the SKU is always stored on the same pallet type as observed when determining the capacity. This assumption is for some parts not very reasonable, since it is known for some parts that the pallet type differs from time to time. However, when this gives capacity problems, it is always possible to repack the contents of the unusual pallet type to the pallet type as described here. So in the end, the assumption can be made.

Since there are multiple pallet types possible at the shared locations of the TR parts and the overflow locations and not all pallet types fit together on a shelf, only limited configurations are possible at these locations. All possible configurations are visible in the corresponding cells in the table. For example, one possible configuration at the overflow locations is 8/6/12. This means that 8 large pallets, 6 big pallets and 12 small pallets fit together at the overflow locations. Suppose that one shelf with 3 small pallets becomes empty, than it is possible to fill it with 3 small pallets again or fill it with 2 big pallets or fill it with 1 large pallet. The configuration then changes to 8/6/12, 8/8/9 or 9/6/9 respectively.

2.5. Conclusion on current situation

The current situation regarding the planning and control of the A/B series assembly and the required parts has been analysed in this chapter. To this end, two knowledge questions have been formulated. These knowledge questions have been subdivided into two sub-questions each, which are answered in the previous four sections. The knowledge questions will be answered now by summarizing the main findings from these sections.

The first knowledge question is *“What is the current situation regarding the planning and control of the A/B series assembly line and corresponding buffer inventory?”*.

Concerning the planning of the A/B series assembly, this takes place according to a week planning. This planning is based on a list with orders to be built (outstanding orders), which is updated from time to time. The main criteria on which they decide on which orders to produce is the build week. The number of mechanics is chosen dependent on the production planning. With the current production planning approach, there is only a production schedule on a weekly basis but not on a daily basis. This indefinite way of scheduling orders has as result that it is not exactly known in advance when each order is started. Consequently, the part requirements are not known on a day level.

Regarding the inventory control of the A/B series, inventory of the basic machines and all possible parts which go onto these machines is kept. These parts are stored at external warehouses and at the production location itself. Inventory at the production location can be divided into buffer

inventory and bulk inventory. Buffer inventory is the inventory of the A/B series parts close to the assembly line, which is directly available for production. Bulk inventory is the inventory which is stored far away and often high on the shelves, not easily available for assembly. The control of the inventory at the external warehouses and the bulk inventory is done by the inbound department. The control of the buffer inventory is done by the logistics employees of the A/B series. To determine whether parts need replenishment, they observe the buffer inventory. When they see that a large number of the allocated pallet locations of a part is empty, they decide to replenish that part. They then search for the part in the WMS. According to the company's procedure, they should order the lot with the oldest lot date. When the lot is located in the bulk inventory, the employees can pick this lot themselves and use it to replenish the buffer inventory. When the lot is located in an external warehouse, the employees should fill in a so called "move file". Everything that is filled in before 12:00 in the move file of the current day will be delivered at the production location on the same day. In principle, only complete lots are moved from the external warehouses to the production location. Between the buffer inventory and the bulk inventory (both at the production location) portions of a lot can be moved, which is called a partial move.

The second knowledge question is *"Which parts of the A/B series should be included in the inventory control system and what are the values of their relevant characteristics?"*.

An analysis is made of the parts which should be included in the inventory control system. These parts are called the critical parts. A basic distinction can be made between C and L parts. C parts can be regarded as complete subassemblies whereas L parts are components which go into a subassembly. The relevant C part families are CM, CL and CD. The CM part family concerns the TR parts which are mounted to the basic machines. Each of the 3 types of machines have their own types of TR parts. The CL part family consists of the kits and MD parts. These are options which the customer can order on his machine. A CD number is not a SKU itself, it is a top number which is linked to a specific Bill Of Materials (BOM). These BOM's include all the parts which every machine needs regardless of the TR parts and options ordered by the customer, this are the L parts. Not all L parts are equally important for the company. To determine which L parts should be included in the inventory control system, an ABC inventory classification method has been used. Only the 3 parts classified as class A and the 12 parts classified as class B will be taken into account in the inventory control system. In that way, 94.26% of total usage value is covered by only dealing with 4.45% of total units. All C SKU's and L SKU's tabulated in Appendix A: Critical parts and values of their relevant characteristics, are going to be included in the inventory control system. This boils down to 75 C parts and 15 L parts.

Concerning the characteristics of the critical parts, the following characteristics are determined: part commonality, storage locations, lead times, lot sizes and buffer capacity. Values for the characteristics were found, which are presented in Table A.6. There appeared to be a certain degree of part commonality among the critical parts. 72 of the 90 critical parts are used exclusively for one of the machine types under study. Next to that, 18 critical parts are used by multiple machine types. Three of these parts are used by the D series, which is not a machine type under study. Locations where the critical parts are stored, were retrieved from the WMS. These locations were determined on a hall level. This makes it possible to differentiate between the three inventory types. The lead time is defined as the time period between the initiation of a replenishment order and the delivery of that order. The lead time of a part is dependent on the location from where it is replenished. Replenishments from the bulk inventory to the buffer inventory are carried out instantaneously and a zero lead time can therefore reasonably be assumed. Replenishments from the external warehouses to the production location require more time, but for simplicity we will always assume a lead time of one day. This is always on the safe side. The lot size is defined as the number of parts

that a lot of a critical part contains. This data is also retrieved from the WMS. It appears that some SKU's have only one specific lot size whereas other SKU's have a lot of different lot sizes. The last characteristic that is determined is the buffer capacity of the critical parts. It is defined as the maximum quantity of the SKU that can be stored in the buffer inventory. The buffer design of the employees will thereby be respected. This means that there is only limited flexibility in changing the capacity of the SKU's in the buffer inventory. Each critical part has one or more pallet locations for its own. In addition, there are shared pallets locations where some SKU's share locations. Next to that, there are overflow locations where a lot of different SKU's can be stored in case there are more pallets on hand than the individual and shared capacities of these SKU's can hold.

3. Literature review

This chapter discusses relevant theory from literature. To this end, the theoretical perspective and theoretical framework will be defined in Section 3.1. Subsequently, several elements of the theoretical framework will be treated in different sections. In Section 3.2, the used inventory classification method is discussed. Section 3.3 covers different kinds of demand models by answering the knowledge question: *“What demand models are proposed in inventory management literature?”*. The definition of an inventory control system is given in Section 3.4. Furthermore, a large variety of inventory models that an inventory control system can use, are discussed here by answering the knowledge question: *“What inventory models are used in inventory control systems according to literature?”*. Finally, in Section 3.5 a conclusion on the discussed elements of the theoretical framework is given. Besides, answers to the knowledge questions are provided in this section.

3.1. Theoretical perspective and theoretical framework

Theoretical perspective

The theoretical perspective from which we approach this project is the Operations perspective, since our project is purely focussed on operational processes. Within this Operations perspective, we look from an Inventory Management perspective, both from an Operations Management and Operations Research point of view. The Operations Management perspective is needed as we are going to design an inventory control system. Since Operations Management is mainly concerned with planning and organizing processes in a production environment, this seems to be the appropriate perspective for our project. Furthermore, the Operations Research perspective is taken to make the replenishment decisions. These are complex decision making problems, which require mathematical models to make good decisions. Since Operations Research tries to find optimal or near-optimal solutions to these kind of problems, this perspective also seems to be appropriate.

Theoretical framework and its elements

Given the theoretical perspective defined in the previous paragraph, the developed theoretical framework will now be presented and its elements will be explained. The theoretical framework is visible in Figure 3.1. It consists of rectangles, ovals and arrows. The rectangles represent objects, the ovals represent methods/systems and the arrows denote the relation that the objects have to these methods/systems.

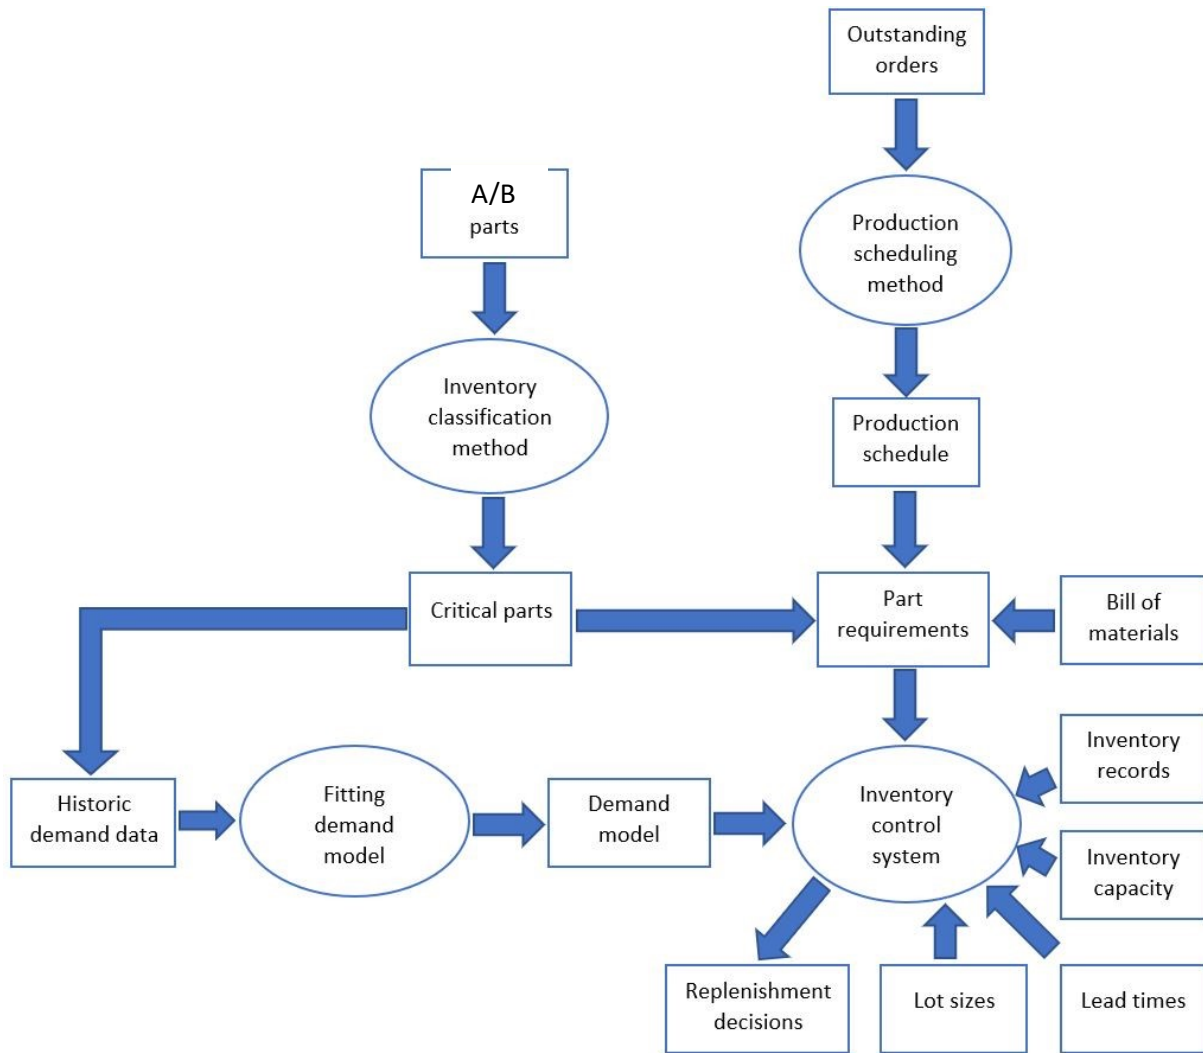


Figure 3.1 The theoretical framework

From the scope of this project, as defined in Section 1.3, we know that an inventory control system is needed for the parts of the A/B series. Therefore, A/B parts form the beginning of the theoretical framework (at the top left). These parts can be found in the company's database. Not all these parts are going to be included in the inventory control system since not all parts are equally important to the company. Only parts with high importance are going to be included. To determine the importance of the parts, an Inventory classification method can be used. A discussion on different inventory classification methods follows in Section 3.2. With the chosen Inventory classification method, the parts which are going to be included in the inventory control system were determined in Subsection 2.3.2. These parts are called the Critical parts. Besides, the Historic demand data of these critical parts need to be collected. These data will be used to determine which Demand model will fit to these data. Different kind of demand models will be covered in Section 3.3. The demand model is eventually an input for the inventory control system.

Another input for the inventory control system will be based on the Outstanding orders, i.e. known orders that need to be built in the short term. The outstanding orders can be retrieved from the company's Warehouse Management System (WMS). A short-term Production schedule of end-products can be made from these outstanding orders using a Production scheduling method. Eventually, Part requirements can be derived from the production schedule. This boils down to: which parts are needed, in which quantities and on which moments. This is also a demand input for

the inventory control system. To come from the production schedule of end-products to the requirements on part level, a Bill of materials is needed. The bill of materials gives information on what parts are required for each end-product. It is a product structure with different 'levels of assembly'. Level 0 assemblies are the end products, level 1 assemblies are the assemblies of which the end-products consist, level 2 assemblies are the sub-assemblies of which level 1 assemblies consist, and so on (Slack, Brandon-Jones, & Johnston, 2016). The part requirements are only needed for the critical parts, so information on which part is critical is also needed to determine the part requirements.

The Inventory control system is the most important system in this theoretical framework. All objects are centred around this system and are, directly or indirectly, an input to the system. The ultimate goal of the inventory control system is to find the Replenishment decisions. These decisions determine what should be ordered and when. This is the output of the inventory control system. Section 3.4 is devoted to the inventory control system.

Several kinds of input are needed into the inventory control system next to the earlier discussed demand model and part requirements. Firstly, Inventory records are needed to determine if parts are sufficiently in stock. Slack et al. (2016) states that three main inventory records should be kept: the item master file, the transaction file and the location file. The first file contains the unique identification code for each part, the second keeps a record of receipts into stock, issues from stock and a running balance and the third describes where inventory is located. All this information is available from the company's WMS. Secondly, the Inventory capacity should be included in the inventory control system. It should not be possible to hold more inventory (according to the system) than is physically possible in the buffer inventory. Therefore, we determined the buffer capacity in Section 2.4. Thirdly, Lead times are required, which are defined as the time periods between placement of replenishments and delivery of replenishments. Lastly, Lot sizes are required, which are defined as the number of parts that a lot of a part contains. Both the lead times and lot sizes were discussed in Section 2.4.

3.2. Inventory classification methods

The first element of the theoretical framework that will be treated is the inventory classification method. The inventory of a company usually consists of a large variety of items. These items are generally not equally important to the company. More important items have more negative consequences if mistakes are made in their inventory control. Therefore, items with high importance require tighter inventory control. An inventory classification method is used to determine the importance of items. Based on the importance, it is decided whether an item will be included in the inventory control system or not.

Various inventory classification methods exist and they all classify based on another criteria. Praveen, Simha & Venkataram (2016) list several inventory classification methods. They make the distinction between single criteria classifications and multi criteria classifications (Praveen, Simha, & Venkataram, 2016). As the names suggest, these classifications discriminate based on a single criteria and multiple criteria respectively. Single criteria classifications are easy to understand and to perform. However, when multiple factors determine the importance of a part, multi criteria classifications are more comprehensive and efficient to use. Since in this project it is for a large part clear which parts are important to control, the inventory classification need not be very extensive. Nevertheless, it will be used to check if our presumptions are right and to determine a boundary of which parts are just important enough to include and which are just not. Therefore, a single criteria classification method will be sufficient.

We will now list some of the single criteria classification methods that Praveen et al. (2016) propose in Table 3.1, together with the criteria on which the methods discriminate.

Table 3.1 Single criteria classification methods according to Praveen et al. (2016)

| Inventory classification method | Criteria |
|--|---|
| ABC analysis | Usage value (usage rate*unit price) |
| XYZ analysis | Variability of the demand |
| VED (Vital Essential Desirable) analysis | Criticality |
| FSN (Fast Slow Non-moving) analysis | Usage rate |
| HML (High Medium Low) analysis | Unit price |
| SDE (Scarce Difficult Easy) analysis | Lead time of procurement |
| SOS (Seasonal Off-Season) analysis | Seasonality in procurement prices |
| MUSIC-3D (Multi-Unit Selective Inventory Control) analysis | Combination of usage value, lead time and criticality |

In this project, there is no seasonality in procurement prices and therefore the SOS analysis will not be a good candidate. Furthermore, the lead times of procurement are very small for all parts, so a classification based on lead time is also not appropriate. Therefore, the SDE and MUSIC-3D analyses are not considered. There is difference in the variability of the demand for different SKU's. However, it is chosen to let this factor not decide whether a part should be included in the inventory control system, so the XYZ analysis will also not be used. On the other hand, an inventory classification based on the criticality of parts is meaningful in this project. However, when using the VED analysis, criticality is not a quantifiable criteria. One or more persons should therefore determine for each part whether it is vital, essential or only desirable. This is an exhaustive but above all subjective task, so we rather choose a classification method based on a quantitative criteria.

The remaining candidates are therefore the ABC, FSN and HML analyses. These inventory classification methods use the criteria usage rate (FSN), unit price (HML) or a combination of both (ABC). The usage rate is an important criteria in this project, since a stockout of an often-used part causes more backorders than a stockout of a part used only a few times a year. Besides, the unit price is also an important criteria in this project, since there is a large variety in the complexity and size of the parts. There are parts which are large and already partly assembled and there are parts as simple as a nut or bolt. It is obviously more important to keep enough of these large and complex parts as they are not as interchangeable as small and simple parts. The unit price reflects the complexity and size well and is therefore a factor that should be taken into account next to the usage rate. The ABC inventory classification uses the product of these factors (the usage value) as criteria. Items with either a high usage rate or a high unit price or both are given a high importance. Therefore, we conclude that the ABC inventory classification is the best one to use in this project. The additional benefit is that the usage rates and unit prices are easily available from the data.

The ABC inventory classification is also the most common inventory classification method and is often mentioned in the literature. Slack et al. (2016) also propose the ABC classification to "discriminate between different stocked items, so that a degree of control can be applied to each item which is appropriate to its importance". The ABC inventory classification thus discriminates between parts based on their usage value, which is defined as the usage rate multiplied by unit price. It appears that, according to this classification, a small proportion of the parts will generally account for a large proportion of the total usage value. This phenomenon is according to the Pareto law. This is also called the 80-20 rule, because typically 80 per cent of the sales are accounted for by only 20

per cent of the item types. The division of Class A, B and C items is then typically as follows, with the items ranked from highest to lowest usage value (Slack et al. (2016)):

- Class A: the highest 20% of items which account for 80% of total usage value
- Class B: the next 30% of items which account for 10% of total usage value
- Class C: the lowest 50% of items which account for 10% of total usage value

3.3. Demand models

In the next section on the inventory control system, the demand is needed as input to the inventory models. The demand can be modelled in different ways and some models are very specific in this. It is therefore important to have prior knowledge regarding different ways to model demand. For this reason, the following knowledge question is treated in this subsection: *“What demand models are proposed in inventory management literature?”*.

The first distinction that can be found in literature, is whether demand is modelled as a deterministic or as a stochastic process. Deterministic demand means that the demand realisations (output) are only dependent on the demand parameters (input). In other words, there is no randomness involved in the realisation of demand. No matter how often you repeat the demand process, the demand realisations stay the same if the input is not changed. This means that you can completely predict the future demand if you know the deterministic demand model and its inputs. On the other hand, under stochastic demand you do not know the future demand with certainty since there is randomness involved in the demand model. This means that there are stochastic variations in the demand realisations. Given certain demand parameters as input, the demand realisations vary in a certain range if you repeat the demand process.

In the real world, there is often uncertainty involved in the realisation of demand. Therefore, the assumption of stochastic demand is often more realistic. However, inventory models that assume deterministic demand are often simpler to use and easier to understand. Besides, the assumption of deterministic demand turns out to be very reasonable in general (Axsäter, 2006). Therefore, it is meaningful to consider inventory models that assume deterministic demand and models that assume stochastic demand. Both types of demand models will be discussed further now in different subsections.

3.3.1. Deterministic demand models

When demand is assumed to be deterministic, a further assumption can be made regarding the variability of the demand. The demand can be assumed constant, i.e. it is assumed that demand does not (significantly) change over time. For example, the yearly demand is a known figure. The monthly demand is then assumed to be one twelfth of the yearly demand. It is the most simple form of deterministic demand: you just need one demand rate and assume that it applies to all periods.

If it is not reasonable to model the deterministic demand at a constant rate, it can be modelled as time-varying demand. This means that the demand rate can vary from period to period. However, these variations in demand are precisely known in advance (it is still deterministic). We therefore need to know the demand for a finite number of successive, future periods. The demand during such periods can be either continuous with time or only at discrete points in time with equal distance. The former means that demand is a stream of small demands during the whole period. The latter means that demand comes in large quantities only at certain times during the whole period. However, the inventory models for this type of demand, which will be discussed in the next section, can handle both demand that is continuous and discrete with time. The only input that is needed, is the total

demand in a period. It is common to assume a constant demand rate during the period, which only changes from period to period (Silver, Pyke, & Thomas, 2017).

3.3.2. Stochastic demand models

Whereas the distinction between constant and time-varying demand was made with deterministic demand, this distinction will not be made with stochastic demand. We will only consider the case where the average demand is approximately constant and we will not cover time-varying stochastic demand.

Stochastic demand is most commonly modelled by a probability distribution (the demand distribution). A probability distribution accounts for the random nature of stochastic demand by assigning probabilities of occurrence to different demand outcomes. Generally, there are two classes of probability distributions: discrete and continuous probability distributions. A discrete probability distribution gives discrete (countable) outcomes and a continuous probability distribution gives outcomes in a continuous (uncountable) range. In the real world, demand during a certain time is generally a nonnegative integer. According to Axsäter (2006), it is natural to use a discrete demand model since it reflects the real demand better. However, this is on the condition that demand is reasonably low. If the demand is relatively high, it is better to approximate the real demand by a continuous demand model from a practical point of view. Although Axsäter (2006) makes this proposal, he does not give a clear rule for what should be considered a high or low demand. Silver et al. (2017) also suggest to use a discrete distribution if the demand is low and a continuous distribution when the demand is high. However, they provide a rule of thumb: If the average demand during the lead time is lower than 10 units, then use a discrete distribution to model stochastic demand and otherwise use a continuous distribution. The average demand in a lead time should be taken if a continuous review inventory model is considered. If a periodic review inventory model is considered, then the average demand in a review interval plus a lead time should be taken. We will now present some discrete and continuous demand distributions which are commonly used to model demand.

3.3.2.1 Discrete demand distributions

The discrete demand distributions that will be discussed are the Poisson distribution and the compound Poisson distribution. Within the class of compound Poisson distributions, we will only discuss the logarithmic compounding distribution.

Poisson distribution

Silver et al. (2017) propose to use a Poisson distribution to model the demand as a stochastic process. This means that customers arrive according to a Poisson process with a so-called intensity λ , which is the rate at which customers arrive (number of customers per unit time). As a consequence, the number of customers in a time interval of length t has a Poisson distribution, which is described by Equation 3.1.

$$P(k) = P(D(t) = k) = \frac{(\lambda t)^k}{k!} e^{-\lambda t} \quad \text{for } k = 0, 1, 2, \dots \quad (3.1)$$

Where $P(k)$ is the probability of k customers in time interval t . When we assume that the demand follows a (pure) Poisson distribution, the demand is equal to the number of customers. This means that Equation 3.1 also describes the probability of total demand k in time interval t ($P(D(t)=k)$), so Equation 3.1 gives the Poisson demand distribution. This distribution has as only parameter λt , which is equal to both the mean (μ) and variance (σ^2) of the demand during time t . Silver et al. (2017) note that due to this property, the Poisson distribution should only be used to model demand which has in reality a mean and variance which are quite close to each other. They give an operational definition

of 'quite close': the variance should be within 10% of the mean. Axsäter (2006) supports the reasoning of Silver et al. (2017) by stating that the Poisson distribution is reasonable when the ratio $\frac{\sigma^2}{\mu}$ is approximately equal to 1. He also provides the same operational definition: use the Poisson distribution when $0.9 \leq \frac{\sigma^2}{\mu} \leq 1.1$. However, it is according to Axsäter (2006) also quite common to use the Poisson distribution to model the demand when the ratio $\frac{\sigma^2}{\mu}$ is smaller than 0.9, even though this results in an overestimation of the variance. This is due to the assumption under the Poisson distribution that the standard deviation is equal to the mean. However, when the ratio is small, the standard deviation is smaller than the mean in reality, hence the overestimation in the model.

Compound Poisson distribution

If the ratio $\frac{\sigma^2}{\mu}$ is larger than 1.1, more difficult probability distributions are used. Silver et al. (2017) propose the negative binomial distribution or a compound Poisson distribution. Axsäter (2006) is more specific and proposes a negative binomial distribution, which is in fact a form of a compound Poisson distribution (with a logarithmic compounding distribution). We will explain these concepts further now.

In the case of a Poisson demand distribution, we have assumed that all individual customer demands had a size of 1 and therefore the number of customers was equal to the total demand. When demand is modelled as a compound Poisson distribution, this assumption no longer holds. Now, the size of a customer demand is a stochastic variable in itself. This variable is independent of other customer demands and also independent of the distribution of the customer arrivals. Since the demand size is a stochastic variable, its outcomes are described by a probability distribution. This specific distribution is called the compounding distribution. A wide range of (discrete) probability distributions can be used as compounding distribution. However, customers still arrive according to a Poisson process, so the number of customers and thus the number of demands is still described by a Poisson distribution. Hence, the compound Poisson demand distribution is given by Equation 3.2.

$$P(D(t) = j) = \sum_{k=0}^{\infty} \frac{(\lambda t)^k}{k!} e^{-\lambda t} f_j^k \quad \text{for } j = 0, 1, 2, \dots \quad (3.2)$$

Where $P(D(t)=j)$ is as before the probability of total demand j during time interval t . You can also recognize the Poisson distribution from Equation 3.1, which gives the probability of k customers during time interval t (just as before). This probability is multiplied by f_j^k , the probability that k customers give total demand j . f_j^k is related to the demand size and can be calculated recursively from the compounding distribution using the relation given by Equation 3.3.

$$f_j^k = \sum_{i=k-1}^{j-1} f_i^{k-1} f_{j-i} \quad \text{for } k = 2, 3, 4, \dots \quad (3.3)$$

Where $f_0^0 = 1$, $f_j^1 = f_j$ and f_j is the probability of demand size j ($j = 1, 2, \dots$). The f_j 's are then determined by the compounding distribution and form the start of the recursion.

As mentioned before, a compound Poisson distribution can make use of a wide range of compounding distributions. We will only discuss the use of a logarithmic distribution as compounding distribution as it is proposed by Axsäter (2006) to use this distribution when $\frac{\sigma^2}{\mu} > 1.1$ (according to his decision rule).

The demand size has then a logarithmic distribution. The probability f_j of demand size j is in that case given by Equation 3.4.

$$f_j = -\frac{\alpha^j}{\ln(1-\alpha)j} \quad \text{for } j = 1, 2, 3, \dots \quad (3.4)$$

Note that a demand size of zero is not possible. However, it is still possible to have zero demand in a time interval as it is possible to have zero demand arrivals in a time interval (see Equation 3.2). The distribution has as single parameter α , which should take a value between zero and one. If we know the mean (μ) and variance (σ^2) of the demand during time t , we can determine α according to Equation 3.5.

$$\alpha = 1 - \frac{\mu}{\sigma^2} \quad (3.5)$$

Subsequently, the intensity λ of the Poisson process can be calculated with Equation 3.6.

$$\lambda = -\mu \frac{(1-\alpha) \ln(1-\alpha)}{t\alpha} \quad (3.6)$$

Please refer to Axsäter (2006) for the derivations of Equations 3.5 and 3.6.

So based on the μ and σ^2 of the demand, we can determine the only parameter α of the compounding distribution and on the basis of that we can determine the only parameter λ of the Poisson process. We then have all information needed to fully specify the applicable compound Poisson distribution to model the demand. The corresponding probabilities can be calculated by using Equations 3.2 and 3.3.

However, Axsäter (2006) presents a simpler way to calculate the probabilities. It turns out that when the compounding distribution is logarithmic, the demand during time t ($D(t)$) has a negative binomial distribution. Instead of using Equations 3.2 and 3.3, the probabilities of the demand distribution can be determined with the much simpler function given by Equation 3.7.

$$P(D(t) = k) = \begin{cases} (1-p)^r & \text{for } k = 0 \\ \frac{r(r+1) \dots (r+k-1)}{k!} (1-p)^r p^k & \text{for } k = 1, 2, \dots \end{cases} \quad (3.7)$$

Where the parameter p is a number between zero and one and the parameter r can be any positive number. The values for p and r can be calculated directly from the μ and σ^2 of the demand by Equations 3.8 and 3.9 respectively.

$$p = 1 - \frac{\mu}{\sigma^2} \quad (3.8)$$

$$r = \mu \frac{(1-p)}{p} \quad (3.9)$$

Note that the parameter p of the negative binomial distribution is equal to the parameter α of the logarithmic compounding distribution (see Equation 3.5). When the parameters p and r have been determined, the probability of a demand equal to k during time t , can be calculated using Equation 3.7. This probability is equal to the corresponding probability calculated with Equation 3.2. However, it is computationally much faster to use Equation 3.7 as only one relatively simple function needs to be calculated instead of multiple summations as is the case with Equation 3.2 and 3.3. The computational efficiency is especially present when the probability of large demand outcomes need

to be calculated. However, we have seen before that a discrete distribution should only be used when the demand is relatively low.

To conclude, we will give another representation of the compound Poisson distribution which might be more insightful. We have seen that when a compound Poisson distribution is used to model demand, it makes use of two probability distributions: the Poisson distribution to model the demand *arrivals* and a compounding distribution to model the demand *sizes*. Consider a fixed time interval in which demand occurs. Let:

- X be the total demand during the time interval
- N be the number of demands that arrive in the time interval
- Y_n be the size of the n th demand, where Y_1, Y_2, Y_3, \dots are independent and identically distributed

If N has a Poisson distribution and the Y_n 's have a (discrete) compounding distribution, then

$$X = \sum_{n=1}^N Y_n \quad (3.10)$$

where X has a compound Poisson distribution.

As we discussed above, if the Y_n 's follow a logarithmic distribution (Equation 3.4), then X follows a negative binomial distribution (Equation 3.7). If the Y_n 's follow a different compounding distribution, then the distribution of X is described by Equation 3.2. Nevertheless, N always follows a Poisson distribution and is thus described by Equation 3.1.

3.3.2.2 Continuous demand distributions

If the average demand is relatively high, it is more common to use a continuous demand distribution to model demand. This is computationally efficient and therefore more practical. The continuous demand distributions that will be discussed are the normal distribution and the gamma distribution.

Normal distribution

The normal distribution is without a doubt the most prominent continuous distribution to model the demand. This has multiple reasons according to Axsäter (2006). First, according to the central limit theorem, the sum of many independent random variables has approximately a normal distribution, even if the individual random variables do not follow a normal distribution. As the demand normally comes from multiple independent customers, it is reasonable to model this demand with a normal distribution. Besides, when you consider the demand over a long time period, discrete demand (which follows for example a distribution from the previous subsection) can also be approximated with a normal distribution. A third reason to use the normal distribution is because it is easy to use and common in practice.

The normal distribution is bell-shaped, as can be seen as the black line in Figure 3.2. However, the shape of this bell (the distribution) depend on the parameters. The normal distribution has two parameters: mean (μ) and variance (σ^2). If you know the mean and variance of the demand, then you can find the specific normal distribution that fits this demand by using this mean and variance as parameters. Probabilities for this unique normal distribution can easily be found by transforming it to a standard normal distribution of which the probabilities are easily available.

However, the normal distribution has two properties which do not (always) represent real demand well. The first one is the fact that there always is a (small) probability of negative demand, which is obviously not realistic. This is due to the left tail of the distribution (see Figure 3.2), which always

crosses the zero point on the x-axis. The second property is the symmetry around the mean of the normal distribution (see Figure 3.2). It might not be reasonable to assume that demand of e.g. more than one standard deviation higher than the mean is equally likely as demand of more than one standard deviation lower than the mean. In the case that this is not equally likely, a skewed distribution can better be used to model the demand.

Gamma distribution

A continuous probability distribution that address both these problems is the gamma distribution. It is only defined for nonnegative values, so demand will never be negative. Furthermore, the gamma distribution is skewed to the right. This means that the probability of very high demands is relatively large. This is in contrast to a symmetric distribution (e.g. a normal distribution). See also Figure 3.2 for an example of a gamma distribution and its properties. The gamma distribution has two parameters: r and λ . Both parameters can be determined when the mean (μ) and variance (σ^2) of the demand are known, using Equations 3.11 and 3.12.

$$r = \left(\frac{\mu}{\sigma}\right)^2 \tag{3.11}$$

$$\lambda = \frac{\mu}{\sigma^2} \tag{3.12}$$

When these parameters have been determined, the gamma distribution that fits the demand can be specified. Unfortunately, the distribution function of the gamma distribution cannot be expressed in closed form, but it is included in Excel. Hence, the probabilities of the gamma distribution can be calculated via Excel.

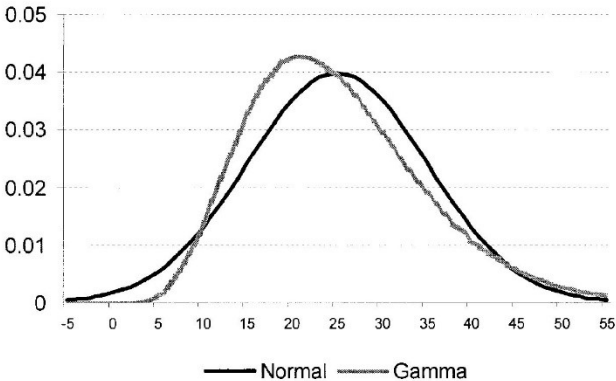


Figure 3.2 Density functions of the normal and gamma distribution, based on the same mean (25) and standard deviation (10) of demand (Axsäter (2006))

We described that it is common to use the normal distribution as continuous demand distribution. Besides, we described the shortcomings of the normal distribution in modelling demand and proposed the gamma distribution instead. We did not yet discuss a decision rule to choose between the two distributions. Silver et al. (2017) base this decision on the coefficient of variation (CV): $\frac{\sigma}{\mu}$. This ratio has influence on the shape of the distribution. If the CV is small, than a lot of the probability mass is centred around the mean and less at extremes. In the case of a normal distribution, this means a high and narrow peak of the normal curve and short tails. This means that the probability of demands close to the mean is high and the probability of demands that are very high or low is small. The gamma distribution with a low CV shows the same characteristics. When the CV is large, the probability mass is more spread out over the whole range of possible demands. In the case of a normal distribution, this means that the probabilities of both very low and very high

demand become higher with the same rate (because of the symmetry). However, since the normal distribution takes on negative values, the probability of negative demand becomes higher, which is undesirable. On the contrary, the gamma distribution does not show this behaviour when the CV is large. It becomes more skewed to the right, which means that the probability of very high demand becomes larger. However, since the gamma distribution is not defined for negative values, negative demand is still not possible. These demand characteristics are more desirable. Therefore, we can conclude that the gamma distribution is the better alternative for a high CV. Silver et al. (2017) set the boundary on 0.5: if the CV is less than 0.5, then the normal distribution is appropriate. If the CV is greater than 0.5, then consider the use of a gamma distribution. The probability of negative demand under a normal distribution is then at least 2.3%.

3.4. Inventory control system

The main goal of this project is to devise an inventory control system. Axsäter (2006) defines its purpose as follows: “The purpose of an inventory control system is to determine *when* and *how much* to order. These decisions should be based on the stock situation, the anticipated demand and different cost factors.”

The stock situation does not only include the stock that is physically available (on hand). It is also important to consider the outstanding replenishment orders, which are going to replenish the stock once they arrive. Furthermore, backorders need to be considered. These are orders which have not been delivered yet. Therefore, the stock situation is described by the inventory position, see Equation 3.13 (Axsäter (2006)).

$$\text{inventory position} = \text{stock on hand} + \text{outstanding orders} - \text{backorders} \quad (3.13)$$

Apart from the stock situation, the anticipated demand should also be taken into account by the inventory control system. This anticipated demand is often based on forecasts and often follows a certain distribution. An elaboration on different types of demand distributions can be found in Section 3.3.

Lastly, the inventory control system bases its decisions on different cost factors. The most common costs in inventory models are listed here:

- *Holding costs*, which are all costs related to holding the inventory and are variable with the inventory level. This is mainly the capital cost, but also consist of, for example, the costs of storage space, handling, damage, obsolescence, insurance and taxes.
- *Ordering costs* (also called setup costs), which are all costs associated with a replenishment. These are usually fixed costs such as transportation and material handling costs. However, also administrative costs associated with the handling of the orders themselves could be considered.
- *Shortage costs*, which occur when an item is demanded but cannot be delivered due to a shortage. These costs depend on whether the order is backlogged or the sale is lost. In the case of backlogging, extra administrative costs or costs due to price discounts can occur. In the case of lost sales, the financial benefit that the sale would have is lost. In both cases, there is usually a loss of goodwill which leads to lost sales in the future. In a production environment, such as in this project, a shortage can cause a chain of negative consequences with accompanying costs. However, a lot of the aforementioned costs are (very) difficult to estimate. For this reason, a service level constraint is often used instead of shortage costs. It can also be difficult to determine a suitable service level, but in practice it turns out to be simpler.

One last comment will be made on the cost factors. Whereas the ordering decisions were based on the inventory position, the holding and shortage costs depend on the so-called inventory level, see Equation 3.14 (Axsäter 2006).

$$\text{inventory level} = \text{stock on hand} - \text{backorders} \quad (3.14)$$

It is important to consider the environment in which the inventory control system will work. Different inventory models were developed for different environments. The books of Axsäter (2006) and Silver et al. (2017) both make the distinction between single-echelon systems and multi-echelon systems. A single-echelon system is a system that is only concerned with the control of inventory at a single location. On the other hand, a multi-echelon system is concerned with the control of inventories at multiple locations. These inventories are coupled to each other in different stages (echelons), for example inventory at a central warehouse and multiple inventories at retailers. In order to get efficient control of all these inventories, it is important to use an inventory model which takes the connection between different inventories into account, i.e. a multi-echelon inventory control system.

Given the scope of this research, as described in Section 1.3, we will confine ourselves to single-echelon systems for the rest of this research. Within the class of single-echelon systems, a further distinction can be made between systems with individual control of single items and systems with coordinated control of multiple items. Of course, individual control of single items could be done in any case. The coordinated control of multiple items through coordination of replenishments makes sense in one of the following cases according to Silver et al. (2017):

- Several items are purchased from the same supplier
- Several items share the same mode of transportation
- Several items are produced on the same piece of equipment

The replenishments on which I will focus in this research take place between different warehouses of Mainfreight and can therefore be considered as ‘purchases’ from the same supplier. Also, the items included in the replenishment are all transported in the same trailer(s). The replenishments under consideration thus satisfy the first two conditions. Therefore, it is meaningful to consider coordinated control of multiple items. Nevertheless, it is still important to consider the individual control of items as it is still a possibility. Besides, as we will see later, coordinated inventory control models make sometimes use of inventory models for single items.

Recall that the inventory control system is concerned with determining *when* and *how much* to order. For these replenishment decisions an inventory model is used. We will now discuss a large variety of inventory models by answering the knowledge question: “*What inventory models are used in inventory control systems according to literature?*”. We split the discussion up in two parts. First we describe the most common single item inventory models in Subsection 3.4.1. Subsequently, we will turn to the coordinated multi-item inventory models in Subsection 3.4.2.

3.4.1. Single item inventory models

This subsection covers single item inventory models. A single item inventory model makes the replenishment decisions for each item independently. As mentioned before, an inventory model needs a demand model as input. We discussed a variety of demand models in the previous section. An important distinction is the one between deterministic and stochastic demand models. Therefore, this subsection is divided in single item inventory models that consider deterministic demand and models that consider stochastic demand.

3.4.1.1 Single item inventory models under deterministic demand

This subsection covers the following single item inventory models under deterministic demand: Basic Economic Order Quantity (EOQ) model and Dynamic lot-size models.

Basic Economic Order Quantity (EOQ) model

The most common deterministic single item inventory model and probably the most well-known model in inventory management is the classical Economic Order Quantity (EOQ) model. This model was already derived by Harris (1913), but its importance and practicality persist to this day, according to Axsäter (2006). Essentially, the model minimizes the Total Relevant Costs (TRC) per unit time by choosing an order quantity, see Equation 3.15.

$$TRC(Q) = \frac{KD}{Q} + \frac{QH}{2} \quad (3.15)$$

Where:

- Q is the replenishment order quantity in units
- D is the demand rate of the item in units per unit time
- K is the fixed ordering cost per order
- H is the holding cost per unit per unit time

The first term in this formula corresponds to the ordering costs and the second term corresponds to the holding costs. The TRC function is convex in Q and the optimal Q can therefore be obtained by setting the first order derivative of Equation 3.15 to 0 and solving for Q . This boils down to solving Equation 3.16.

$$Q = \sqrt{\frac{2KD}{H}} \quad (3.16)$$

The corresponding Q is then the optimal order quantity. Furthermore, it can be shown that the TRC reaches its minimum when the order costs and holding costs are equal. The immediate result of the EOQ model is thus *how much* should be ordered. However, it follows directly *when* the EOQ should be ordered, namely when the inventory is zero. This becomes clear when looking at the graph with the stock level, as can be seen in Figure 3.3.




Figure 3.3 Stock level according to the basic EOQ model (Axsäter (2006))

The graph shows a typical sawtooth pattern. Every tooth corresponds to one replenishment cycle and has length Q/d . When the stock level hits zero, the EOQ is ordered and a new replenishment cycle starts. The graph shows a repetitive pattern: the teeth all have the same height and same length.

However, this repetitive pattern only shows up because the EOQ model is based on quite a list of assumptions, some of which are quite strong. Silver et al. (2017) list the following assumptions:

- A constant and deterministic demand rate
- The order quantity can be non-integer and does not have a minimum or maximum restriction
- The unit variable cost does not depend on the order quantity
- The costs do not change significantly with time
- The items are treated independently of each other
- The replenishment lead time is zero
- Shortages are not allowed
- The complete order quantity is delivered at the same time
- A very long planning horizon is considered

As Silver et al. (2017) also point out, all these assumptions can be relaxed. A simple example is the assumption of a zero replenishment lead time. When this assumption is not realistic, the model can easily be changed to the case of a nonzero replenishment lead time, provided that the lead time is known and constant and the demand remains deterministic. Since the costs are unchanged, the EOQ does not change. However, the moment *when* the EOQ is ordered, changes. It is not when the inventory hits zero, but when the inventory level is equal to the demand during the lead time. In this way, the replenishment order arrives when the inventory hits zero. As a consequence, the stock level still follows the pattern of the graph of Figure 3.3.

Dynamic lot-size models

For the case where the constant demand rate assumption is violated, separate models covering time-varying demand, are developed. These are dynamic lot-size models. These models still consider deterministic demand, so the variations in demand are known in advance. Given the variations in demand over time, there is no longer a constant optimal lot size such as in the basic EOQ model.

The classical dynamic lot-size problem is then as follows. A finite number of periods in which the demand is known, is considered. Orders are placed to fulfil demand for one or more periods. The objective is then to choose the order quantities such that the sum of the ordering and holding costs is minimized. The optimal solution should satisfy two properties according to Axsäter (2006):

1. A replenishment must always cover the demand in an integer number of consecutive periods
2. The holding cost for a period demand should never exceed the ordering cost

A consequence of the first property is that a delivery should only occur when the inventory is zero. Besides, the optimal solution has a quite simple structure.

It is possible to solve the dynamic lot-size problem exactly. Most commonly, this is achieved by using Dynamic Programming. Wagner and Whitin (1958) first proposed this so-called Wagner-Whitin algorithm (Silver et al. (2017)). If such exact methods are impractical or computationally inefficient, simpler heuristics can give an approximate solution. A well-known heuristic is the Silver-Meal heuristic, which considers the average per period costs. A quantity equal to the demand for the first n periods is ordered, if including the demand for period $n + 1$ gives higher average per period costs. Another heuristic uses the principle of the EOQ model that the holding costs and ordering costs are equal in the optimal solution. A quantity equal to the demand for the first n periods is ordered, if including the demand for period $n + 1$ lets the holding costs exceed the ordering costs. However, there exist many more dynamic lot-sizing heuristics.

3.4.1.2 Single item inventory models under stochastic demand

In the previous subsection, single item inventory models assuming deterministic demand were covered. These models give order quantities as output that exactly fulfil the demand that you give as input to these models. This is preferable in the case of deterministic demand since you assume that demand does not show stochastic variations. It is therefore optimal to order exactly what is needed. However, it also means that these models are not concerned with stockouts and their associated costs. In practice, there is often randomness involved in the demand, i.e. stochastic demand, which means that stockouts are possible.

To mitigate the risk of stockouts safety stock is kept. Silver et al. (2017) defines safety stock as “the average inventory level just before a replenishment arrives. A positive safety stock provides a cushion or buffer against larger-than-average demand during the effective replenishment lead time.” There is a trade-off to be made regarding the size of the safety stock. If the demand during the effective replenishment lead time is so large that even the safety stock cannot cover it, then a stockout occurs. On the other hand, when this demand is lower than expected, excess inventory is carried. Both events are not desirable and therefore need to be balanced. Two common approaches to achieve this are:

- An approach based on minimizing cost. This approach specifies certain shortage costs. The objective is then to minimize the total cost, which includes the shortage costs.
Several types of shortage costs exist, for example:
 - Fixed cost (B_1) per stockout occasion
 - Fractional charge (B_2) per unit short
 - Fractional charge (B_3) per unit short per unit time
 - Charge (B_4) per customer line item short
- An approach based on customer service. This approach specifies a certain service level for a certain service measure. The objective is then to minimize total cost, subject to the specified service level constraint.
Several service measures exist, for example:
 - Probability (P_1) of no stockout per replenishment cycle
 - Fraction (P_2) of demand to be satisfied routinely from available inventory
 - Fraction of time (P_3) during which net stock is positive
 - Average time (TBS) between stockout occasions

We will now turn to the inventory models themselves. The models presented here can use one of the above described approaches to determine the amount of safety stock. The four most common single item inventory models under stochastic demand according to Silver et al. (2017) are visible in Table 3.2. Subsequently, they will be treated one by one in detail using theory as presented in Silver et al. (2017).

Table 3.2 The most common single item stochastic inventory models according to Silver et al. (2017)

| | | When to order | |
|-------------------|-------------------|-------------------|----------------------|
| | | Continuous review | Periodic review |
| How much to order | Fixed quantity | (s, Q) | - |
| | Variable quantity | (s, S) | (R, S) (R, s, S) |

All four inventory models make use of the inventory position. The first distinction that can be made between the models is whether the inventory position is reviewed continuously or periodically. In a

continuous review model, the stock level is always known. In a periodic review model, the stock level is only reviewed at certain moments. The time interval between two consecutive review moments is called the review interval R . Both review forms have their advantages and disadvantages, these will become clear when the specific models are discussed. The second distinction that can be made is whether a fixed quantity is ordered or a variable quantity is ordered. In a fixed order quantity model, the order quantity Q is independent of the current inventory position. In a variable quantity model, the order quantity depends on the current inventory position. Namely, the order quantity will be the difference between the current inventory position and the order-up-to-level S . The arguments for choosing a fixed or variable order quantity will also be discussed at the particular model.

(s,Q) model

The (s,Q) model is a continuous review model where a fixed quantity Q is ordered whenever the inventory position reaches the reorder point s or lower. Since the inventory is continuously reviewed, a replenishment order is immediately placed if demand lets the inventory drop from a level higher than s to exactly level s or to a level below s . In the latter case, the difference between the attained level below s and s is called the undershoot. This undershoot can vary in size, depending on the demand transaction size that triggered the replenishment. As a consequence, the inventory position after placing the replenishment is not fixed.

An advantage of the (s,Q) model is that it is simple to understand and thus mistake will less likely happen. Besides, if a replenishment takes place, its quantity is known, so the supplier knows what to expect. The main disadvantage is that the (s,Q) model will not work in an environment where transaction sizes are large and replenishment quantities are small. In that case, large undershoots can occur and the replenishment quantity might not be enough to raise the inventory above the reorder point.

(s,S) model

The (s,S) model is also a continuous review model where a replenishment order is placed whenever the inventory position drops to the reorder point s or lower. However, in contrast to the (s,Q) model, here the inventory position after placing the replenishment is fixed. This is called the order-up-to-level S . This means that when the inventory position drops to or below s , we order up to level S . As a consequence, the order quantity is variable and depends on the undershoot. This is not the case if all demand transactions have size one. In that case, a replenishment is always triggered exactly at inventory level s and the order quantity is always equal to size $S - s$. The (s,S) model is then identical to the (s,Q) model with $Q = S - s$. An advantage of the (s,S) system is that the inventory position is always between the levels s and S , except when an undershoot just occurred. Furthermore, it can be shown that the best (s,S) model does not have larger total costs than the best (s,Q) model. However, it is harder to find the best parameters of the (s,S) model. Besides, the suppliers do not face a predictable, fixed order quantity when the (s,S) model is used.

(R,S) model

The (R,S) model is a periodic review model where every R units of time (the review interval) enough is ordered to raise the inventory position up to the order-up-to-level S . This periodic review aspect has advantages and disadvantages. Since the inventory position is checked only at certain moments and not continuously, less time and money needs to be spent on reviewing. However, the foremost advantage of a periodic review model is that replenishments can be coordinated. This is the case since replenishments can only happen at every R units of time and not immediately as with a continuous model. This comes in handy when multiple items are ordered from the same supplier, since they can be combined in one order, which provides significant reduction in the order costs. We will make use of this benefit when we cover the coordinated multi-item inventory models. Besides, a

periodic review model gives a rhythmic replenishment pattern, which is easier to manage and makes the level of workload easier to predict. Another advantage is that the order-up-to-level S can be adjusted at every R to cope with a changing demand pattern. However, there are also disadvantages with periodic review models. The most important disadvantage is that more safety stock is required, which leads to higher holding costs. This is due to the fact that the period over which safety protection is required, is longer under periodic review. This is because the inventory position can drop significantly between two review moments without the possibility to immediately react to that. Furthermore, the order quantities are also variable here, which is less convenient for the suppliers.

(R,s,S) model

The (R,s,S) model is a combination of the (s,S) and the (R,S) model. Every R units of time the inventory position is reviewed (just like with the (R,S) model). When the inventory position is equal to or lower than the reorder point s , enough is ordered to raise the inventory to the order-up-to-level S (just like with the (s,S) model). Therefore, the order size is variable with this model as well. Scarf (1960) has shown that under some general assumptions the best (R,s,S) model achieves the lowest total costs in comparison to other models (Silver et al. (2017)). However, it is very difficult to find the best values of the parameters in comparison to other models. This complexity results partly from the fact that undershoots of the reorder point are possible even if all demand transactions have size one. This is due to the fact that the inventory is only reviewed every R units of time. This does not happen with a (R,S) model, since there is no reorder point there (there is always an order if the demand is nonzero).

All four models use parameters to control the inventory. Every parameter of a certain model should be determined in order for the model to work. These parameters could be determined sequentially or simultaneously, where the latter is more complex but gives better results (due to the dependence of the parameters). In an (s,Q) model, it is common practice to determine reorder point s and order quantity Q sequentially. First, The Q is calculated according to the EOQ model. Then, the reorder point s is determined given the predetermined Q . The reorder point s should be equal to the demand during the replenishment lead time plus a possible safety stock. This safety stock should be related to the lead time demand. It turns out that the (R,S) model is equivalent to the (s,Q) model, if the transformations from Table 3.3 are made (Silver et al. (2017)).

Table 3.3 The transformations from a (s,Q) to a (R,S) model

| (s,Q) model | (R,S) model |
|---------------|---------------|
| s | S |
| Q | DR |
| L | $R + L$ |

If these substitutions are made in decision rules for the parameters of the (s,Q) model, the parameters R and S can be determined. Whereas the Q in a (s,Q) model was predetermined, here the review period (R) should be predetermined. It could be determined as the EOQ expressed as time supply. The substitutions result from the fact that, in a (R,S) model, the safety stock should protect against a shortage in the replenishment lead time (L) and the review period (R). This is due to the fact that if an order is placed now, the next chance to place an order is only after time R later and that order will be received only after time L . So protection is needed over the time period $R + L$. This is in contrast to the (s,Q) model, where only the replenishment lead time (L) should be protected. The parameters in (s,S) and (R,s,S) models do not show such simple analogies with the (s,Q) system and will not be covered here.

3.4.2. Coordinated multi-item inventory models

We are now going to elaborate on inventory models concerned with the coordinated control of multiple items. This means that the inventories of individual items are not controlled independently anymore. These inventory models are concerned with the so-called Joint Replenishment Problem (JRP): the problem of determining an inventory replenishment policy where multiple items are ordered from a single supplier. Silver et al. (2017) list multiple advantages and disadvantages of coordinated control.

Advantages include:

- Savings on unit purchase costs: This is the case when a supplier gives certain quantity discounts when an order has a certain size. It might be economical to jointly replenish multiple items to make use of these quantity discounts. However, these quantity discounts do not exist in this project.
- Savings on unit transportation costs: This is the case when multiple items are shipped with the same form of transportation. Savings could be achieved by ordering in full truckloads or containers. This form of savings is definitely relevant in this project, since replenishments from external warehouses take place in the same truck. The transportation costs can be reduced by transporting fuller trucks less frequently.
- Savings on ordering costs: If there are costs involved in placing a replenishment order, savings could be achieved when multiple items are replenished in the same order. However, there are no costs involved in placing a replenishment order in this project. The order costs solely exist of transportation and handling costs.
- Ease of scheduling: Replenishments can be controlled per supplier instead of per individual item. This makes it easier to schedule the replenishments and to manage the workload resulting from receiving and inspecting the goods. The company will certainly benefit from this, since a clear replenishment schedule does not exist at present.

Disadvantages include:

- A possible increase in the average inventory level: Under coordinated inventory control, some items will be replenished earlier than if they were controlled by an individual inventory policy. This means that the average inventory level of these items will increase, just like the associated holding costs.
- An increase in system control costs: Coordinated control of items is more complex than the control of individual items. This likely results in higher costs of, for example, reviewing the system and computing the parameters.
- A reduction in flexibility: Since the inventory of an item is not controlled individually anymore, it is harder to deal with unforeseen circumstances.

In the remainder of this subsection, some examples of coordinated multi-item inventory models will be given. Again, the distinction will be made between models that consider deterministic demand and models that consider stochastic demand.

3.4.2.1 Coordinated multi-item inventory models under deterministic demand

This subsection describes coordinated multi-item inventory models under deterministic demand. First, it covers a model with replenishment cycles using EOQ time supply. Then, Dynamic-demand joint replenishment models are discussed. Finally, Rolling horizon systems are described in this subsection.

Replenishment cycles using EOQ time supply

The first model that will be described, is proposed by Silver et al. (2017). This model assumes constant, deterministic demand and makes use of the EOQ expressed as time supply. Whereas the EOQ is the optimal *quantity* to order, the EOQ time supply is the *time period* in which the EOQ gets used up. The EOQ time supply is given by Equation 3.17, with the variables K, D, H as before.

$$T_{EOQ} = \sqrt{\frac{2K}{DH}} \quad (3.17)$$

It also means that the assumptions of the EOQ model, as described before, also apply to this model. The only exception is the assumption of the independent treatment of items.

Furthermore, the order costs will be split in two types of order cost. The first is the major order cost, the fixed cost of placing an order independent of the number of items in the order. The second is the minor order cost for item i , which will be incurred if item i is included in the order. This minor order cost can differ from one item to another.

The idea of the model is to determine a time interval (T) between successive replenishments. Next to that, for each item i , an integer m_i is determined. This is the number of time intervals (T) that the order quantity of item i will last. After this time period of $m_i * T$, the inventory should just hit zero and this item is then replenished again. In this way, replenishment cycles of length T arise, where some items are ordered in every replenishment cycle ($m_i = 1$) and some items have replenishment cycles with lengths as multiples of T ($m_i > 1$). The objective is to choose the T and the set of m_i 's in such a way that the total costs are minimized. Several procedures exist for finding the best m_i 's. Subsequently, the best T can be found using Equation 3.18.

$$T(m_1, \dots, m_n) = \sqrt{\frac{2(K + \sum \frac{k_i}{m_i})}{\sum m_i D_i H_i}} \quad (3.18)$$

Where:

- n is the number of items under coordinated control
- K is the fixed major order cost per order
- k_i is the minor order cost when item i is included in the order
- D_i is the demand rate of item i in units per unit time
- H_i is the holding cost for item i per unit per unit time

In this formula, the EOQ time supply formula is clearly visible. Once the replenishment cycles are known, the replenishment quantities (Q_i 's) automatically follow. This is due to the fact that demand is assumed to be deterministic and constant. The quantities are then given by Equation 3.19.

$$Q_i = m_i T D_i \text{ for } i = 1, 2, \dots, n \quad (3.19)$$

Dynamic-demand joint replenishment models

For the case where demand can be assumed deterministic but not constant, other models have been developed. These models are concerned with the Dynamic-Demand Joint Replenishment Problem (DJRP). According to Boctor, Laporte and Renaud (2004), in the DJRP n items must be replenished to satisfy the demand at T different periods. The demand of each item for each period is known for all T periods of the planning horizon. However, these demands can vary with time, i.e. the demand of an item for one period can significantly differ from the demand of the same item for another period.

Again, a major and minor order cost are assumed, next to holding costs per unit per period. Each order consists of a subset of item types. The DJRP is solved by determining the order quantities of each item in each period, such that the sum of order and holding costs over the whole planning horizon are minimized (Boctor, Laporte, & Renaud, 2004).

The optimal solution of the DJRP has the following properties (Boctor et al. (2004)):

1. It does not occur in an optimal solution that a replenishment quantity is ordered in the current period when the ending inventory of the previous period is greater than zero.
2. The optimal replenishment quantity is equal to the demand of the current period or the sum of the demands of multiple, successive, future periods. (The same property as the first property of the dynamic lot-size optimal solution.)
3. The optimal inventory level at the end of the period is equal to zero or the demand of the next period or the demands of multiple, successive, future periods.
4. It is not optimal to replenish the demand of item type i for period q at the beginning of period t , if the holding costs of that demand from period t up to period q are higher than the order costs of that item in period q . (The same property as the second property of the dynamic lot-size optimal solution.)

The DJRP can be solved by exact algorithms. Boctor et al. (2004) subdivide these algorithms in four categories: dynamic programming, branch-and-bound, branch-and-cut and Dantzig-Wolfe decomposition. However, these exact algorithms can only handle small sized problems. For larger sized problems, heuristics are used. Boctor et al. (2004) describe six classical heuristics and came up with a new heuristic that outperforms them all. This is a local search procedure which tries to escape from a local optimum by making a small modification to a feasible solution. They propose to obtain a feasible solution by means of the Fogarty and Barringer heuristic followed by the Silver-Kelle improvement heuristic (two of the classical heuristics).

The Fogarty and Barringer heuristic is one of the earliest and best heuristics. It makes the additional requirement that when a replenishment is made, it should cover precisely all demand until the next replenishment. This makes the problem simpler and the solution can be found by solving the forward dynamic program with the recursion formula as in Equation 3.20.

$$f_t = \min_{q \leq t} \{f_{q-1} + c_{qt}\} \quad (3.20)$$

Where:

- f_t is the cost of the optimal solution for the first t periods only (to cover the demands of the first t periods)
- c_{qt} = the total cost of a replenishment made at period q that covers the demand of all item types from period q to period t

The Silver-Kelle improvement heuristic tries to improve the solution obtained by the dynamic program. It consecutively checks for every item type ordered at any given period whether a cost saving could be achieved by including that order quantity in the previous replenishment. Subsequently, the new heuristic of Boctor et al. (2004) tries to improve the solution even further.

Rolling horizon systems

A short comment will be made on the use of rolling horizon systems. The models concerning dynamic-demand joint replenishment problem make use of a finite number of periods, where the demand in each period is known. However, the horizon over which the replenishment decisions

should be made, is infinite. In order to make optimal decisions, it is important to consider a finite horizon which is long enough to be a good approximation. However, it could be difficult to determine the demand in such a long term. Therefore, it is common in practice to use the rolling horizon concept for dynamic-demand joint replenishment problems (Sahin, Narayanan, & Robinson, 2013). Besides, it is also commonly used in practice for (single item) dynamic lot-size problems (Axsäter (2006)).

Sahin et al. (2013) describe the concept of a rolling horizon as follows. The idea is to iteratively solve a series of linked short-term stationary problems for which the demand is fairly well known. As time passes, only a certain part of the first replenishment decisions is executed. Then, after a certain time, the demand information is updated, a new short-term stationary problem is solved and new replenishment decisions are made accordingly. These replenishment decisions might interfere with the previous decisions. Therefore, a so-called frozen interval and free interval are introduced. Only orders inside the free interval may be rescheduled.

The objective is twofold: minimise the total schedule cost and minimise schedule instability (resulting from schedule changes due to updated demand data). According to Sahin et al. (2013), the heuristic from Boctor et al. (2004) gives the lowest schedule costs in rolling horizon systems. At the same time, the schedule stability is relatively high.

3.4.2.2 Coordinated multi-item inventory models under stochastic demand

Coordinated control of multiple items is more complex when the demand is stochastic. Recall that we have seen four single item inventory models under stochastic demand: the (s,Q) , (s,S) , (R,S) and (R,s,S) model. Individual replenishments of items were triggered by either a reorder point ((s,Q) and (s,S) models) or an inventory review ((R,S) model) or a combination of both ((R,s,S) model). Under coordinated control, it is no longer the case that an item is always replenished at its reorder point. This is due to the fact that another item can trigger a replenishment in which the former item is also ordered. From Silver et al. (2017) it becomes clear that this complicates the problem in two ways. First, the average inventory level of an item is more difficult to determine. Second, the service levels given certain reorder points are more difficult to compute. Nevertheless, coordinated multi-item inventory models under stochastic demand are developed, some of them will be explained now.

(S,c,s) model (can-order model)

A well-known model is the (S,c,s) model, also called the can-order model, first proposed by Balintfy (1964), according to Silver et al. (2017). It is a continuous review system which works as follows: if the inventory of an item drops to the must-order point s or lower, a replenishment is triggered with a size such that the order-up-to-level S is achieved (just like a (s,S) model). In addition, there is a can-order point c specified for every item, which is higher than the must-order point. The idea is that if a replenishment is triggered by one item's must-order point, then all other items with inventory positions below their corresponding can-order points, are also included in the replenishment and ordered up to their corresponding order-up-to-levels. The benefit of this model is that it combines items which almost need replenishment into one replenishment, instead of a few replenishments just after each other. On the other hand, an item is not included if replenishing that item is not efficient (far from reorder point). The downside of the (S,c,s) model is that it is difficult to compute optimal values for the parameters for every item. Besides, it has been shown by Ignall (1969) that can-order policies need not be the optimal policy (Silver et al. (2017)).

Periodic review joint replenishment models

Given the shortcomings of the continuous (S,c,s) inventory model, several periodic review joint replenishment models were developed. One of these models was proposed by Atkins & Iyogun

(1988). Their model performs better than the (S,c,s) model as the major order cost increases, up to an improvement of 20% (Atkins & Iyogun, 1988). Besides, it is easier to compute and likely easier from a practical point of view due to the periodic review property. The principle of this model is similar to the deterministic replenishment cycle model (as described before) and it again uses the EOQ time supply. The model assumes stochastic, discrete demand generated by independent Poisson processes. Other assumptions are as before: a constant lead time, a major order cost per order, a minor order cost per item and holding costs per item per unit time. Furthermore, the computations are based on proportional shortage costs. The idea is to devise (R,S) models for every item, where coordination is achieved by setting the R 's as some integer multiple of a base period. The procedure is to first calculate all the values for the R 's and do this as if it were a deterministic problem. First all the EOQ time supplies are calculated using only the minor order cost. Recall that this is the expected runout time in which the EOQ gets used up. Then start to allocate a part of the major order cost to the item with the shortest runout time. This lengthens the runout time. Allocate as much so that the runout time is equal to that of the second shortest runout time. Then start allocating to both these items until the third shortest runout time is reached, whereby the time supplies of the first two items should be balanced. Continue in this fashion until the complete major order cost is allocated. The result is a set of the same (shortest) runout times (for items that got allocated a part of the major order cost) and a set of different, (longer) runout times for the other items. The former runout time is called the base period and the associated items are called the base set. The items in the base set get the base period as R and are thus ordered every base period. The other items have a greater runout time than the base period and get the nearest multiple of the base period as R . These items are thus ordered in not every base period. Now that the values for the R 's have been determined, the values for the S 's need to be determined and here the stochastic element comes in. These are calculated as if they were part of (stochastic) single item (R,S) inventory models, where the R is predetermined. Atkins & Iyogun (1988) determine the values of the R 's such that the expected holding and shortage costs are minimised.

Atkins & Iyogun (1988) conclude the following: "Natural elementary extensions to this work are to compound Poisson demand and to systems working under service-level constraints". These extensions are both provided by the model of Fung, Ma & Lau (2001). Their model makes use of stochastic, discrete demand generated by independent compound Poisson processes. Instead of shortage costs they consider the service measure P1: the probability that the cycle ends with no backorders (Fung, Ma, & Lau, 2001). However, the assumptions of the lead time, major and minor order cost and holding costs are the same as in the model of Atkins & Iyogun (1988). The principle of the model is also the same: devise individual (R,S) policies for every item, where some items are ordered every base period and some items are ordered in certain multiples of the base period. However, the approach to find the best (R,S) parameters differs. In the model of Fung et al. (2001), a non-linear mixed-integer programming problem is formulated to find the optimal parameters. Since it is non-linear, it is too difficult to solve directly. Nevertheless, Fung et al. (2001) describe a heuristic approach, consisting of two algorithms, to find the near-optimal (R,S) parameters that solve the programming problem. The first algorithm assumes a certain base period and as a consequence the non-linear mixed-integer program becomes equivalent to a subproblem for every item. These subproblems can be solved by the second algorithm, so that the locally minimising R and S values for every item can be found for this base period. This approach is repeated for a range of base period values, which should be determined in advance by the user. The first algorithm stops just before the expected costs increase for the first time or otherwise when the end of the range is reached. It has then found the optimal base period and the corresponding (R,S) parameters for every item have been determined by the other algorithm. Fung et al. (2001) compared their model with the can-order

model and found that their model performed better as the lead time increases. However, the models perform equally well on average when the lead time is zero.

The last periodic review joint replenishment model that will be discussed, is proposed by Viswanathan (1997). In his model, individual (R,s,S) models are devised for each item and a common R is found (Viswanathan, 1997). This is in contrast to the models of Atkins & Iyogun (1988) and Fung et al. (2001), which apply individual (R,S) models for every item with different R 's. In other words, their models do not check every item at every R , but always order when an item is checked, whereas Viswanathan's model does check every item at every R , but not always orders every item at every R . The assumptions regarding the demand, lead time, order costs, holding costs and shortage costs are the same as in the model of Atkins & Iyogun (1988). The approach of Viswanathan (1997) to find the individual (R,s,S) policies and the common R is as follows. First, find the optimal individual (s,S) policies for each item for a fixed review period r , using only the minor order cost. This is a complex and time-consuming operation, although it could be efficiently done using the algorithm of Zheng & Federgruen (1991) (Silver et al. (2017)). The initial value of r that is used, is the base period as found by the model of Atkins & Iyogun (1988). Second, the cost of the (R,s,S) policies with the base period can be found with Equation 3.21, where the r should be equal to the base period.

$$C(r) = \frac{K}{r} + \sum_i C_i(r) \quad (3.21)$$

Where $C_i(r)$ is the cost per unit time of the optimal (s,S) policy for item i , when a review period of r is used. Here also the major order cost comes in, this cost is incurred whenever a review is carried out. This is not quite correct since the major order cost are not incurred if no order is placed at a review moment. Therefore, the actual cost is less than the one according to Equation 3.21, although it has been verified that the difference is negligible. Now that the cost for the initial value of r has been found, the search for the best value of r starts. Evaluate Equation 3.21 for values of r in small timesteps from the initial value in both directions, until the minimum cost has been found. The corresponding r will be the R of the (R,s,S) policies. The parameters s and S then follow from the optimal (s,S) policies given a review period R . Viswanathan (1997) found that his model performs slightly better than the model by Atkins & Iyogun (1988). In addition, the former model is able to dominate other models (such as the can-order model) in more problem instances than the latter model. However, the downside is that the required computational effort is substantially larger.

3.5. Conclusion on the literature review

Relevant literature was reviewed in this chapter. In order to determine which literature is relevant, a theoretical perspective and theoretical framework were defined in the first section. We found that both an Operations Management and an Operations Research perspective fit this project best. Furthermore, we defined the theoretical framework of Figure 3.1. In this framework, the inventory control system is the key system where all other objects are centred around and which directly or indirectly form an input to the system. The only output of the inventory control system are the replenishment decisions, the ultimate goal of the project.

The second section covered inventory classification methods. An inventory classification was needed in the analysis of the current situation to determine which parts should be included in the inventory control system. Since important parts require tighter inventory control than less important parts, it is useful to classify the parts based on importance. An inventory classification does this based on a single or on multiple criteria. We discussed several single criteria classification methods and the criteria on which they decide. It turned out that the most commonly used method, the ABC classification, is the most appropriate classification method for this project. This classification uses

the product of usage rate and unit price (the usage value) as criteria. Both these factors reflect the importance of a part in this case.

In the third section, different kinds of demand models were treated. The knowledge question that was proposed there, “*What demand models are proposed in inventory management literature?*”, will therefore be answered now. A basic distinction can be made between deterministic and stochastic demand models. Deterministic demand models can be subdivided in models that assume constant demand and models that assume time-varying demand. Regarding stochastic demand models, we only looked at models assuming approximately constant average demand. These models are commonly modelled by a probability distribution. We discussed both discrete and continuous demand distributions and a decision rule to decide between a discrete or continuous distribution. Besides, a rule for deciding between the (discrete) Poisson and Compound Poisson distribution has been given. Furthermore, we discussed a rule for deciding between the (continuous) Normal and Gamma distribution. An overview of the proposed demand models and decision rules can be found in Table 3.4.

Table 3.4 Overview of the discussed demand models with the decision rules in red

| Deterministic demand models | | Stochastic demand models | |
|--|--|--|--|
| Constant demand | Time-varying demand | Constant average demand | |
| One constant demand rate for all periods | Constant demand rate during period and changing demand rates between periods | Discrete distribution $D_{L(+R)} < 10$ | Continuous distribution $D_{L(+R)} \geq 10$ |
| | | Poisson $\frac{\sigma^2}{\mu} \leq 1.1$ | Normal $CV < 0.5$ |
| | | Compound Poisson $\frac{\sigma^2}{\mu} > 1.1$ | Gamma $CV \geq 0.5$ |

The subject of the fourth section is the inventory control system. The inventory control system and relevant concepts are defined there, where the focus is on single-echelon systems only. An inventory control system uses an inventory model to determine the replenishment decisions. These inventory models can be divided in models with individual control of single items and models with coordinated control of multi items. We have seen that it is useful to explore both types of models. This has been done by means of the knowledge question “*What inventory models are used in inventory control systems according to literature?*”, which will be answered now. We have seen that the demand model is an important input to the inventory model and we have seen the basic distinction between deterministic and stochastic demand models. Therefore, this distinction is also useful for classifying the inventory models. Both the single item inventory models and the coordinated multi-item inventory models were subdivided in deterministic and stochastic variants. The deterministic models were further subdivided in models assuming constant and time-varying deterministic demand. The stochastic models were further subdivided in models under continuous and periodic review. Inventory models for each class of this classification were studied and described in the fourth section. An overview of these models is given in Figure 3.4, where the blue boxes give the classification as described above. The white boxes underneath give the discussed models that have these properties.

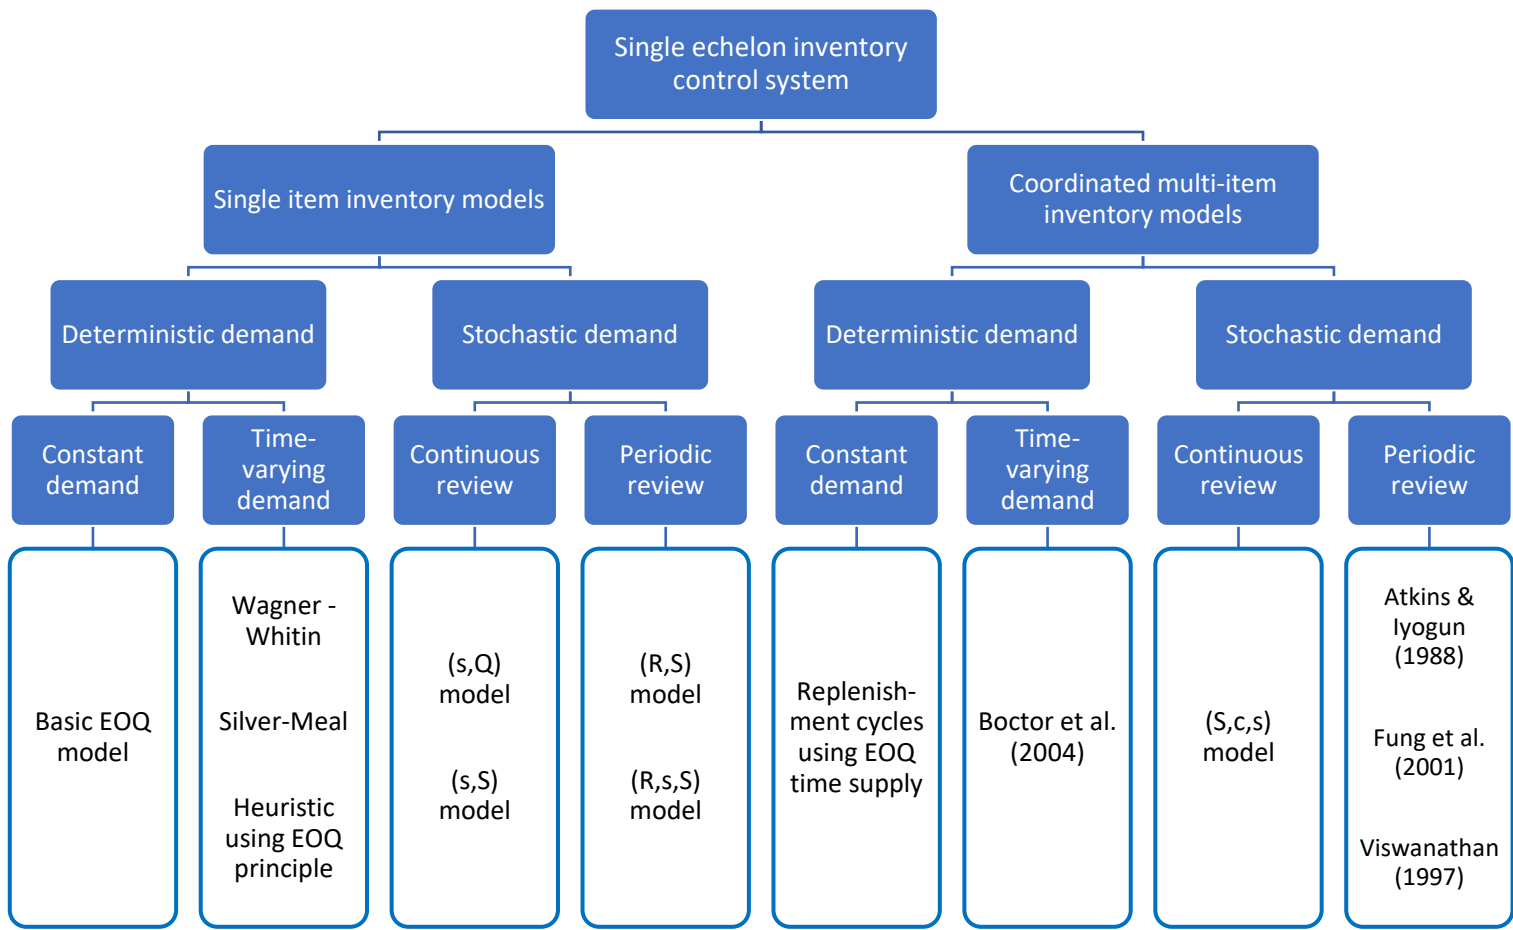


Figure 3.4 Overview of the discussed inventory models

4. Inventory control system methodology

This chapter covers the inventory control system that will be proposed to the company. In Section 4.1, the foundation of the inventory control system will be chosen by choosing the most appropriate demand model and inventory model. The methodology behind the inventory control system will then be described in more detail in Section 4.2: the demand model, inventory model and extensions of the inventory control system will all be described in detail in this section. In Section 4.3, the implementation of the inventory control system (as described in Section 4.2) will be treated. The inventory control system has been implemented in Excel, where the model behind the inventory control system is developed in VBA. Finally, Section 4.4 contains the determination of values for the parameters of the inventory control system.

4.1. Choice of the inventory control system

In order to devise an appropriate inventory control system, we need to choose an appropriate demand model and an appropriate inventory model. Therefore, this section is subdivided in the choice of the demand model and the choice of the inventory model. Both choices will be based on the literature from Chapter 3.

4.1.1. Choice of the demand model

It is important to first decide on the most appropriate demand model. The demand model is namely an input to the inventory model, so the demand model has influence on the choice of the inventory model. Therefore, the most appropriate demand model will be chosen now by means of Table 3.4 from the conclusion of the literature review (Section 3.5).

The first consideration that needs to be made according to this table, is whether a deterministic or stochastic demand model should be chosen. We know from the problem identification and analysis of the current situation that it should be possible to make a production schedule based on outstanding orders. When the company adheres to this predetermined schedule, the demand for parts can in principle be assumed as deterministic. However, this schedule could only be made for the short term, i.e. for something like one/two week(s). In order to make optimal decisions using a deterministic model, it is important to know the demand for a long planning horizon. We have seen that this problem can be solved by using a rolling horizon. Nevertheless, it is not likely that this approach is going to work in this project. This results from the 'truncated horizon effect' (Sahin et al. (2013)). This effect occurs when the planning horizon is too short to find the optimal early replenishment decisions that would have been found if the demand beyond the planning horizon had been known. This is due to the fact that the weekly demand of a lot of critical parts is very low. Consider for example a weekly demand equal to 2 in the first week and equal to 3 in the second week. If a rolling horizon of one week is chosen, then it is optimal to order 2 units at the start of the week (and have zero inventory at the end of the week). However, when the demand for two weeks had been known immediately it would likely be optimal to order 5 units at the start of the first week. Since this will likely happen for a lot of the critical parts, it does not seem appropriate to use a deterministic model. On the other hand, a stochastic demand model can be used to model (uncertain) demand beyond the planning horizon, since there is randomness involved in the demand model. In this way, replenishment decisions can be made based on demand over a longer period, which likely gives more optimal decisions than when using a completely deterministic demand model. We will therefore only consider stochastic demand models from now on.

Stochastic demand is commonly modelled by a probability distribution. Several demand distributions that are proposed in literature, were discussed in the literature review (Chapter 3), together with decision rules. The first consideration that needs to be made regarding the demand distributions, is

whether a discrete or continuous probability distribution should be chosen. We have seen that when the average demand during the (replenishment) lead time D_L is lower than 10 units, then a discrete distribution is advised. This is in the case of a continuous review inventory model. In the case of a periodic review inventory model, the average demand during a review interval plus a lead time D_{R+L} should be taken. We have seen in the analysis of the current situation (Chapter 2) that the replenishment lead time L is simply assumed to be one day, since that will always be enough. Obviously, we do not know yet whether a periodic review inventory model will be used, nor do we know what the length of the review interval R will be in that case. However, from the average week demand of last year, we can calculate the average demand for different time intervals and draw conclusions from that. We see in Table 4.1 that for short time intervals, a lot of SKU's should be modelled with a discrete demand distribution according to the decision rule from literature.

Table 4.1 Percentage of SKU's which should use a discrete distribution according to literature, for different time intervals.

| Length of $R + L$ | Percentage of SKU's with $D_{L+R} < 10$ |
|-------------------|---|
| 1 week | 94.62% |
| 2 weeks | 72.04% |
| 1 month | 53.76% |
| 2 months | 33.33% |

It will therefore make more sense to use a discrete demand distribution as stochastic demand model. Even for the (small) portion of parts for which a continuous distribution is advised, a discrete distribution can be used. This is the case since the continuous distribution provides an approximation of the deterministic distribution in that case. Thus, a deterministic distribution will be more precise, but it will make it computationally harder and therefore less practical. However, since it applies to only a small portion of parts, it is expected that the impact will be manageable.

The next consideration is which discrete demand distribution should be used. We have seen that the Poisson and logarithmic compound Poisson distribution have been proposed. The Poisson distribution is advised when the ratio $\frac{\sigma^2}{\mu} \leq 1.1$ and the compound Poisson distribution with logarithmic compounding distribution is advised otherwise. Note that it is not even possible to use the logarithmic compound Poisson distribution to model demand of SKU's with a ratio $\frac{\sigma^2}{\mu} \leq 1$. In that case, there will exist negative probabilities and also the sum of the probabilities will be larger than 1. Since this should not be possible, the logarithmic compound Poisson distribution will in any case not be used for these SKU's. However, it is possible to use the Poisson distribution in this case and that is also the suggested distribution according to the decision rule. For the SKU's with a ratio $\frac{\sigma^2}{\mu} > 1$, both the Poisson and logarithmic compound Poisson distribution can be used to model the demand. We will stick to the decision rule to decide which of the two distributions will be used in that case: The ratio $\frac{\sigma^2}{\mu}$ has been estimated for every SKU using the sample mean and sample variance of the week demand of last year. It appears that 90% of the critical parts have a ratio $\frac{\sigma^2}{\mu}$ higher than 1.1. This means that it is best to use a logarithmic compound Poisson distribution for these 90% of the critical parts and a Poisson distribution for the remaining critical parts. We will now test this presumption by doing goodness of fit tests for some SKU's of both categories.

Goodness of fit

We will test whether the suggested probability distribution provides a good fit with the historic demand data. To this end, a comparison between the observed frequencies from the past and

expected frequencies according to the suggested probability distribution has been carried out for some critical parts. If possible, the other probability distribution was also included in the comparison. Histograms were made from the historic demand in Excel to find the observed frequencies. The expected frequencies of the Poisson distribution and compound Poisson distribution with logarithmic compounding distribution were also calculated and plotted in the same histogram in Excel. The implemented Poisson distribution (POISSON.DIST) has been used to find the expected Poisson frequencies. The expected frequencies according to the logarithmic compound Poisson distribution have been calculated using the negative binomial distribution of Equation 3.7. The first term of the equation has been replaced by the binomial coefficient $\binom{k+r-1}{k}$, where k is integer and r can be non-integer. Since $k+r-1$ can be non-integer, the binomial coefficient has been calculated as $\frac{\Gamma(k+r)}{\Gamma(k+1)\Gamma(r)}$. Here, $\Gamma(x)$ is the gamma function, which is implemented from Excel 2013 onwards.

The expected frequencies were graphically compared with the observed frequencies to see if the suggested distribution provides a good fit. Furthermore, if both probability distributions are possible, the expected frequencies of the probability distribution that is not suggested are also graphically compared. The histograms are visible in Appendix B: Goodness of fit tests. The findings have been summarised in Table 4.2.

First, two SKU's with a low average demand (0.96 and 0.35 per week) and a ratio $\frac{\sigma^2}{\mu} \leq 1.1$ (0.90 and 0.78) have been tested. According to literature, the Poisson distribution should be used to model the demand of these SKU. We have plotted the observed frequencies and the expected frequencies according to the Poisson distribution in the histograms of Figure B.1 (first SKU) and Figure B.2 (second SKU). Since the ratio $\frac{\sigma^2}{\mu}$ is for both SKU's lower than 1, it is not possible to plot the expected frequencies of the logarithmic compound Poisson distribution. When we compare the observed frequencies with the expected frequencies, we see that the Poisson distribution fits quite good for both SKU's. The goodness of fit tests therefore confirm the decision rule.

Subsequently, we tested two SKU's with a low average demand (3.58 and 0.10 per week) and a ratio $\frac{\sigma^2}{\mu} > 1.1$ (4.09 and 1.33). A logarithmic compound Poisson distribution is advised for both SKU's in literature. However, these demands can both also be modelled with a Poisson distribution. Therefore, the observed frequencies and the expected frequencies according to both distributions are plotted in the histograms of Figure B.3 (third SKU) and Figure B.4 (fourth SKU). We can conclude from these histograms that for both SKU's a logarithmic compound Poisson distribution indeed fits best. In the case of the third SKU, the expected frequencies according to the logarithmic compound Poisson distribution are generally far closer to the observed frequencies than the expected frequencies according to the Poisson distribution. For the fourth SKU, the differences are much smaller, but still the logarithmic compound Poisson distribution fits better than the Poisson distribution.

Finally, we tested the fit of a discrete and a continuous distribution for one SKU with a high average demand (11.90 per week). Since this demand is higher than 10 units and the CV of this SKU is higher than 0.5 (0.63), a gamma distribution is proposed in literature. The corresponding gamma distribution was plotted in the histogram of observed frequencies of this SKU. This histogram is visible in Figure B.5 (fifth SKU). The expected frequencies according to the logarithmic compound Poisson distribution were also plotted in this histogram, since this is the proposed discrete distribution for this SKU ($\frac{\sigma^2}{\mu} = 4.70 > 1.1$). We observe in the histogram what we expect based on literature: the continuous gamma distribution is a good approximation of the discrete logarithmic

compound Poisson distribution. Namely, the line of the expected frequency of the gamma distribution follows the tops of the expected frequency bars of the logarithmic compound Poisson distribution closely. This means that both distributions fit the historic demand data of this SKU equally well and it does not matter much which distribution we use from a goodness of fit point of view.

Table 4.2 The tested SKU's with their demand properties, the tested distributions, the distribution with the best fit (only when there were multiple distributions tested) and the suggested distribution from literature.

| SKU | Average week demand | $\frac{\sigma^2}{\mu}$ | CV | Tested demand distribution(s) | Demand distribution with the best fit | Suggested demand distribution in literature |
|-----|---------------------|------------------------|------|--|---|---|
| 1 | 0.96 | 0.90 | 0.97 | Poisson | N/A | Poisson |
| 2 | 0.35 | 0.78 | 1.50 | Poisson | N/A | Poisson |
| 3 | 3.58 | 4.09 | 1.07 | Poisson & logarithmic compound Poisson | logarithmic compound Poisson | logarithmic compound Poisson |
| 4 | 0.10 | 1.33 | 3.72 | Poisson & logarithmic compound Poisson | logarithmic compound Poisson | logarithmic compound Poisson |
| 5 | 11.90 | 4.70 | 0.63 | gamma & logarithmic compound Poisson | Both gamma and logarithmic compound Poisson | gamma |

Conclusion

We can draw some conclusions from the findings of Table 4.2. First, we can conclude that the decision rule of Axsäter (2006), which states that the Poisson distribution should best be used when the ratio $\frac{\sigma^2}{\mu} \leq 1.1$ and the logarithmic compound Poisson distribution otherwise, indeed applies to this project. Namely, when the logarithmic compound Poisson distribution is suggested by this decision rule, this distribution indeed has the best fit according to the tests. On the other hand, when the Poisson distribution is suggested, it is often the only distribution possible. Nevertheless, the Poisson distribution fits reasonably good in the tested cases. Furthermore, we can conclude that if a continuous distribution is proposed by literature, we can also use a discrete distribution instead, since it gives equally good results.

Based on these conclusions, we can model the demand of all the critical parts with (at least) two discrete probability distributions: the Poisson distribution and the compound Poisson distribution with logarithmic compounding distribution. The demand model should therefore consist of both these distributions. To be more precise, the demand model will look as in Table 4.3.

Table 4.3 The demand model

| Demand characteristic | Demand distribution to be used | Number of SKU's it applies to |
|---------------------------------|--------------------------------|-------------------------------|
| $\frac{\sigma^2}{\mu} \leq 1.1$ | Poisson | 9 |
| $\frac{\sigma^2}{\mu} > 1.1$ | logarithmic compound Poisson | 81 |

4.1.2. Choice of the inventory model

Now that we know how the most ideal demand model should look like, we can use this information in the choice of the inventory model. This choice will be made on the basis of the classification as visible in Figure 3.4 from the conclusion of the literature review (Section 3.5).

The first choice that needs to be made is whether a single item or a coordinated multi-item model should be chosen. We observed that the replenishments can be considered as orders from the same supplier. Besides, the items share the same mode of transportation, namely they are all transported in the same trailer(s). This means that a coordinated multi-item inventory model can be chosen. We also found that this approach has advantages for the company, mainly savings on unit transportation costs and ease of scheduling. It is expected that these advantages outweigh the disadvantages of coordinated control of multiple items. Therefore, it is best to choose this kind of inventory model.

The next distinction is the one between deterministic and stochastic coordinated multi-item inventory models. We have seen in Subsection 4.1.1 that we will only consider stochastic demand models. We will therefore focus on the coordinated multi-item inventory models under stochastic demand only.

Within these stochastic inventory models the choice should be made for continuous or periodic review. We discussed one continuous and three periodic review models. We learned that it is difficult to compute the parameters of the continuous model. Besides, it is more practical to the company to use an inventory model that needs periodic review as opposed to continuous review. Furthermore, we have seen that the periodic review models generally perform equally well or even better than the continuous model. For these three reasons, we will choose a periodic review joint replenishment model.

We studied three of these models, each of these models uses a different approach. The model of Viswanathan (1997) determines individual (R,s,S) models and a common R . He found that his model outperforms a lot of coordinated multi-item models in a lot of instances. However, the computational effort that this model requires, is regarded as too much. This is mainly due to the optimal (s,S) policies that are required. The remaining two models make use of individual (R,S) models and multiple R 's. Of these two models, the model of Fung et al. (2001) is preferred to the model of Atkins & Iyogun (1988) because of two reasons. The first one is the use of a service measure in the model of Fung et al. (2001). As mentioned before, it is often difficult to estimate shortage costs (as in the model of Atkins & Iyogun (1988)). On the other hand, it is often simpler to determine a suitable service level. This makes the model of Fung et al. (2001) more desirable. The second reason to use this model, is the use of compound Poisson demand in this model. As we have seen in Subsection 4.1.1, the inventory model should at least be able to handle both the Poisson and the compound Poisson distribution with a logarithmic compounding distribution. The compound Poisson distribution is especially important, since 81 of the 90 critical parts should be modelled with that distribution. The model of Fung et al. (2001) is focussed in particular on compound Poisson demand. The compounding distribution is a separate input in the model. It is therefore easy to choose a logarithmic compounding distribution to model logarithmic compound Poisson demand. This is in contrast to the model of Atkins & Iyogun (1988), which only assumes (unit-sized) Poisson demand and thus uses no compounding distribution for the demand size. The model of Fung et al. (2001) is therefore the best choice.

Changes to the inventory model

We will not implement the inventory model as it is in our inventory control system. We have seen that the model of Fung et al. (2001) assumes only compound Poisson demand. However, apart from

the compound Poisson distribution, the “normal” (unit-sized) Poisson distribution should also be used to model the demand of some SKU’s. This distribution should therefore also be implemented in the model of Fung et al. (2001). Given the separate input of a compounding distribution in this model, it is easy to incorporate the Poisson distribution in this model. This can be done by defining the compounding distribution as follows: set the probability for a demand transaction size of one equal to one and all probabilities of other demand sizes equal to zero. A more detailed description will be presented in Subsection 4.2.1. This notion has not explicitly been made in the article, but it is useful to consider it as it makes it possible to use this model for items that should be modelled by Poisson demand and items that should be modelled by compound Poisson demand.

Besides, it turned out that one algorithm proposed by Fung et al. (2001) (the HCTSP algorithm), is not giving the results it should give. Therefore, this algorithm had been adjusted and the new adjusted HCTSP algorithm will be used in the implementation. There is a further elaboration on this matter in Subsection 4.2.2.

Conclusion

We have reasoned in this subsection by means of Figure 3.4 from the conclusion of the literature review (Section 3.5) which inventory model is applicable for the inventory control system for the company. We can conclude that a periodic review stochastic coordinated multi-item inventory model is the most appropriate inventory model for this project. Of the discussed models in this category, we found that the model of Fung et al. (2001) fits the demand model as defined in Subsection 4.1.1 the best. Besides, it uses a service measure which can easily be determined. Therefore, this model will be chosen as the inventory model. However, some adjustments will be made to the inventory model. The model has been shortly presented in Chapter 3, but it will be described in more detail in Section 4.2, together with the adjustments.

4.2. Methodology of the inventory control system

The inventory control system will be based on the demand model as chosen in Subsection 4.1.1 and the inventory model as chosen in Subsection 4.1.2. The demand model and the inventory model are described in Subsection 4.2.1 and Subsection 4.2.2 respectively. The result of that part of the inventory control system will be that the (R,S) parameters of every critical part will be known. However, it is still not known *when* (date) and *how much* (quantity) to order, the replenishment decisions, for every critical part. The reason behind this is that additional information is required to calculate the replenishment decisions. Therefore, an extension is made to the inventory control system that calculates the replenishment decisions. Besides, it will advise the user which lots of the critical parts to pick exactly to satisfy these replenishment decisions. Furthermore, the inventory control system will include two other extensions, so-called checks. Together with the extension of the replenishment decisions, these checks are not described in the inventory model of Subsection 4.2.2, but are tailor-made extensions to the inventory control system based on the results that the inventory model provides. The extensions are described in different subsections:

- Extension 1: Replenishment decisions (the required order quantities and an advice on the lots to pick to satisfy these quantities) in Subsection 4.2.3
- Extension 2: Buffer inventory capacity check in Subsection 4.2.4
- Extension 3: Outstanding orders check in Subsection 4.2.5

The complete inventory control system will therefore consist of the demand model, the inventory model and the three extensions. We will describe the methodology behind these parts of the inventory control system in the next subsections. We first provide a list in which we define the

parameters, random variables, decision variables and output together with their notations. They are taken from Fung et al. (2001), since the inventory model is based on their model.

Parameters

- m the number of items in the inventory control system
- R_0^{min} the minimum base period to consider (in weeks²)
- R_0^{max} the maximum base period to consider (in weeks)
- ΔR_0 the time steps between the minimum and maximum base period to consider (in weeks)
- L_i the lead time of item i (in weeks)
- μ_i the mean week demand of item i
- σ_i^2 the variance of the week demand of item i
- λ_i the demand arrival rate of item i (in arrivals per week)
- $t_{i,max}$ the largest possible demand transaction size for item i
- P_i the service level of item i (according to service measure P_1 : the probability that the replenishment cycle ends with no backorders)
- K the major order cost (in euro)
- k_i the minor order cost for including item i in the order (in euro)
- h_i the holding cost rate of item i (in euro per unit per week)

Random variables

- t_i the demand transaction size of item i
- x_i the total demand of item i for the time period $L_i + R_i$

Decision variables

- R_0 the base period (in weeks)
- n_i the number of base periods that the review period of item i consists of
- S_i the order-up-to-level of item i

Output

- R_0 the base period (in weeks)
- $R_i = n_i R_0$ the review period of item i (in weeks)
- S_i the order-up-to-level of item i
- $RealP_i$ the real service level of item i
- EC_i the expected cost per week of item i (in euro)
- $Mean P_i$ the mean service level of the inventory control system
- EC the expected cost per week of the inventory control system (in euro)

4.2.1. Description of the demand model

The demand is modelled by the demand arrival rate λ_i and the discrete random variable t_i that describes the demand transaction size. This discrete random variable is described by probability distribution $p_{t_i}(t_0)$ (the compounding distribution). Here, $p_{t_i}(t_0)$ is the probability that t_i takes on the value t_0 for item i . We know from Subsection 4.1.1, that the λ_i and t_i should be able to model either the Poisson or logarithmic compound Poisson demand for a part, depending on the ratio $\frac{\sigma^2}{\mu}$. We also noted that the chosen inventory model can handle this by changing the compounding distribution. We can incorporate it in the inventory control system as follows:

If $\frac{\sigma_i^2}{\mu_i} \leq 1.1$ for item i , then:

² A week consists of 5 (working) days in this research, since the company only works on weekdays.

$$\lambda_i = \mu_i \quad (4.1)$$

$$p_{t_i}(t_0) = \begin{cases} 1 & \text{for } t_0 = 1 \\ 0 & \text{for } t_0 = 2, 3, \dots, t_{i \max} \end{cases} \quad (4.2)$$

Items with this demand characteristic should be modelled by a Poisson distribution. This distribution has as only parameter $\lambda_i t$, which is equal to both the mean (μ_i) and variance (σ_i^2) of the demand during time t . Since we defined everything in weeks, we will set t equal to one week. Therefore, λ_i is expressed in arrivals per week and is simply equal to the mean demand per week (μ_i). This is because all demands are unit sized. The simple relation between λ_i and μ_i is given by Equation 4.1. The unit sized demands should also become clear from the compounding distribution: the probability of a demand size of one (unit size) is equal to one and the probabilities of other demand sizes are zero. This behaviour is described by the probability distribution of Equation 4.2. Of course, this probability distribution is not defined for a demand size of zero ($t_0 = 0$) and demand sizes greater than the defined maximum ($t_0 > t_{i \max}$).

If $\frac{\sigma_i^2}{\mu_i} > 1.1$ for item i , then:

$$\lambda_i = -\mu_i \frac{\frac{\mu_i}{\sigma_i^2} \ln\left(\frac{\mu_i}{\sigma_i^2}\right)}{1 - \frac{\mu_i}{\sigma_i^2}} \quad (4.3)$$

$$p_{t_i}(t_0) = -\frac{\left(1 - \frac{\mu_i}{\sigma_i^2}\right)^{t_0}}{\ln\left(\frac{\mu_i}{\sigma_i^2}\right) t_0} \quad \text{for } t_0 = 1, 2, \dots, t_{i \max} \quad (4.4)$$

Items with this demand characteristic should be modelled by a compound Poisson distribution with logarithmic compounding distribution. Therefore, the demand arrival rate λ_i is given by Equation 4.3, which is derived from Equations 3.5 and 3.6 (Axsäter (2006)), where t is again equal to one week. The compounding distribution is the logarithmic distribution as described by Equation 4.4, which is derived from Equations 3.4 and 3.5 (Axsäter (2006)). Again, this probability distribution is not defined for a demand size of zero ($t_0 = 0$) and demand sizes greater than the defined maximum ($t_0 > t_{i \max}$).

In conclusion, the Poisson demand can be modelled with Equations 4.1 and 4.2 and the logarithmic compound Poisson demand can be modelled with Equations 4.3 and 4.4. This means that the complete demand model is described by these 4 equations.

4.2.2. Description of the inventory model

As mentioned before, the demand model is an input to the inventory model. Based on the demand model of Equations 4.1 through 4.4, we can define the probability distribution of total demand x_i of item i for the time period $L_i + R_i$. Let $p_{x_i}(x_0)$ be the probability that x_i takes on the value of x_0 for item i . The probabilities $p_{x_i}(x_0)$ can be calculated recursively by Equation 4.5³ (Fung et al. (2001)).

³ The second equation has been slightly adjusted but is still equivalent to the corresponding equation from Fung et al. (2001).

$$p_{x_i}(x_0) = \begin{cases} e^{-\lambda_i(L_i+R_i)} & \text{for } x_0 = 0 \\ \frac{1}{x_0} \sum_{j=\max(0, x_0-t_{i \max})}^{x_0-1} (x_0-j)\lambda_i(L_i+R_i)p_{t_i}(x_0-j)p_{x_i}(j) & \text{for } x_0 = 1, 2, \dots \end{cases} \quad (4.5)$$

Now that the probability distribution of x_i can be completely determined, we can formulate the service level constraint for item i as in Inequality 4.6 (Fung et al. (2001)).

$$\sum_{x_0=0}^{S_i} p_{x_i}(x_0) \geq P_i \quad (4.6)$$

Observe that the left-hand side of the inequality denotes the probability that the demand x_i of item i during time period $L_i + R_i$ is less than or equal to order-up-to-level S_i . Recall that under a (R, S) inventory model, the S should cover the demand during time period $L + R$. If this demand is lower than or equal to S , then all demand is fulfilled (no backorders). The probability of this event is given by the left-hand side of Inequality 4.6. This probability is equal to the service level P_i , the probability that the replenishment cycle ends with no backorders. The service level P_i is set by the company and should be achieved. Therefore, Inequality 4.6 imposes the inventory model to choose decision variable S_i so large that the determined service level is (at least) met.

However, the objective of the inventory model is not only to meet the defined service level, it is the objective to meet the defined service level at minimal costs. The costs to be minimised are the total expected costs (EC) of the inventory control system per unit time (week), given by Equation 4.7 (Fung et al. (2001)).

$$EC = \frac{K}{R_0} + \sum_{i=1}^m EC_i \quad (4.7)$$

The first term of Equation 4.7 represents the major order cost per week. Suppose that the base period R_0 is equal to 2.5 days ($R_0 = 0.5$ week). Then two times the major order cost per order K are incurred according to Equation 4.7, corresponding with 2 orders per week. The second term represents the minor order cost and holding cost per week. It is a summation of all the expected costs EC_i per item i . To be more specific, EC_i is the long-run average cost of item i per unit time (week) under an independent (R_i, S_i) inventory model. EC_i can be expressed as in Equation 4.8 (Fung et al. (2001)).

$$EC_i = \frac{k_i}{R_i} + h_i \left[\frac{\lambda_i E(t_i) R_i}{2} + \sum_{x_0=0}^{S_i} (S_i - x_0) p_{x_i}(x_0) \right] \quad (4.8)$$

Where $E(t_i) = \sum_{t_0=1}^{t_{i \max}} t_0 p_{t_i}(t_0)$ is the expected demand size and $R_i = n_i R_0$ (where n_i is a positive integer).

The first term of Equation 4.8 represents the minor order cost per week and works in the same way as the major order cost in Equation 4.7, where k_i is the minor order cost which is incurred if item i is included in the replenishment order. The difference is that in Equation 4.7 the order cost is for *the complete order* and R_0 is the period between *orders* and in Equation 4.8 the order cost is for *item i only* and R_i is the period between *orders with item i*. The second term represents the holding cost for item i per week. h_i is here the holding cost rate of item i per unit per week and the terms between brackets correspond to the expected average inventory level. To be more precise, the first term

between brackets corresponds to the expected average cycle inventory level and the second term to the expected safety stock.

We have now discussed all elements of the inventory model. The goal is to find values for the decision variables R_0 , n_i and S_i for $i=1,2,\dots,m$ that minimize Equation 4.7 subject to Inequality 4.6. This results in mathematical programming Problem 4.9 (Fung et al. (2001)).

$$\min EC = \frac{K}{R_0} + \sum_{i=1}^m \left[\frac{k_i}{R_i} + h_i \left[\frac{\lambda_i E(t_i) R_i}{2} + \sum_{x_0=0}^{S_i} (S_i - x_0) p_{x_i}(x_0) \right] \right] \quad (4.9)$$

$$s. t. \sum_{x_0=0}^{S_i} p_{x_i}(x_0) \geq P_i \quad \text{for } i = 1, 2, \dots, m$$

Problem 4.9 is a non-linear mixed-integer programming problem, which makes it extremely difficult to solve the problem exactly. The non-linearity results from the fact that with $R_i = n_i R_0$, the decision variables n_i and R_0 are multiplied. However, we can solve Problem 4.9 heuristically when we separate R_0 from the n_i 's and S_i 's. If R_0 is fixed, solving Problem 4.9 is equivalent to solving Equation 4.8 subject to Inequality 4.6 for every item i . This means solving m times the mathematical programming sub Problem 4.10 (Fung et al. (2001)).

$$\min EC_i = \frac{k_i}{R_i} + h_i \left[\frac{\lambda_i E(t_i) R_i}{2} + \sum_{x_0=0}^{S_i} (S_i - x_0) p_{x_i}(x_0) \right] \quad (4.10)$$

$$s. t. \sum_{x_0=0}^{S_i} p_{x_i}(x_0) \geq P_i$$

This sub problem can be solved by enumerating n_i to find the local minimizers n_i and S_i for every item i . To this end, Fung et al. (2001) propose a heuristic algorithm to solve Problem 4.10 with given R_0 . They call this heuristic algorithm: HSTSP, please refer to the paper of Fung et al. (2001) for the pseudocode of the HSTSP algorithm. A graphical representation of the HSTSP algorithm is given in the flowchart of Figure 4.1.

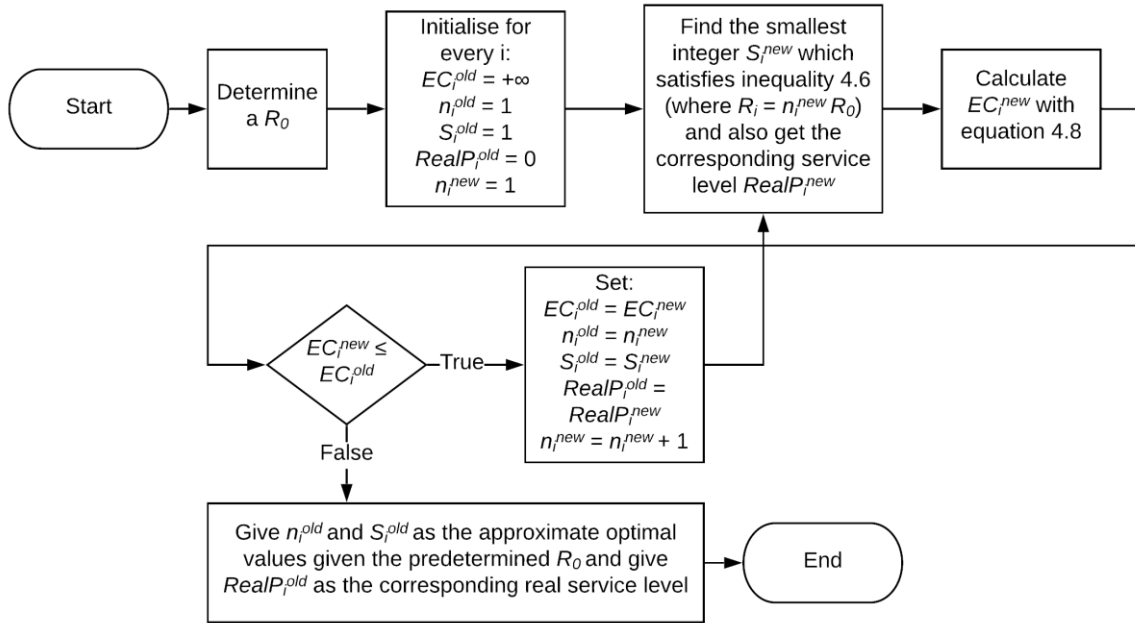


Figure 4.1 The HSTSP algorithm of Fung et al. (2001) represented in a flowchart.

The result of executing the HSTSP algorithm for every item i is that the parameters (R_i, S_i) of the inventory model are found for every item as follows: $R_i = n_i^{old} R_0$ and $S_i = S_i^{old}$. The resulting service level per item is $RealP_i^{old}$. Nevertheless, these will not be the final parameters of the inventory model. This is due to the fact that we minimised the expected cost EC_i of each item i independently given a user-defined, fixed R_0 . However, it does not mean that the expected cost EC of the complete inventory control system is minimised as a different R_0 might result in a lower EC . Therefore, Fung et al. (2001) propose a second heuristic algorithm, called the HCTSP, which tries to find the optimal R_0 by increasing R_0 step by step. It thereby uses the n_i , S_i , $RealP_i$ and EC_i from the HSTSP output for every item i . The HCTSP thus solves Problem 4.9 heuristically by enumerating R_0 to find the minimum EC . Please refer to the paper of Fung et al. (2001) for the pseudocode of the HCTSP algorithm. A graphical representation of the HCTSP algorithm is given in the flowchart of Figure 4.2.

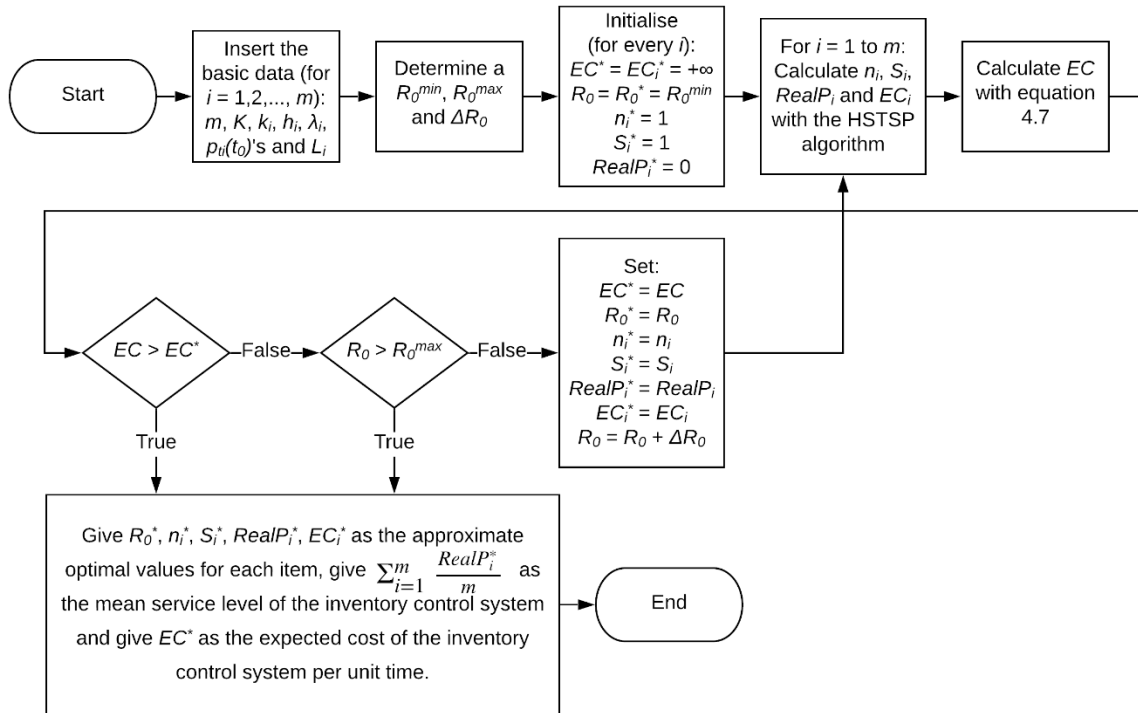


Figure 4.2 The HCTSP algorithm of Fung et al. (2001) represented in a flowchart.

We see that the HCTSP algorithm requires a range of values for R_0 for the enumeration. This range should be given as a minimum base period R_0^{min} , a maximum base period R_0^{max} and a step length ΔR_0 . The algorithm starts evaluating the EC at R_0^{min} and then evaluates EC with R_0 one step larger ($R_0^{min} + \Delta R_0$). This process continues until EC does not decrease anymore or when R_0^{max} is reached. The algorithm then returns the final parameters of the (R_i, S_i) inventory model for item i as follows: $R_i = n_i^* R_0^*$ and $S_i = S_i^*$. The service level and expected cost for item i corresponding to these parameters are equal to $RealP_i^*$ and EC_i^* respectively. The service level and expected cost of the complete inventory control system are equal to $\sum_{i=1}^m \frac{RealP_i^*}{m}$ and EC^* respectively. The former will be called *Mean P_i* in this project.

Discussion

In conclusion, the inventory model starts with executing the HCTSP algorithm, which enumerates solutions for a certain number of R_0 's. Within each R_0 , the HSTSP algorithm needs to be executed m times. Within each m , the HSTSP algorithm enumerates solutions for a certain number of n_i 's. Note however that, in this way, the inventory model only gives a heuristic solution to Problem 4.9 and therefore EC^* is most likely not the absolute minimum. This is due to the fact that the R_0 , n_i and S_i are not simultaneously determined, but they are determined in two steps (two algorithms). It's also in the nature of both (heuristic) algorithms that the absolute minimum will not be attained. Namely, both algorithms rely on enumeration to find the best solution. In this case, there are two flaws with this approach:

1. Not every solution is enumerated, this only applies to the enumeration of R_0 (HCTSP algorithm). Of course, every positive non-integer number can be a feasible base period R_0 . However, it is not possible to enumerate every positive non-integer number and evaluate the corresponding EC . Therefore, a range for R_0 should be given (as $R_0^{min}, R_0^{min} + \Delta R_0, \dots, R_0^{max}$) and only the EC 's for values of R_0 in this range are evaluated. Most likely, the value of R_0 for

which EC attains its absolute minimum will not be evaluated in this way. Nevertheless, the question is whether this is a problem, since many solutions will not be very convenient. For example, when the optimal base period appears to be 3.15 days, it is likely not practicable to do a replenishment every 3.15 days. The company will probably replenish every 3 days in such a case. Specifying a range with practical R_0 's is therefore probably more important than evaluating all possible (and many unpractical) R_0 's. Besides, it is much more efficient to evaluate only practical R_0 's as it saves a lot of computational time. This flaw does not apply to the enumeration of the n_i 's (HSTSP algorithm). This is the case since n_i can only be a positive integer and the enumeration starts with $n_i = 1$ and is incremented with 1. Therefore, no feasible n_i is omitted with the enumeration.

2. The enumeration might stop too early, this applies to both the enumeration of R_0 (HCTSP algorithm) and n_i (HSTSP algorithm). Both algorithms evaluate the EC function from left to right and stop when the expected cost is greater for the first time. When the EC function is convex, it means that the algorithm has found the global minimum (the blue point in the left graph of Figure 4.3), since all the next expected costs will be higher. However, when the EC function is not convex, it could mean that only a local minimum has been found when the algorithm stops (the left blue point in the right graph of Figure 4.3). However, if the algorithm would continue long enough after the local minimum, it could possibly have found a lower local minimum or even the global minimum (the right blue point in the right graph of Figure 4.3). It is therefore not easy to determine when the algorithm should stop in the case of a non-convex function.




Figure 4.3 The graph on the left shows a convex function with one (global) minimum. The graph on the right shows a non-convex function with two minima. The left minimum is merely a local minimum whereas the right minimum is the global minimum of the depicted range. The figure is taken from (He, Rexford, & Chiang, 2010).

After initial testing with the HCTSP and HSTSP algorithm on the data from this project, it appeared that the EC and EC_i functions were both non-convex. This means that the behaviour of a non-convex EC function, as described above and depicted in the right graph of Figure 4.3, is also present in our case.

It is very difficult to let the HSTSP algorithm search for the global minimum instead of a local minimum, because the user should define a value for n_i as a sign that the HSTSP stops searching (otherwise it will continue forever). However, as the HSTSP has to enumerate n_i 's for each of the m items included in the inventory control system, the user should define m of such values. Furthermore, the evaluated range of values for n_i is also dependent on the value of R_0 , since the review period of an item i is $R_i = n_i R_0$. It is therefore difficult to determine m values of n_i for which the HSTSP algorithm should stop evaluating. Hence, we will not alter the (search procedure of the) HSTSP algorithm and accept this flaw.

On the other hand, the heuristic search procedure of the HCTSP algorithm can be improved more easily. This is the case since the user already needs to define a value for which the HCTSP algorithm should stop searching for the global minimum, namely R_0^{max} . This is visible in the flowchart of the HCTSP algorithm (Figure 4.2) as the second decision diamond. The first decision diamond can thus easily be removed from the HCTSP algorithm and the HCTSP algorithm will still stop at some point. It is the first decision diamond that causes the problem of stopping the search for the global minimum too early in the case of a non-convex EC function. Therefore, the HCTSP can be improved by removing this first decision diamond and the second decision diamond makes it possible to do this without many negative consequences. The result will be that every possible solution within the given R_0 range will be enumerated. Subsequently, the global minimum within this enumerated list can be chosen. Of course, the solution with the least EC will be chosen. Therefore, this choice procedure should also be added to the HCTSP algorithm.

The main difference between the original HCTSP algorithm and the HCTSP algorithm with these adjustments is then as follows: the original HCTSP stops evaluating solutions when a (local or global) minimum has been found *even if not all solutions in the given R_0 range are evaluated yet*, whereas the adjusted HCTSP algorithm *always evaluates all solutions in the given R_0 range*. Only then, the adjusted HCTSP algorithm will certainly choose the lowest minimum of the given R_0 range, which means that it will always find the global minimum of the given R_0 range. The advantage of this adjusted algorithm is that it will find the global minimum of the given R_0 range for *both* convex and non-convex EC functions. The only disadvantage of the adjusted HCTSP algorithm is that the computational effort (and thus time) is almost always larger, as the adjusted algorithm always enumerates all possible solutions within the given R_0 range. However, when the user gives a reasonable R_0 range (i.e. a R_0^{max} reasonably close to R_0^{min} and a ΔR_0 reasonably large), the number of solutions to enumerate (and with that the computational time) stays within reason. We therefore choose to develop the adjusted HCTSP algorithm and use that algorithm instead of the original HCTSP algorithm in the implementation of the inventory control system. The development and the details of the adjusted HCTSP algorithm will be given now.

Adjusted HCTSP algorithm

To find the global minimum of a specified range instead of a local minimum, the adjusted HCTSP algorithm needs to be conceptually different from the original HCTSP algorithm. The original HCTSP algorithm stores two values for EC , R_0 , n_i , S_i , $RealP_i$ and EC_i : one 'current' value and one optimal value (denoted with an asterisk). The optimal values are initialised first. Subsequently, the current values are calculated. If the current values are better than the optimal values (i.e. the corresponding EC is lower), then the optimal values are overwritten by the current values. This means that the old optimal values will not be stored anymore and the current values have become the new optimal values. Then, R_0 is incremented and the main loop starts again. This process continues until the first time that the current values are not better than the optimal values anymore or when the incremented R_0 exceeds R_0^{max} . The optimal values at that moment are given as output of the HCTSP algorithm.

However, the adjusted HCTSP algorithm should be fundamentally different: it should not stop at the moment that the EC rises for the first time. Instead, the algorithm should evaluate all the values for EC in the given R_0 range and should determine the optimal values only then. This has major implications on the original HCTSP algorithm. The first one is that the "If $EC > EC^*$ Then ..." statement of step 2.3 is removed. This means that the main loop only stops when the R_0^{max} is exceeded, which in turn means that all possible R_0 's are being evaluated. The second major implication is that a new step is added between the main loop (step 2) and the output (step 3), so the new step becomes step

3 and the output becomes step 4. This new step consists of searching the minimum EC and finding the corresponding values for R_0 , n_i , S_i , $RealP_i$ and EC_i . This leads also to another major implication: the concept of storing two values for each of the aforementioned variables (a current and an optimal value) is not sufficient. The problem is that the values get overwritten, so they are not all known anymore at step 3. When the minimal EC is chosen at step 3, it should be possible to retrieve the values of R_0 , n_i , S_i , $RealP_i$ and EC_i belonging to that EC . Therefore, all values need to be stored in the process and they should be linked to each other. We therefore introduce two new variables: $index_R_0$ and $index_R_0^{max}$. Since new values are calculated for every possible R_0 , values can be retrieved and linked to each other by an index given to a specific R_0 . The first R_0 gets index 1 ($index_R_0 = 1$), the second R_0 gets index 2 ($index_R_0 = 2$) etc. The final R_0 in the R_0 range gets index $index_R_0^{max}$. Since one value for EC needs to be stored for every $index_R_0$, the EC values need to be stored in a one-dimensional array with length $index_R_0^{max}$. As n_i , S_i , $RealP_i$ and EC_i also have an index indicating the item, values for these variables need to be stored in two-dimensional arrays with length m (first dimension) by $index_R_0^{max}$ (second dimension). This makes it easy to find the right values in step 3 of the adjusted HCTSP algorithm: when the $index_R_0$ of the minimum EC is found, then the corresponding values can be found in their arrays under array index $index_R_0$. Here the notation with the asterisk comes in, because these values (found in the arrays) will be stored separately as the optimal values. Suppose that $EC(5) = EC^*$, then 5 will be the $index_R_0$ that is used to find the other optimal values. In that case, the value of $n_i(5)$ is looked up in the $n_i(index_R_0)$ array and is then given to the variable n_i^* . This is done for every i and for every of the following variables: n_i^* , S_i^* , $RealP_i^*$ and EC_i^* . Of course, the value for EC^* has already been found when determining the minimum EC . Furthermore, the value for R_0^* can easily be found with the following relation: $R_0^* = R_0^{min} + (index_R_0 - 1) * \Delta R_0$.

The conceptual difference between storing a current value and an optimal value (original HCTSP) or storing an array with values and an optimal value (adjusted HCTSP) does not have implications for step 0 and step 3 of the original HCTSP algorithm. Therefore, they are taken over literally from the original HCTSP except that step 3 becomes step 4 in the adjusted HCTSP. However, step 1 of the original HCTSP (initialisation) changes due to this difference. The first thing to do at this step is to calculate $index_R_0^{max}$ which determines also the number of R_0 's of which the EC should be evaluated.

$index_R_0^{max}$ can be calculated as $\left\lfloor \frac{(R_0^{max} - R_0^{min})}{\Delta R_0} + 1 \right\rfloor$, with the floor function $\lfloor x \rfloor$ returning the largest integer less than or equal to x . The last thing that should be done is the initialisation of the variables. The variables storing the optimal values are initialised in the original HCTSP and this initialisation remains the same as these variables are still used by the adjusted HCTSP. The same goes for the initialisation of R_0 , it should be initialised as the given value of R_0^{min} . In addition, the new variable $index_R_0$ should be initialised as 1, since R_0^{min} is the first R_0 that should be evaluated.

The pseudocode of the adjusted HCTSP algorithm, containing all the adjustments described above, is visible in Figure 4.4. This is the algorithm that will be implemented in the inventory control system as opposed to the original HCTSP algorithm given in the paper of Fung et al. (2001).

Adjusted HCTSP

Step 0 Input basic data.

Given $m, K, k_i, h_i, \lambda_i, p_{ti}(t_0), L_i$ for $i = 1, 2, \dots, m$.

Step 1 Initialize.

Given the minimum, the maximum base review time R_0^{min}, R_0^{max} and the step length ΔR_0 , calculate the number of R_0 's to evaluate with: $index_R_0^{max} = \left\lceil \frac{(R_0^{max} - R_0^{min})}{\Delta R_0} + 1 \right\rceil$

Furthermore, set (for $i = 1$ to m where applicable):

$$\begin{array}{llll} R_0 = R_0^{min} & R_0^* = R_0^{min} & n_i^* = 1 & RealP_i^* = 0 \\ index_R_0 = 1 & EC^* = +\infty & S_i^* = 1 & EC_i^* = +\infty \end{array}$$

Step 2 Main loop (Calculate the EC for every R_0 in the range R_0^{min} to R_0^{max} with steps of ΔR_0).

Step 2.1

For $i = 1$ to m

Get $n_i(index_R_0), S_i(index_R_0), RealP_i(index_R_0), EC_i(index_R_0)$
as $n_i^{old}, S_i^{old}, RealP_i^{old}, EC_i^{old}$ from HSTSP.

Step 2.2

Get $EC(index_R_0)$ as EC from Equation 4.7.

Step 2.3

If $index_R_0 > index_R_0^{max}$ then

Go to step 3

Else

$$R_0 = R_0 + \Delta R_0$$

$$index_R_0 = index_R_0 + 1$$

Go to step 2.1

Step 3 Search the minimum EC^* and find the corresponding values of $R_0^*, n_i^*, S_i^*, RealP_i^*, EC_i^*$.

$$EC^* = \min_{i \in \{1, 2, \dots, index_R_0^{max}\}} EC(i)$$

For $j = 1$ to $index_R_0^{max}$

If $EC(j) = EC^*$ then

$$R_0^* = R_0^{min} + (j - 1) * \Delta R_0$$

For $i = 1$ to m

$$n_i^* = n_i(j)$$

$$S_i^* = S_i(j)$$

$$RealP_i^* = RealP_i(j)$$

$$EC_i^* = EC_i(j)$$

Go to step 4

Step 4 Output.

Output $R_0^*, n_i^*, S_i^*, RealP_i^*, EC_i^*$ as the approximate values for each item.

Furthermore, $\sum_{i=1}^m \frac{RealP_i^*}{m}$ is the mean service level of the inventory control system and EC^* is the expected cost of the inventory control system per unit time.

Figure 4.4 The pseudocode of the adjusted HCTSP algorithm. The adjustments are made to the pseudocode of the original HCTSP algorithm of Fung et al. (2001).

4.2.3. Extension which calculates the replenishment decisions

The inventory model as described in Subsection 4.2.2 only returns the parameters (R_i, S_i) for every item i . With these parameters, we can determine the *replenishment schedule*, i.e. when each item needs to be reviewed (every R_i weeks) and to which level each item should be ordered up to at the review moments (level S_i). However, we cannot yet determine the *replenishment decisions*, i.e. the exact dates of the review moments and the quantities that should be ordered at these dates. We will therefore develop an extension that calculates these replenishment decisions.

The first thing that needs to be determined, are the replenishment dates. The replenishment dates are the dates at which the inventory is reviewed and replenishment orders are placed. The time period between consecutive replenishments is known as the base period R_0 . It is equal to the lowest R_i that appears among the SKU's. It means that SKU's with that R_i are reviewed at every replenishment and that the other SKU's are not reviewed at every replenishment. Suppose that the lowest R_i that appears among the SKU's is 1 week, then $R_0 = 1$ week and there will be 1 week between consecutive replenishments, see Figure 4.5. The only thing that needs to be determined, is the date of the first replenishment. The rest of the replenishment dates will follow automatically, since they will always be scheduled one base period (1 week) after the last replenishment. In Figure 4.5, an arbitrary date was chosen for the first replenishment to show how the rest of the replenishment dates are determined. However, the extension which calculates the replenishment decisions should not randomly pick a replenishment date for the first replenishment. It is chosen to set the first replenishment date equal to the date on which the replenishment schedule is made. Namely, the assumption is that when there is a need for a replenishment schedule, there is a need for a replenishment immediately. In conclusion, the replenishment dates are fixed when you have determined the date of the first replenishment and the length of the base period R_0 .




Figure 4.5 Timeline which shows an example of replenishment dates when $R_0 = 1$ week

We mentioned that not every SKU is reviewed at every replenishment date. However, from the (R_i, S_i) inventory model that is determined for every SKU i , we know that there are repetitive replenishment cycles for each item. Furthermore, the R_i 's are always multiples of each other, meaning that replenishments for multiple parts coincide. Therefore, a schedule with replenishments can be constructed solely based on the parameter R_i . Several schedules can be constructed, two of them are given in Table 4.4 and Table 4.5, again using a base period of 1 week as example.

Table 4.4 An example of a schedule with replenishments solely based on the parameter R_i ($R_0 = 1$ week), with replenishments at the beginning of the first base period of a cycle.

| SKU | R_i | Replenishment date (date at the beginning of the base period) | | | | | | | | | | | |
|-----|---------|---|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| | | 01-1-2020 | 08-1-2020 | 15-1-2020 | 22-1-2020 | 29-1-2020 | 05-2-2020 | 12-2-2020 | 19-2-2020 | 26-2-2020 | 04-3-2020 | 11-3-2020 | 18-3-2020 |
| 1 | 1 week | x | x | x | x | x | x | x | x | x | x | x | x |
| 2 | 2 weeks | x | | x | | x | | x | | x | | x | |
| 3 | 3 weeks | x | | | x | | | x | | | x | | |
| 4 | 4 weeks | x | | | | x | | | | x | | | |
| 5 | 5 weeks | x | | | | | x | | | | | x | |
| 6 | 6 weeks | x | | | | | | x | | | | | |

Table 4.5 An example of a schedule with replenishments solely based on the parameter R_i ($R_0 = 1$ week), with replenishments at the beginning of the last base period of a cycle.

| SKU | R_i | Replenishment date (date at the beginning of the base period) | | | | | | | | | | | |
|-----|---------|---|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| | | 01-1-2020 | 08-1-2020 | 15-1-2020 | 22-1-2020 | 29-1-2020 | 05-2-2020 | 12-2-2020 | 19-2-2020 | 26-2-2020 | 04-3-2020 | 11-3-2020 | 18-3-2020 |
| 1 | 1 week | x | x | x | x | x | x | x | x | x | x | x | x |
| 2 | 2 weeks | | x | | x | | x | | x | | x | | x |
| 3 | 3 weeks | | | x | | | x | | | x | | | x |
| 4 | 4 weeks | | | | x | | | | x | | | | x |
| 5 | 5 weeks | | | | | x | | | | | x | | |
| 6 | 6 weeks | | | | | | x | | | | | | x |

From the schedules in Table 4.4 and Table 4.5, it becomes clear that the replenishment cycles of all 6 SKU's have different lengths, as all SKU's have a different review period R_i . Thus, the length of the review period is equal to the length of the replenishment cycle. More specifically, every replenishment cycle consists of n_i base periods ($R_i = n_i * R_0$). At the beginning of one of these n_i base periods, there should be one review moment for the SKU. There is no choice when $n_i = 1$, but when n_i is a larger integer, then there are multiple possibilities. That is why there are several schedules possible. Two schedules are given in Table 4.4 and Table 4.5. Here, the black rectangles are the replenishment cycles. You can see that the width of the rectangle is related to the length of the review period R_i : the black rectangles span n_i columns (base periods). In one of the columns within the black rectangle there should be a review moment of the inventory of that SKU. However, it should be in the same column for every rectangle of that SKU, otherwise there will not constantly be a time period of R_i weeks between successive review moments. For example, if there is a review moment at the beginning of the first base period of the first replenishment cycle, there should also be a review moment at the beginning of the first base period of the second replenishment cycle. Only then there is a time period equal to R_i weeks between the reviews.

In Table 4.4 and Table 4.5, the replenishment dates at the top of the columns indicate the starting dates of the base periods (columns). If there is a review moment planned on a date, then it is represented by a black cross in the corresponding column. It becomes clear from the two schedules that given certain review periods R_i , different positions of the crosses are possible, each unique combination giving a different schedule. This means that even if the (R_i, S_i) parameters and the replenishment dates are fixed, there are still multiple replenishment schedules possible. For

example, if a SKU has a review period of 3 base periods (so the black rectangles consists of 3 columns), there are three choices: plan the review moment at the beginning of the first, second or third base period of the review period. In the schedule of Table 4.4, it was chosen to plan all the review moments at the beginning of the first base period of the review period. In the schedule of Table 4.5, it was chosen to plan all the review moments at the beginning of the last base period of the review period. Obviously, you can make much more different schedules by choosing the crosses in different positions, as long as there is one cross in every rectangle and the crosses in rectangles from the same row are in the same position. The long-run replenishment cost of each schedule is the same, so it does not matter which schedule we choose. We choose the schedule of Table 4.5, where every SKU is reviewed at the beginning of the last base period of the review period. Since all review periods start at the first replenishment date, this schedule has the benefit that not all SKU's are immediately reviewed at the first replenishment date.

Now that the exact dates of the review moments and the SKU's that are reviewed at each moment are determined, we only need the quantities that should be ordered at these moments to arrive at the *replenishment decisions* for these SKU's. From the parameter S_i we know the level to which SKU i should be ordered up to at its review moments. We also know that the order quantities are variable under this kind of inventory model. Namely, the order quantity is the difference⁴ between the order-up-to-level S_i and the current inventory level (which is not the same at each review moment). We therefore need to retrieve the current inventory level for each reviewed SKU from the WMS in order to find the replenishment decisions.

The definition of this current inventory level is important. We define it as follows: the buffer inventory is replenished via the inventory control system by the bulk inventory and the external warehouses. The current inventory level is therefore the buffer inventory level. In order to find this level, we need to retrieve the locations of all the lots of a certain part from the WMS and, based on the location, subdivide them in lots from the buffer inventory and inventory from the bulk/external warehouses.

According to the definition of the current inventory, we should add up all the lot quantities of the buffer inventory lots to find the current inventory level. This level should be compared with the order-up-to-level. If the order-up-to-level is higher than the current inventory level, then the difference should be ordered. We will call this difference the required order quantity. The replenishment order should consist of one or multiple lots from the bulk inventory and/or external warehouses. The user of the inventory control system can determine him- or herself which lot(s) should be picked and if it/they should be completely or partially moved. The only condition is that the lot(s) should consist of enough units to raise the inventory level to at least the order-up-to-level. Ordering more units than necessary is even better for the service level, but at the expense of increased holding costs. These are considerations that still need to be made by the user.

Recall that we saw in the analysis of the current situation (Chapter 2) that the current procedure is to order the lot(s) with the oldest lot date (FIFO). We also noted that this procedure cannot always strictly be followed from a practical point of view. We therefore give the user freedom in choosing the lot(s). However, we will give an advice to the user regarding the best lot(s) to choose based on the lot date. We therefore need to retrieve the so-called FIFO dates from all the bulk/external warehouse lots of a reviewed SKU from the WMS as well. These lots should be ranked based on the FIFO date (lot date). If the lot quantity of the oldest lot is equal to or more than the required order

⁴ If we define the difference as $S_i - \text{current inventory level}$ then the order quantity is equal to the difference only if the difference is positive. If the difference is zero or negative, then the order quantity is equal to zero.

quantity, then only this lot should be advised. If it is not sufficient, we should also look at the second oldest lot. If the combined quantity of both these lots satisfies the required order quantity, then these two lots are advised. Otherwise, the third oldest lot is evaluated. This process continues until the entire required order quantity is covered. The resulting lots are advised to the user of the inventory control system. However, it can also happen that there are no lots or not enough lots to satisfy the required order quantity. In that case, the inventory control system should also give the quantity that is short for the concerning SKU.

The steps described in the last two paragraphs should be executed at every review moment for every SKU that has a planned review at that moment. When a reviewed SKU has a required order quantity of zero (so the current inventory level is equal to or higher than the order-up-to-level), it is not included in the replenishment decisions. Thus, the list with reviewed SKU's is possibly longer than the list with the replenished SKU's (replenishment decisions). However, when the number of SKU's in the inventory control system is large, it is likely that at least some of the reviewed SKU's need replenishment. In that case, a replenishment order, according to the replenishment decisions, is placed by the user of the inventory control system. That is why we call the date on which this all happens "the replenishment date" in the beginning of this subsection.

In conclusion, when this extension of the inventory control system is used on a given replenishment date, the following replenishment decisions are returned as output. The replenishment decisions only apply to the reviewed SKU's that need replenishment (only their required order quantities are greater than zero):

1. Required order quantity, it is equal to the (positive) difference between the current inventory level and the order-up-to-level.
2. Advice on the lots to satisfy the required order quantity, based on the FIFO method. This advice consists of the lot date, lot location and lot quantity of all the advised lots.
3. Quantity short, if the advice and thus the inventory from the bulk/external warehouses cannot satisfy the required order quantity. The available lots are given by the advice of the second point and the quantity that is still short is given by this point. The quantity short is equal to the difference between the required order quantity and the available lot quantity.

The first point is the decision that the user should stick to, whereas the second point gives the best way to fulfil the requirement of the first point (best according to the FIFO method). Nevertheless, the user can neglect the second point if (s)he has a better way to satisfy the required order quantity. The third point is only given if there is not enough inventory in the bulk and external warehouse inventories to satisfy the required order quantity. In that case, a complete advice cannot be given and also the user is most likely not able to find enough inventory to satisfy the required order quantity in another way.

4.2.4. Buffer inventory capacity check extension

This subsection describes the next extension to the inventory control system: the buffer inventory capacity check. This check is tailor-made to the buffer inventory of the A/B series. The check is required since the inventory model does not take into account the capacity of the buffer inventory at the production location. It might be the case that an item has a high demand, higher order cost but low holding cost. In that case, the inventory model decides to order a large quantity of the part on an infrequent basis. This quantity may be so large that the buffer inventory cannot hold it, especially if this applies to many of the controlled SKU's. The capacity of the buffer inventory has been determined in Chapter 2. To check whether the buffer inventory capacity is large enough, we are implementing a capacity check on the replenishment schedule.

Figure 4.6 shows the development of the inventory position (dotted line) and inventory level (solid line) over time when inventory is controlled by an (R,S) inventory model. Only the solid line is relevant in this case, since that line depicts the on hand inventory (*OHI*) and thus the amount of inventory which is physically in the buffer inventory. The peaks in this graph are the maximum amount of inventory that will ever be in the buffer capacity, provided that the order-up-to-level S is never exceeded with replenishments. We will call this maximum the maximum *OHI*. When the observed capacity of the buffer inventory is higher than this maximum *OHI*, then there will not be a capacity problem. Therefore, the capacity check will compare this maximum *OHI* with the observed buffer inventory capacity for every SKU of the inventory control system.




Figure 4.6 The inventory position (dotted line) and inventory level (solid line) over time according to a (R,S) inventory model. The figure is taken from (Jensen & Bard, 2002).

The first task is to find the maximum *OHI* for every SKU of the inventory control system. We can see in Figure 4.6 that the maximum *OHI* is not equal to the order-up-to-level S , but equal to the inventory level after the replenishment lead time L . This means that the demand during the lead time D_L should be subtracted from S . D_L is a random variable and therefore does not take on one value.

Nevertheless, we can determine the expected D_L as $E(y_i) = \sum_{y_0=0}^{\infty} y_0 p_{y_i}(y_0)$ where y_i is the D_L of item i and $p_{y_i}(y_0)$ is the probability that y_i takes on the value of y_0 for item i . The $p_{y_i}(y_0)$'s can be determined in a similar fashion as the $p_{x_i}(x_0)$'s in Equation 4.5, see Equation 4.11.

$$p_{y_i}(y_0) = \begin{cases} e^{-\lambda_i L_i} & \text{for } y_0 = 0 \\ \frac{1}{y_0} \sum_{j=\max(0, y_0 - t_{i \max})}^{y_0 - 1} (y_0 - j) \lambda_i L_i p_{t_i}(y_0 - j) p_{y_i}(j) & \text{for } y_0 = 1, 2, \dots \end{cases} \quad (4.11)$$

With the help of Equation 4.11, $E(y_i)$ can be calculated for every SKU i . Note that a summation until infinity is required for the calculation of $E(y_i)$. We propose to use a self-determined maximum instead of infinity for the maximum of the summation. We will call this maximum $y_{i \max}$, the largest possible D_L for item i . The equation to calculate $E(y_i)$ is then given by Equation 4.12.

$$E(y_i) = \sum_{y_0=0}^{y_{i \max}} y_0 p_{y_i}(y_0) \quad (4.12)$$

We mentioned before that the maximum *OHI* can be calculated by subtracting D_L from S . Since we use the expected value of the random variable D_L , the maximum *OHI* will also be an expected value. The maximum *OHI* should be an integer, so we round up the expected maximum *OHI* to the nearest integer when the expected value is a decimal number. Rounding up will always be safe since it is the worst case in terms of capacity usage. The expected maximum *OHI* for item i can therefore be expressed with Equation 4.13, where the ceiling function $\lceil x \rceil$ returns the least integer greater than or equal to x .

$$\text{maximum } OHI_i = [S_i - E(y_i)] \quad (4.13)$$

Now that the maximum *OHI* can be determined for every SKU, the next step is to determine if the buffer inventory capacity is sufficiently large. An analysis of the buffer inventory capacity is provided in Chapter 2. We came to the conclusion that the buffer inventory capacity can be best divided in three parts:

1. Individual capacities, for (almost) every SKU individually
2. Shared capacities, for certain sets of SKU's
3. Overflow capacity, one large capacity for one large set of SKU's

The order in which we check if the buffer inventory capacity for a certain SKU is large enough, is from 1 to 3. First, we will check if the individual capacity of that SKU (in units) is large enough to store the entire maximum *OHI* of that SKU. If so, there is no capacity problem for that SKU. If not, the maximum *OHI* will not entirely fit and we will check whether the SKU is contained in a shared capacity. If so, we will count the number of pallets that is needed to store the remainder of the maximum *OHI* that could not be stored in the individual capacity. This number is added to the usage of the concerning shared capacity (in pallets). If the SKU is not contained in a shared capacity, we will check if the SKU is contained in the overflow capacity. If so, we do the same the procedure as with the shared capacity, but this time the number of required pallets is added to the usage of the overflow capacity (in pallets). If not, we will simply give the number of pallets short in the individual capacity of the SKU.

Subsequently, we will compare the usages of the shared capacities and the overflow capacity with predetermined, observed sizes of these capacities (in pallets). One of the following two things can occur: Firstly, the size of a shared/overflow capacity can be not sufficient enough. In that case, there is a certain number of pallet places short in that shared/overflow capacity. Secondly, the size of a shared/overflow capacity can be sufficient enough. In that case, there can be still some empty pallet places remaining in that shared/overflow capacity or it is just the right size for the maximum *OHI*.

It could be the case that only the individual capacity of a SKU is already large enough to store the maximum *OHI* of that SKU. Not only is there no need to use a shared/overflow capacity then, but also it could be that there are still empty pallet places remaining in the individual capacity of that SKU. In that case, the extension should also report the number of pallet places remaining in the individual capacity of that SKU.

The capacity checking procedure, as described above, is graphically represented in the flowchart of Figure 4.7.

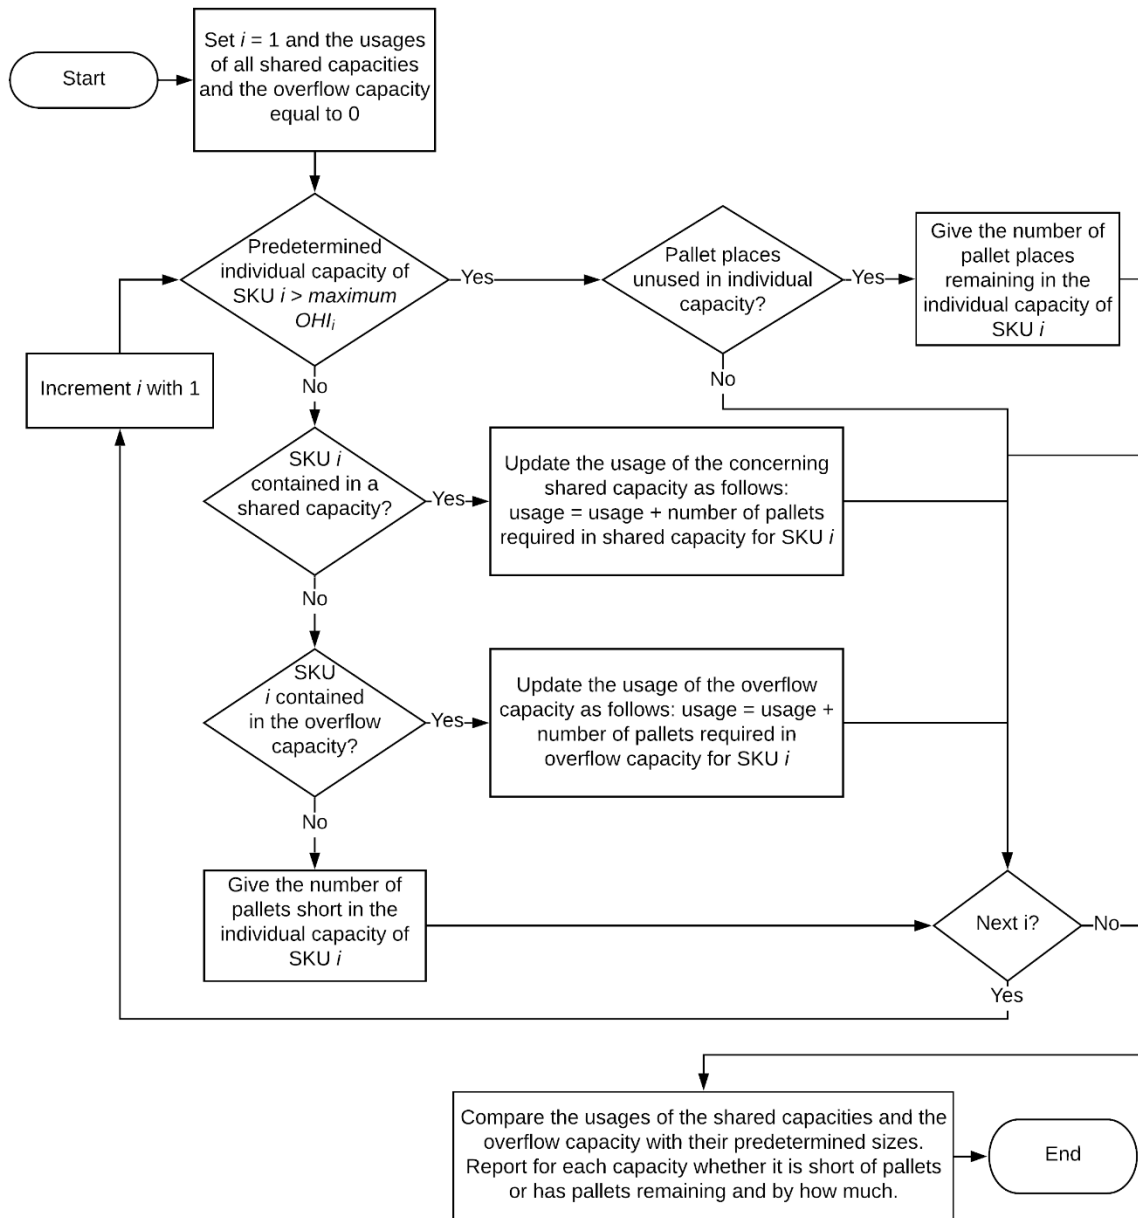


Figure 4.7 The capacity checking procedure of the buffer inventory capacity check.

Discussion

Although this extension checks the buffer inventory capacity, there are factors that affect the accuracy of the check. Namely, some assumptions are made which will not always hold in practice. The check should therefore only be used to get an idea of the usage of the buffer inventory capacity. Besides, it can be seen as recommendation to improve the design and size of the buffer inventory. We will now list several factors that affect the accuracy.

The main factor that affects the accuracy of the check is the assumption that the maximum *OHI* is present in the buffer inventory for all SKU's at the same time. This means that the buffer inventory is filled with the maximum quantity that is expected to be in the buffer inventory for all SKU's simultaneously. This assumption is unreasonable since there exist replenishment cycles with different lengths. If SKU A has a different replenishment cycle length than SKU B, then SKU A is replenished at different times compared with SKU B. Consequently, both SKU's have their maximum

OHI at different times. Hence, the scenario of the assumption will never happen in reality. However, the advantage of the assumption is that it assumes the worst case scenario. Namely, the usage of the capacity can never be worse than when the maximum *OHI* of every SKU is stored. Thus, if the buffer inventory capacity is able to store the worst case scenario, then it is able to store all other scenarios as well. It is therefore not immediately a problem if the check indicates that the buffer inventory capacity is too small, since the check is always on the safe side.

There are also two factors that affect the value of the maximum *OHI* and thus indirectly the usage of the buffer inventory capacity. We have seen in Equation 4.13 that the maximum *OHI* consists of 2 components. The first one is the order-up-to-level S . As already mentioned, the assumption is that the order-up-to-level S is never exceeded with replenishments. At the same time, we give the user the freedom to order a larger quantity than required, which results in exceeding S . This has the consequence that the maximum *OHI* (and thus the capacity usage) will be higher than expected. The second component is $E(y_i)$, which is the expected demand during the lead time. Since this is just an expected value, the real value can be either higher or lower, resulting in a lower or higher maximum *OHI* (and thus capacity usage) respectively.

4.2.5. Outstanding orders check extension

This subsection describes the last extension to the inventory control system: the outstanding orders check. Just like the buffer inventory capacity check, this check is tailor-made to the way that production happens at Mainfreight. There is namely additional information in the form of outstanding orders, which is currently not used by the inventory control system. Since these outstanding orders contain valuable information for the inventory control, a so-called outstanding orders check, in which they are used, is implemented as an extension to the inventory control system.

We analysed the production planning process in Section 2.1. It appeared that production takes place according to a week planning. This planning is based on a list with orders to be built. This list is the outstanding orders file, an Excel file which can be extracted from the WMS. This file is updated with new orders from time to time by the administration department. Although only the outstanding orders for one week of production are currently required, there are generally already more orders included in the outstanding orders file. Not only does this file contain the order numbers of all outstanding orders, it also contains all part requirements for these orders. This information is used in the outstanding orders check.

The principle behind the outstanding orders check is as follows: When you can determine which order numbers are going to be built in a certain period, you exactly know how much of every part you need to produce these orders (from the outstanding orders file). Besides, you already know how the future replenishments are going to look like (from the replenishment schedule), so you know for each part when it is going to be ordered and to which level it is going to be ordered. Thus, you also know the length of the period that the stock should last before it is replenished again. Based on the information of the outstanding orders file, we can predict if the decisions of the replenishment schedule for a certain time period will be sufficient to cover the requirements of the outstanding orders that are scheduled for that time period.

In other words, we use the outstanding orders check to check the results of the replenishment schedule. Namely, this schedule is completely based on stochastic demand (the Poisson and Compound Poisson demand distributions). However, the demand information from the outstanding orders file is deterministic demand (if we know which orders are scheduled for a certain time period). Nevertheless, the orders can only be scheduled for a small time period (1-2 weeks), whereas the

stochastic demand is used to determine the replenishment schedule for a long time period into the future. The outstanding order check will therefore only work for a small time period into the future. Since the time period between consecutive review/replenishment moments (base period) is also small, we propose to do the outstanding order check at every review/replenishment moment for one base period into the future. Therefore, we should be able to give all the order numbers of the orders which are going to be built in the next base period. The outstanding orders check can then be done for every SKU controlled by the inventory control system, whether it is only reviewed, reviewed and replenished or neither.

We will now clarify the outstanding orders check further with the help of an example. Suppose that it appears from the replenishment schedule that the optimal time period between consecutive review/replenishment moments is R_0 weeks. Moreover, suppose that SKU 1 is reviewed at every review/replenishment moment. That is, SKU 1 has a review period of R_0 weeks ($R_1 = R_0$). Suppose that we arrived at a replenishment date, where we are going to review the inventory of SKU 1 to decide whether it needs to be replenished. Additionally, we are doing the outstanding orders check on SKU 1. The goal of this outstanding orders check is to predict whether the inventory level after the current replenishment has arrived, is sufficiently high to cover the demand of the next R_0 weeks. To this end, we define:

- CIL_i , the buffer inventory level of SKU i at the current moment (just before the current replenishment takes place)
- $StartIL_i$, the buffer inventory level of SKU i at the start of the new base period (just after the current replenishment has arrived)
- $EndIL_i$, the buffer inventory level of SKU i at the end of the new base period (just before the new replenishment takes place), as expected according to the selected outstanding orders
- $OOQty_i$, the required quantity of SKU i according to the outstanding orders that are selected for the new base period

The goal is to find $EndIL_i$, which can be calculated with Equation 4.14.

$$EndIL_i = StartIL_i - OOQty_i \quad (4.14)$$

The prediction is then given by the value of $EndIL_i$: when it is zero or positive, then we do not expect a stockout to occur in the new base period. However, when it is negative, then we expect a stockout to occur and should ask ourselves whether $StartIL_i$ is high enough.

Figure 4.8 gives the situation described above. The blue line shows the inventory level of SKU 1 and it becomes clear that SKU 1 is replenished at every replenishment date (review period of SKU 1 $R_1 =$ base period R_0). Please note that we assume a lead time of zero here, which means that there is no time between the placement of the replenishment order and the arrival of that order. Furthermore, the roles of CIL_i , $StartIL_i$, $EndIL_i$ and $OOQty_i$ become clear in the figure. When we review the inventory of SKU 1 at the current replenishment date, we see that the CIL_1 is below S_1 . This means that the inventory of SKU 1 is ordered up to level S_1 . Therefore, the new base period will start with an inventory level of S_1 , so $StartIL_1$ will be equal to S_1 . The outstanding orders check then needs to find $OOQty_1$ for the new base period. It can then predict $EndIL_1$ by subtracting $OOQty_1$ from $StartIL_1$. In Figure 4.8, we have given three possible values for $OOQty_1$ (and thus also for $EndIL_1$). For the two lower values of $OOQty_1$, you can see that $EndIL_1$ stays positive, which indicates that a stockout for SKU 1 is not expected to happen. For the highest value of $OOQty_1$, the $EndIL_1$ becomes negative, which means that a stockout for SKU 1 at the end of the new base period is likely. Since the next replenishment order only occurs after the new base period (and thus the stockout) has passed, it is

wise to take action at this replenishment and order a larger quantity than the required order quantity proposed by the inventory control system. This quantity should have at least the size of the already required order quantity + the quantity that $EndIL_1$ is below zero.

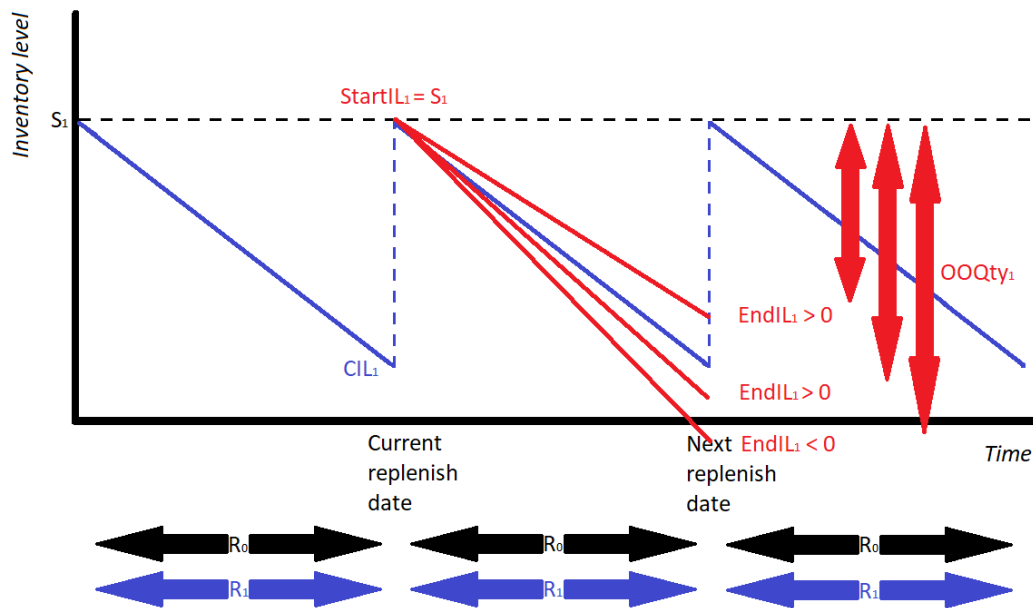


Figure 4.8 The inventory level of SKU 1 over time is given by the blue line. A graphical representation of the outstanding order check is given in red. This is the standard situation where $CIL_1 < S_1$, so $StartIL_1 = S_1$.

Above we have given the standard situation of a SKU that is replenished at every replenishment date. This is a standard situation since the CIL_i of a SKU is expected to be below S_i at a replenishment date. However, there are some exceptional situations where this is not the case. For example, when the inventory control system is used for the first time, it could be that the CIL_i of a SKU is much larger than S_i . This is because the inventory of SKU i is not controlled by the system yet, so it can have any level (thus also very high). A second exceptional situation is when the order-up-to-level S_i is exceeded with the last replenishment by a very large amount. This means that, even with the demand of the last period subtracted from it, CIL_i will be higher than S_i . This situation is described in Figure 4.9, again for the same SKU 1 with review period $R_1 =$ base period R_0 . Again, the blue line shows the inventory level of SKU 1. When we review the inventory of SKU 1 at the current replenishment date, we see that CIL_1 is higher than order-up-to-level S_1 . This means that we will not replenish SKU 1 at this replenishment date. Therefore, no action is taken for SKU 1 at this replenishment date, it remains only an inventory review. Consequently, the new base period will not start with an inventory level of S_1 , but will simply start with CIL_1 , so $StartIL_1 = CIL_1$. The rest of the procedure of the outstanding orders check is the same as with the standard situation, so find $OOQty_1$ and calculate $EndIL_1$. Again, three possible values for $OOQty_1$ and $EndIL_1$ are given in Figure 4.9, where only the highest $OOQty_1$ results in a negative $EndIL_1$ (stockout). However, the probability that a stockout occurs is lower compared to the standard situation, since the $StartIL_1$ is higher than in the standard situation. Nevertheless, if the outstanding orders check predicts a negative $EndIL_1$, it is wise to still order SKU 1 even if it is not proposed by the inventory control system to order the SKU. The order quantity should be at least equal to the amount that $EndIL_1$ is negative.

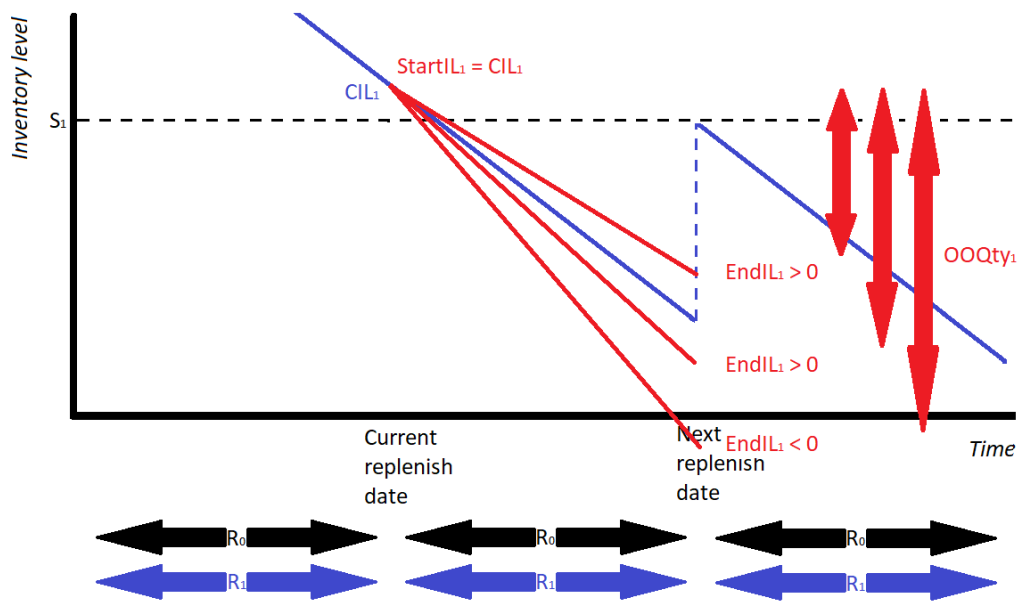


Figure 4.9 The inventory level of SKU 1 over time is given by the blue line. A graphical representation of the outstanding order check is given in red. This is the exceptional situation where $CIL_1 > S_1$, so $StartIL_1 = CIL_1$.

The two situations just described, concerned a SKU, SKU 1, which is reviewed at every replenishment date ($R_1 = R_0$). However, there are also SKU's that are not reviewed at every replenishment date. Suppose that SKU 2 is reviewed only half the time, so at every second replenishment date ($R_2 = 2 * R_0$). Suppose that SKU 2 is not reviewed (and thus not replenished) at the current replenishment date according to the replenishment schedule. Although we do not replenish SKU 2 at this moment, we can still do the outstanding orders check for this SKU. Therefore, we still need to review the inventory level of SKU 2. The $StartIL_2$ will always simply be equal to CIL_2 , since we are not concerned with ordering up to a level of S_2 at this moment. This situation is given in Figure 4.10. The $OOQty_2$ and the $EndIL_2$ are determined in the same way as before. We have again given three possible values for them, with only the highest value of $OOQty_2$ resulting in a negative $EndIL_2$. Only in that case, we propose to order SKU 2 even if it should not be ordered according to the inventory control system. The order quantity should be at least equal to the amount that $EndIL_2$ is negative.

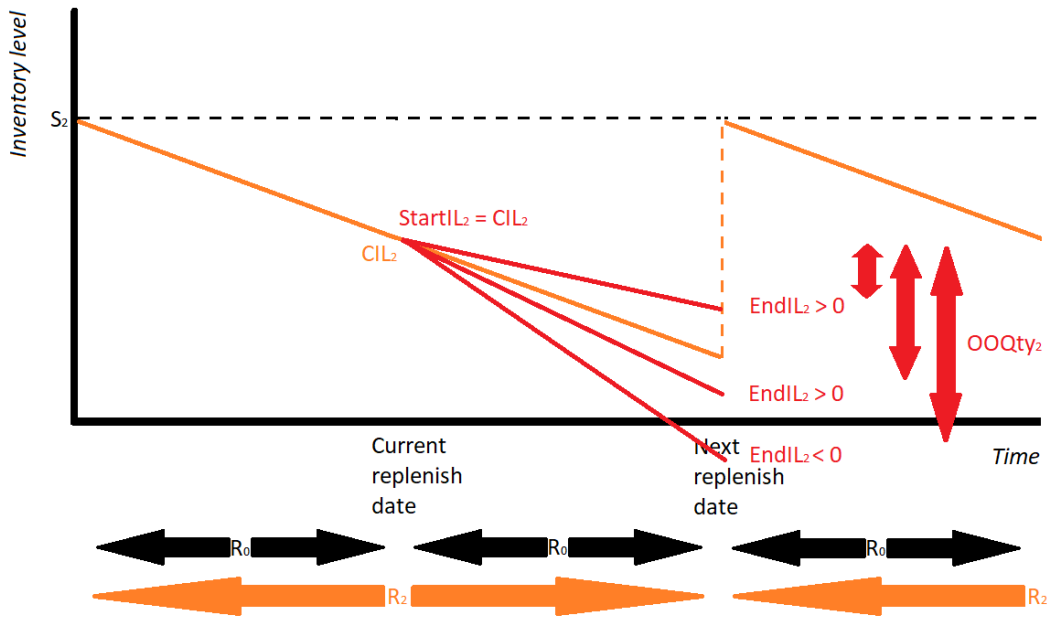


Figure 4.10 The inventory level of SKU 2 over time is given by the orange line. A graphical representation of the outstanding order check is given in red. This is the situation where a SKU is not ordered at the current replenishment date, so $StartIL_2 = CIL_2$.

In the situation above, we have given just one example of a SKU which is not replenished at every replenishment date, but only at every second replenishment date. However, the same applies to all other SKU's which are not replenished at every replenishment date, so also for the SKU's that are replenished at every third/fourth/fifth etc. replenishment date. Therefore, we can conclude that the outstanding orders check can be executed at *every replenishment date* for *every SKU of the inventory control system*, independent of their review/replenishment moments.

In conclusion, the ultimate goal of the outstanding orders check is to predict stockouts by finding the value of $EndIL_i$. The procedure to find it is roughly as follows (it should be executed at each replenishment date):

1. Determine $OOQTY_i$ for every SKU i
2. Determine $StartIL_i$ for every SKU i
3. Find $EndIL_i$ by subtracting $OOQTY_i$ from $StartIL_i$ for every SKU i

The first step is to determine the $OOQTY_i$ for every SKU i . For this, we need to determine the order numbers which are going to be built in the next base period. Furthermore, the outstanding orders file is needed, so the required quantities for every SKU i can be found for the selected order numbers. The sum of the required quantities of a SKU i over the selected order numbers will form the $OOQTY_i$.

The second step is to determine the $StartIL_i$ for every SKU i . As we have seen in the three situations described before, $StartIL_i$ depends on two factors. The first factor is whether the SKU is reviewed or not, which can be derived from the replenishment schedule. If so, we have to look at a second factor: whether the SKU is replenished or not. This depends on the order-up-to-level S_i , which can be found from the replenishment schedule, and the current buffer inventory level CIL_i , which should be calculated from the count per comp file. When CIL_i is less than S_i , then SKU i is replenished till level S_i , so $StartIL_i$ becomes S_i . When CIL_i is greater than or equal to S_i , then SKU i is not replenished. As SKU i is only reviewed in that case, nothing changes to the inventory level, so $StartIL_i$ becomes CIL_i . The same goes for a SKU that is not reviewed in the first place: nothing happens to the inventory level, so $StartIL_i$ becomes CIL_i .

The third step is to find $EndL_i$ for every SKU i by using Equation 4.14. For this, we need the values of $OOQty_i$ (determined in step 1) and $StartL_i$ (determined in step 2) for every SKU i . We can predict stockouts by looking at the value for $EndL_i$. Namely, a stockout of SKU i in the next base period is expected if the value of $EndL_i$ is negative, because $EndL_i$ represents the inventory level of SKU i at the end of the next base period.

The outstanding orders checking procedure, as described above, is graphically represented in the flowchart of Figure 4.11.

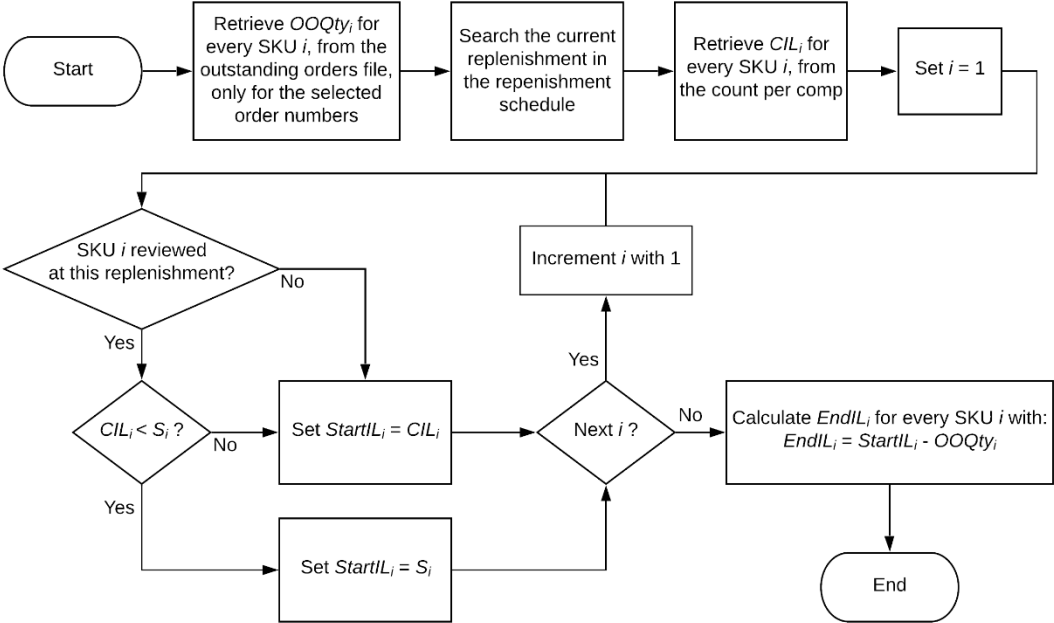


Figure 4.11 The outstanding orders checking procedure of the outstanding orders check.

Discussion

This extension predicts if stockouts will occur when the replenishment decisions that follow from the replenishment schedule are followed. However, there are assumptions that affect the accuracy of this prediction. We will list some of these assumptions now.

The first assumption we made in the outstanding orders check is the assumption of a lead time equal to zero, so we assume no time between the placement and delivery of the replenishment order. This makes the definition and calculation of the variables easier. Besides, it is easier to do the outstanding order check together with a replenishment, for a period until the next replenishment (base period). However, we have seen that lead time is not zero in reality. Nevertheless, the actual lead time is close to zero (one day at most), so this assumption seems reasonable and seems to have only a minor effect on the accuracy of the prediction.

A second assumption that might affect the accuracy of the prediction, is the assumption that all orders that will be built in the next base period, are already known precisely before the base period starts. The outstanding orders check does the prediction based on the $OOQty_i$ resulting from the selected outstanding orders. If changes in the production planning occur during the base period (so other orders are eventually built), then the determined $OOQty_i$ and thus the prediction of $EndL_i$ will not be right. It will therefore be a recommendation to the company to do a proper production planning at a replenishment date for the next base period and, more importantly, stick to this planning.

The last assumption that might affect the accuracy of the prediction, is the assumption that the user orders up to precisely level S_i when SKU i needs replenishment. This assumption is made since we always set $StartIL_i = S_i$ when SKU i is replenished. However, we give the user the freedom to exceed S_i with a replenishment, which means that $StartIL_i$ will be greater than S_i in that case. Nevertheless, this will not give problems since this results in a higher $EndIL_i$ in reality than in the outstanding orders check (with $StartIL_i = S_i$). Hence, if there is a positive $EndIL_i$ (no stockout) according to the outstanding orders check, then there is certainly a positive $EndIL_i$ when S_i is exceeded with the current replenishment. In other words, the outstanding orders check gives the worst case scenario (order up to precisely level S_i).

We can also see in Figure 4.8 why we give the user of the inventory control system the freedom to exceed S_i with a replenishment order. In that case, the $StartIL_i$ level is higher than it should be according to the (R_i, S_i) model. This exceedance of the inventory level has no lasting effect, since it is automatically corrected by the system at the next replenishment, because of the variable order quantity under the (R_i, S_i) model. Namely, $EndIL_i$ will be higher than expected by the (R_i, S_i) model, so the system will propose to order a smaller quantity than expected so S_i is not exceeded again. This is the case since the inventory model strives for an $StartIL_i$ of S_i at every replenishment date. If there would be a fixed order quantity, the order quantity would not be adjusted and S_i would probably be exceeded again. Since the variable order quantity has this nice advantage, we can allow the user to exceed S_i with a replenishment order.

The situation is basically the same in Figure 4.9, where $StartIL_i$ is also greater than S_i . Only here it is not because S_i is exceeded with the replenishment, but because CIL_i is already higher than S_i (so there is no replenishment). Here, we also see that the inventory model automatically corrects this exceptional situation at the next replenishment by proposing to order up to level S_i again. In other words, if the proposed required order quantity is respected at the next replenishment date, then the exceptional situation becomes the standard situation again.

Furthermore, one could argue why we implement an outstanding orders check next to the (R_i, S_i) inventory model in the first place. Namely, the inventory model already determines the replenishment decisions based on a demand probability distribution and subject to a certain service level (probability that the replenishment cycle ends with no stockout). Therefore, one should expect that stockouts will *hardly* occur if the right demand distribution and a high service level are chosen. However, it is not possible to prevent *all* stockouts because of these two factors. Firstly, the (parameters of the) demand distribution might not fit the demand well, which leads to more stockouts if the demand is underestimated by the distribution. Secondly, it is not possible to set the service level at 100%, since it is always possible that an outlier is so high that it is not taken into account. The outstanding orders check does not have these problems, since it makes use of deterministic demand information, so you know the demand (and thus how much inventory you need to prevent a stockout) for sure. The (R_i, S_i) inventory model, on the other hand, tries to find a schedule for the long term and can only estimate the future demand. It therefore determines a policy which prevents stockouts in a large number of cases, but it is not prepared for high outliers. The outstanding order check is able to detect such high outliers beforehand, but only for the length of a base period. However, this is early enough to take action, so the outstanding orders check is a useful addition to the (R_i, S_i) inventory model.

In conclusion, the outstanding orders check should only be used as an additional check to see if the replenishment decisions that follow from the (R_i, S_i) inventory model (still) hold for the next base period. It can be used as a check since the outstanding orders provide an extra (deterministic) demand input, next to the (stochastic) demand input of the (R_i, S_i) inventory model. In that way, it can help to prevent foreseeable stockouts. A stockout can be foreseen because the check can

compare the replenishment decisions (from the stochastic (R_i, S_i) inventory model) with the deterministic demand of the outstanding orders. If a stockout is foreseen, the outstanding orders check will indicate this and the user of the inventory control system can take action accordingly.

4.3. Implementation of the inventory control system in Excel

This section describes the implementation of the inventory control system with all its extensions. The implemented inventory control system will follow the methodology of Section 4.2. It was chosen to do the implementation in Microsoft Excel, for a number of reasons. Firstly, Excel is a program that the company already uses a lot, so the employees are used to work with it. Secondly, I am able to program the models behind the inventory control system in Visual Basic for Applications (VBA), the programming language of Excel (and other Microsoft Office programs). Last but not least, Excel is very appropriate for an implementation of this inventory control system since it consists of both worksheets and a programming language to automate tasks in the worksheets. This means that the worksheets can act as a user interface where the user can give input and receive output of the inventory control system. The model behind the inventory control system can then be developed in the programming language (VBA). Subsequently, the user can control the (model behind the) inventory control system via the worksheets.

The inventory control system has been implemented in a single Excel macro-enabled workbook file. However, two versions of this file are made: a version for the university (UT version) and a version for the company (company version). Both versions are in principle the same, except for the terms used in the worksheets. The UT version contains terms which are in line with the methodology of Section 4.2, whereas the company version contains different terms which are understandable for the company. Therefore, there are only changes in the worksheets (user interface), the underlying VBA code is the same. Hence, we only have one VBA code for the two versions. This code is given in Appendix E: VBA code of the inventory control system implementation in Excel.

Since the company has to work with the company version of the Excel file, the manual has been based on this file. This manual is given in Appendix D: Manual of the inventory control system. The manual gives a very detailed explanation of the user interface and describes how the implemented inventory control system should be used.

In this section, we will focus on the implementation itself, so we are going to explain how the methodology of Section 4.2 has been translated to the VBA code of the inventory control system Excel file. Furthermore, we will show the user interface and its relation with the code. From now on, we will only consider the UT version of the Excel file, since it is more in line with the terminology used in the methodology. However, as we pointed out, the model behind the inventory control system is the same for both versions.

When we say “model behind the inventory control system”, we mean all the steps to come from the input to the output of the inventory control system. It was chosen to implement this model as two workflows in VBA: Workflow 1 and Workflow 2. Both workflows have approximately the same structure: both are programmed in their own VBA Module, both consist of multiple Sub procedures (Subs) and both are controlled by an own user interface (UserForm). In total, this means that the VBA model consists of two user forms and five subs divided over two modules. The division of the subs over the modules is given in Table 4.6. We mentioned before that the implemented inventory control system will be made according to the methodology of Section 4.2. This means that the implemented inventory control system will also consist of the demand model, the inventory model and the three extensions. You can see which parts of the inventory control system are contained in which sub. The methodology subsection in which the part is explained, is given between parentheses.

Table 4.6 The division of the subs over the modules.

| Module | Sub | Contains |
|--------|------------------------|--|
| 1 | Workflow1 | Part of the inventory model (4.2.2) |
| | HSTSP | Demand model (4.2.1) Part of the inventory model (4.2.2) |
| | CapacityCheck | Extension which does the buffer inventory capacity check (4.2.4) |
| 2 | Workflow2 | Extension which calculates the replenishment decisions (4.2.3) |
| | OutstandingOrdersCheck | Extension which does the outstanding orders check (4.2.5) |

Please look at Figure 4.12, where flowcharts of both workflows are given. Here you can see how the parts of the inventory control system are contained in the workflows (given by coloured rectangles). Furthermore, you can see the relations between subs, between a sub and a user form and between the workflows. We will first explain the relation between the workflows on the basis of these flowcharts. Subsequently, we will explain the implementation of the workflows in more detail, also referring to these flowcharts. We will end this section with a short conclusion on the inventory control system implementation.

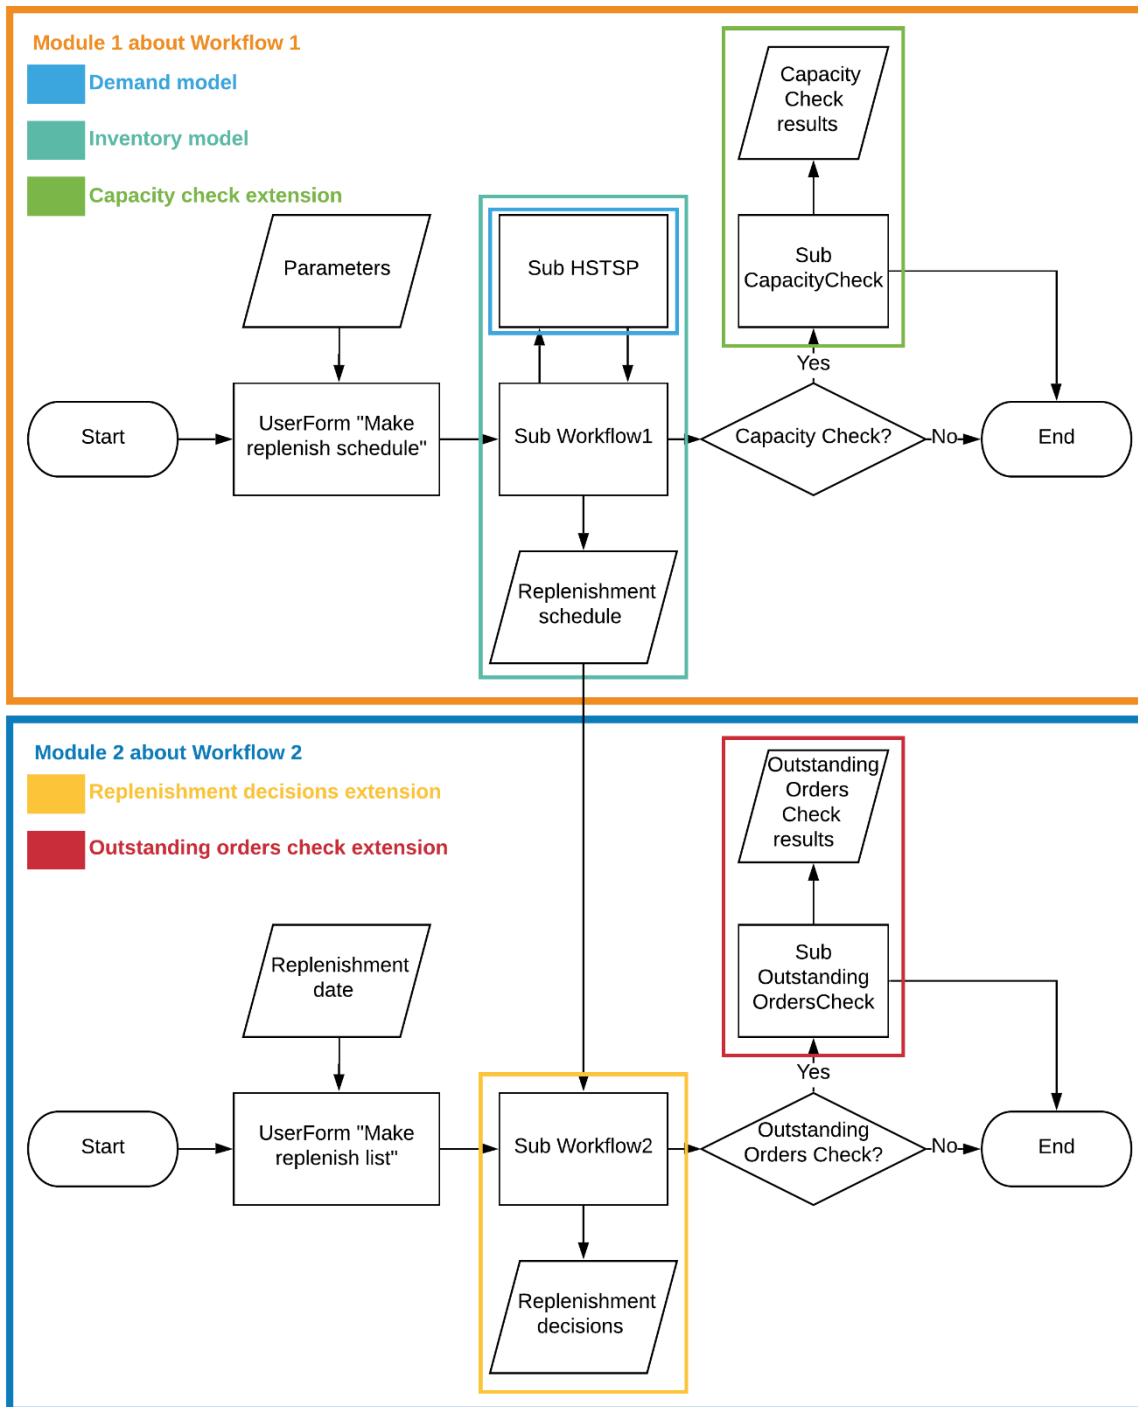


Figure 4.12 The flowcharts of workflow 1 and 2

4.3.1. Relation between the workflows

The model behind the inventory control system consists of two workflows. The first workflow calculates the *replenishment schedule* and the second workflow calculates the *replenishment decisions* that follow from the replenishment schedule. You should view the replenishment schedule as a long-term replenishment strategy, which gives the user a list with replenishment dates and the SKU's which should be ordered at these dates. Furthermore, it gives the (fixed) order-up-to-levels for every SKU. In principle, this replenishment schedule should be made only one time for a time period till far in the future and therefore the workflow should be executed only rarely.

With the main output of the first workflow (the replenishment schedule), the user of the inventory control system can determine the replenishments in principle. However, the replenishment schedule is still not very easy to use for the user, since the user should manually check the schedule for the SKU's to replenish. Furthermore, the user should check the current inventory level for these SKU's one by one and then derive the replenishment decisions him-/herself. Hence, this work is time consuming and prone to error. Therefore, it has been automated by the extension that calculates the replenishment decisions. This extension has been implemented in the second workflow. This workflow evaluates the replenishment schedule at the right replenishment date and automatically checks the current inventory level to determine the exact replenishment decisions. This workflow should therefore be executed at every replenishment date. We can now see the relation between the workflows: the second workflow needs the replenishment schedule to calculate the replenishment decisions. The replenishment schedule is therefore an output of the first workflow and an input for the second workflow, see Figure 4.12. Every time the second workflow is executed (often), it uses the same replenishment schedule (once made by the first workflow), until a new replenishment schedule is made (rarely).

4.3.2. Implementation of workflow 1

We will now have a closer look at the implementation of the first workflow. We will show the user form, subs and the relations between them in more detail. We can see in the flowchart of workflow 1 (Figure 4.12), that the workflow starts with the user form "Make replenish schedule" and the input of parameters. To this end, worksheet "Input replenish schedule" (Figure 4.13) is used.

| i | SKU | μ_i | σ_i^2 | $t_{i,max}$ | k_i | h_i | P_i | L_i | $Y_{i,max}$ |
|----|-----|-------------|--------------|-------------|-------|-------|-------|-------|-------------|
| 1 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 2 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 3 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 4 | 50 | 6.32873113 | 1.373819231 | 0.98 | 0.2 | 50 | | | |
| 5 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 6 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 7 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 8 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 9 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 10 | 50 | 6.32873113 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 11 | 50 | 6.32873113 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 12 | 50 | 4.219154087 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 13 | 50 | 4.219154087 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 14 | 50 | 8.438308174 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 15 | 50 | 8.438308174 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 16 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 17 | 50 | 4.219154087 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 18 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 19 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 20 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 21 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 22 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 23 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 24 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 25 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 26 | 50 | 2.109577043 | 0.025207692 | 0.98 | 0.2 | 50 | | | |
| 27 | 50 | 2.109577043 | 0.025207692 | 0.98 | 0.2 | 50 | | | |
| 28 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 29 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 30 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |
| 31 | 50 | 2.109577043 | 0.126038462 | 0.98 | 0.2 | 50 | | | |

Figure 4.13 Worksheet "Input replenish schedule" contains the user form and parameters (input) for the first workflow.

The user should give a list with the SKU's that should be controlled by the inventory control system in this worksheet. Furthermore, (s)he should give the values of the parameters μ_i , σ_i^2 , $t_{i,max}$, k_i , h_i , P_i , L_i and $y_{i,max}$ for each SKU i and the value of K separately. The next section (Section 4.4) discusses the determination of values for these parameters. When the blue button in the worksheet is clicked, the user form shows up (Figure 4.14). This user form contains the instructions to input the parameters. Besides, the parameters R_0^{min} , R_0^{max} and ΔR_0 need to be set in this user form. The user form is programmed in such a way that only reasonable values with at most one decimal are allowed for these three parameters.

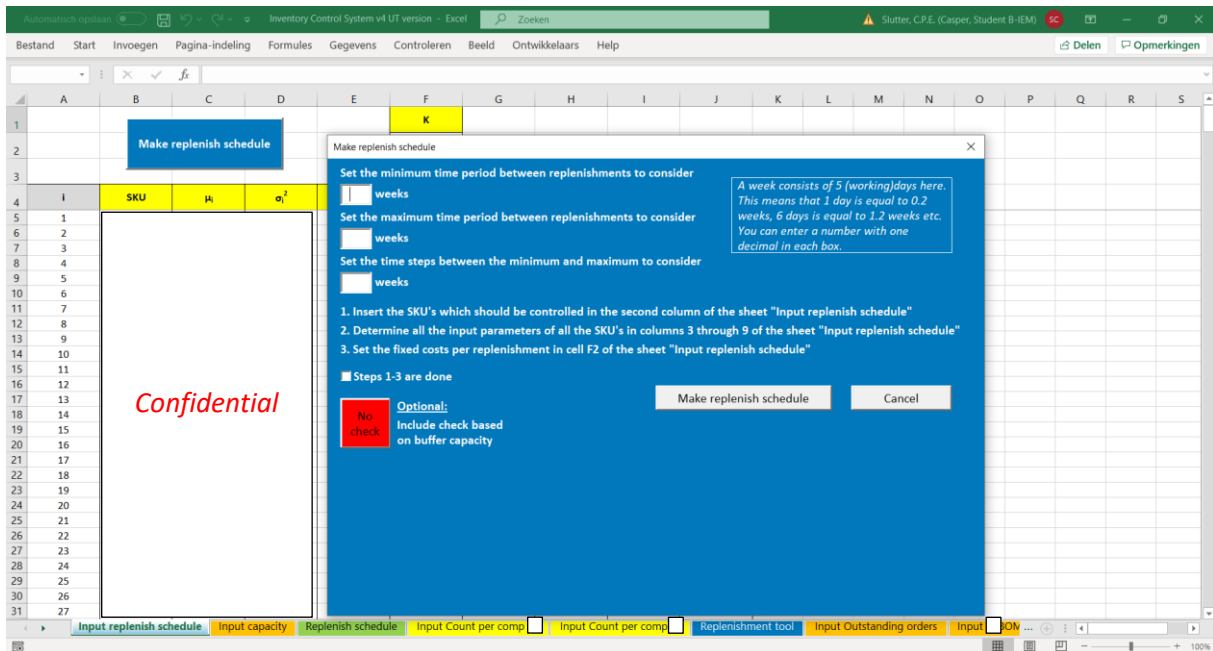


Figure 4.14 The user form of worksheet "Input replenish schedule" with which workflow 1 can be controlled.

Workflow 1 can now be executed (we come back to the optional capacity check later). To this end, the button "Make replenish schedule" should be clicked. The user form then calls sub Workflow1, as can also be seen in the flowchart of workflow 1 in Figure 4.12. When calling sub Workflow1, the user form passes on the values for R_0^{min} , R_0^{max} , ΔR_0 and a True/False for *CapacityCheckYesOrNo*, which concerns the capacity check (False for now).

Sub Workflow1 is the main sub of workflow 1. We can see in the flowchart of workflow 1 in Figure 4.12 that sub Workflow1 is part of the inventory model. In this role, it is responsible for calling sub HSTSP and making the replenishment schedule. Furthermore, it should call the sub CapacityCheck (or not). You can find the code of sub Workflow1 in Appendix E: VBA code of the inventory control system implementation in Excel. The code is graphically represented in the flowchart of Figure 4.15.

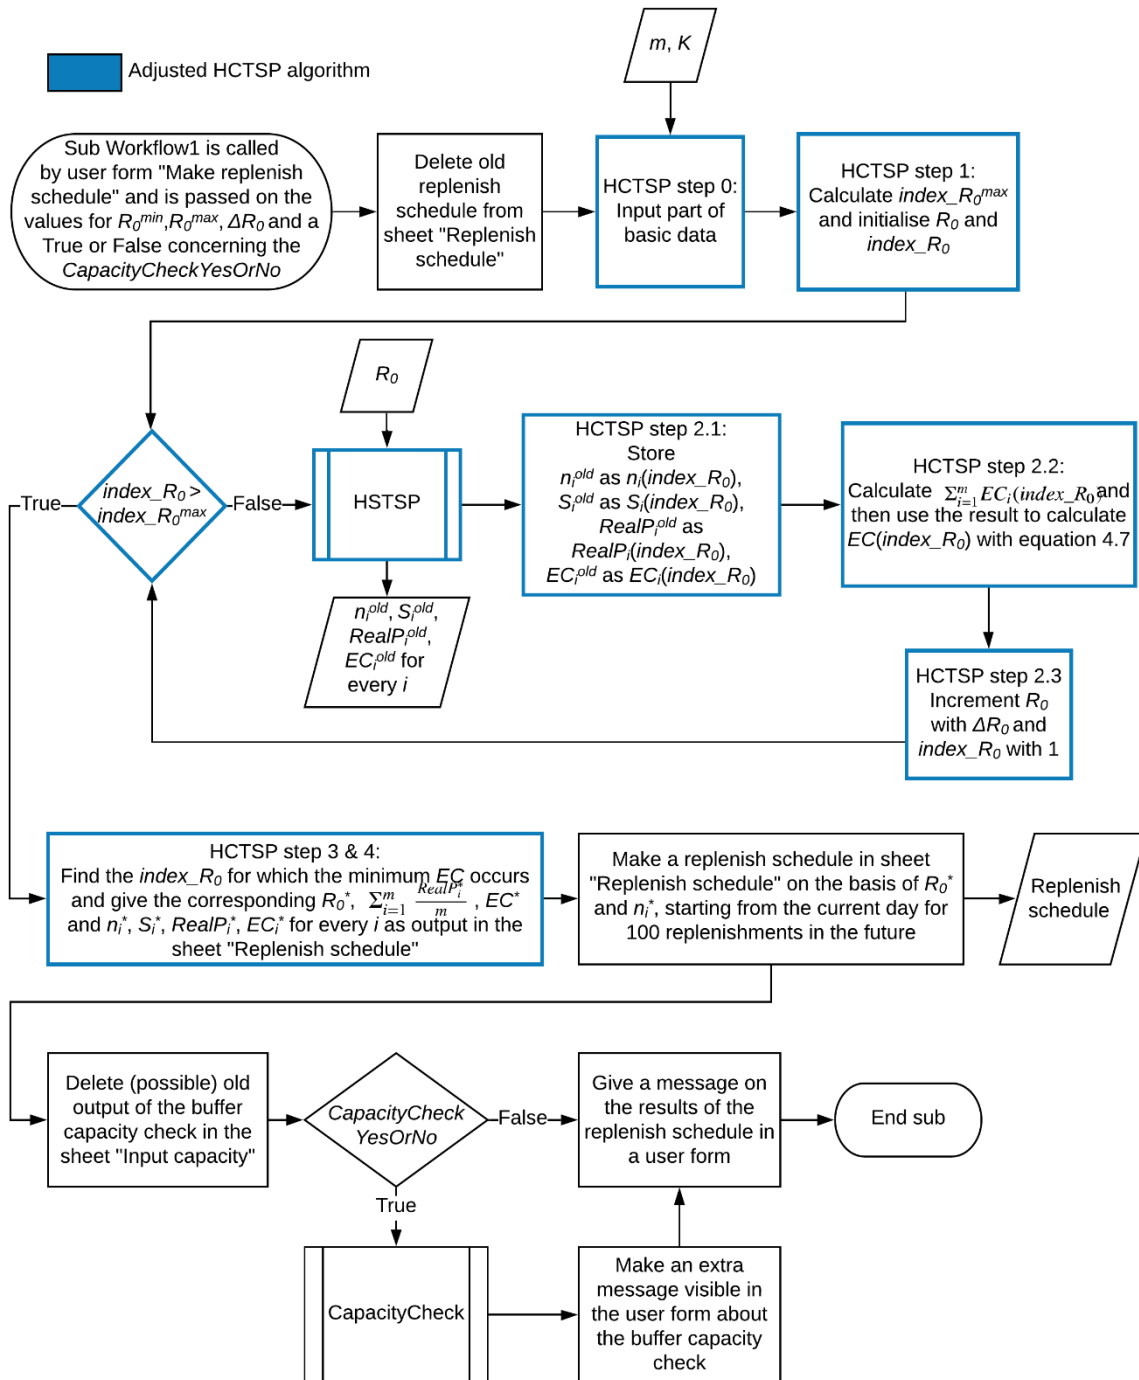


Figure 4.15 The flowchart of sub Workflow1

We have seen in Subsection 4.2.2 about the inventory model, that the inventory model can be implemented as two heuristic algorithms, the HCTSP and HSTSP algorithm from Fung et al. (2001). We also proposed an adjusted HCTSP algorithm, which performs better than the original one. Therefore, we have implemented the adjusted HCTSP algorithm in sub Workflow1, see the blue rectangles in Figure 4.15. When we look at Figure 4.15, we see that the sub starts (after deleting the old output) with all steps from the adjusted HCTSP algorithm. These steps follow the pseudocode of Figure 4.4 closely. One difference is visible in step 0, where we only input the parameters which are not SKU specific (so without index i). Another difference is step 1, where a part of the initialisation is

not needed. The last difference is the fact that step 3 and 4 of the pseudocode are combined in one step in the code.

For each iteration of R_0 in the adjusted HCTSP algorithm, we need to execute the HSTSP algorithm for all SKU's separately (so m times) in step 2.1. For this reason, it is convenient to implement the HSTSP algorithm in a different sub (called sub HSTSP), which can be called multiple times by sub Workflow1. However, we made sub HSTSP in such a way, that it executes the algorithm for all m SKU's at once, instead of just for one SKU. This means that sub HSTSP should be called only once per iteration of R_0 , instead of m times per iteration of R_0 .

When the HSTSP algorithm needs to be executed, sub HSTSP should be called by sub Workflow1. Since it should be executed for a certain value of R_0 , sub Workflow1 passes on this value of R_0 when calling sub HSTSP, which is also visible in Figure 4.15. Sub HSTSP is then executed with this value of R_0 . We can see in the flowchart of workflow 1 in Figure 4.12 that sub HSTSP contains the demand model and a part of the inventory model (HSTSP algorithm). In this role, it is responsible for calculating the demand and the optimal (R_i, S_i) for every SKU i given this demand and a certain R_0 . You can find the code of sub HSTSP in Appendix E: VBA code of the inventory control system implementation in Excel. The code is graphically represented in the flowchart of Figure 4.16.

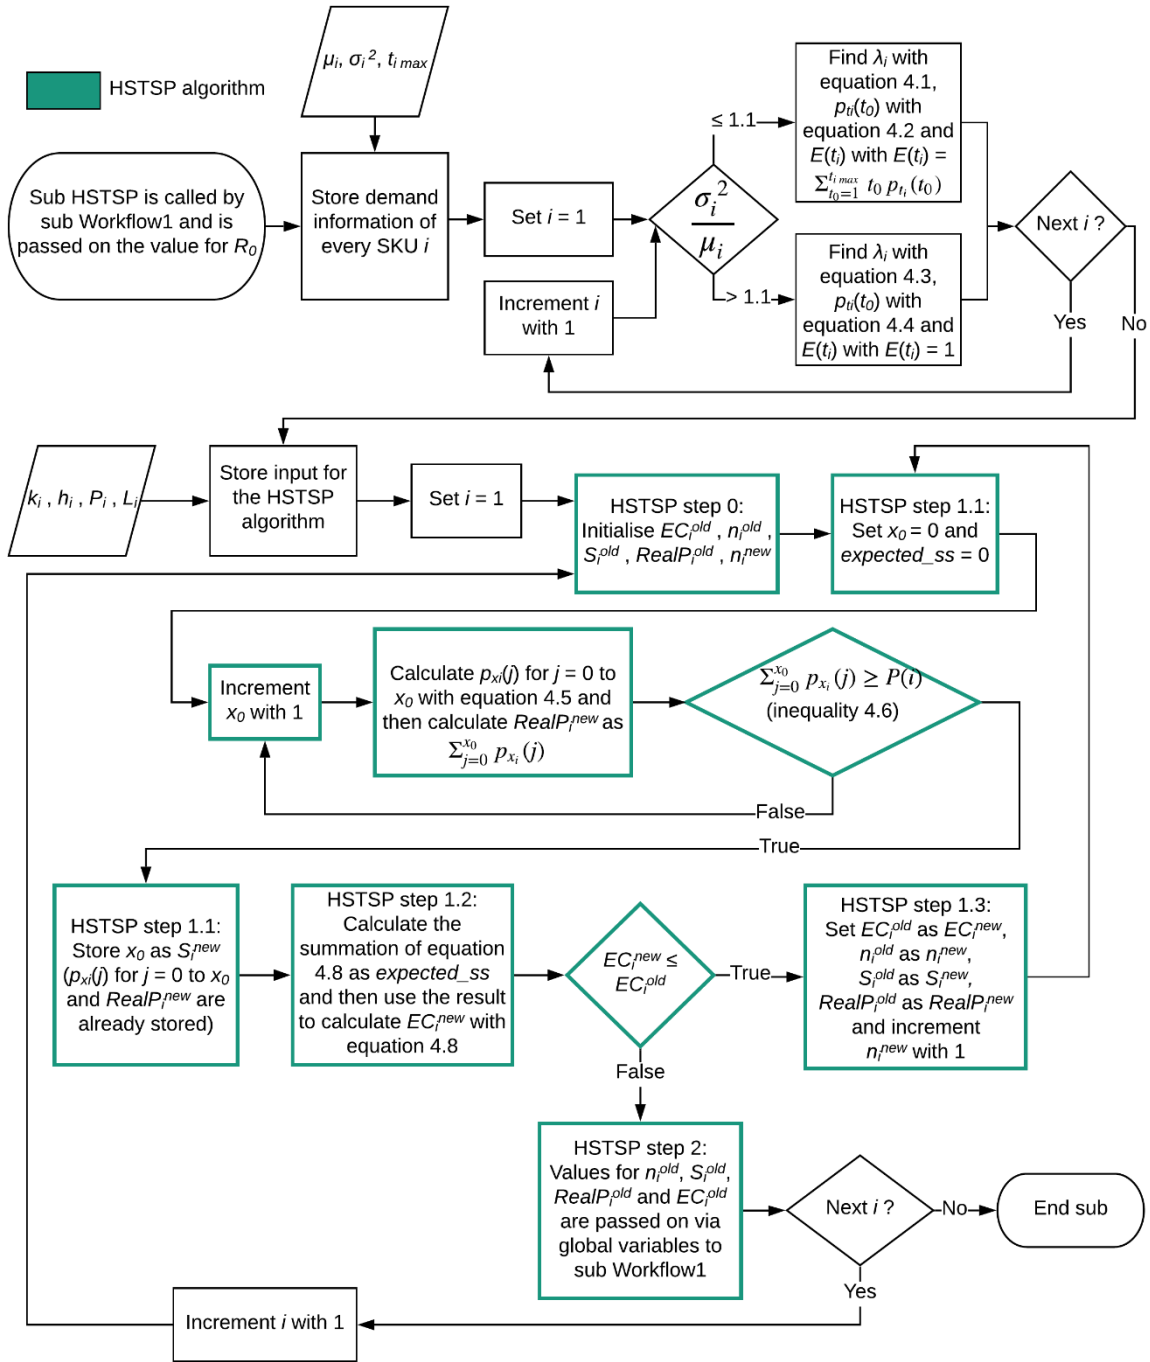


Figure 4.16 The flowchart of sub HSTSP

First of all, we want to note why we did not input all the parameters of step 0 of the HCTSP pseudocode (Figure 4.4) in step 0 of the implemented HCTSP algorithm. These parameters are namely just used in the HSTSP algorithm (k_i, h_i, L_i) or just calculated in the demand model ($\lambda_i, p_{it}(t_0)$). Therefore, they will be just inputted in sub HSTSP and not in sub Workflow1.

The demand model is implemented in sub HSTSP for the same reason. Namely, it determines the values of λ_i and $p_{it}(t_0)$ for every SKU i , which are both inputs for the HSTSP algorithm. We can see the demand model in the first row of the flowchart of sub HSTSP in Figure 4.16. After the input of the demand parameters, we see that based on the decision rule of the demand model, the SKU's are

given either the Poisson or compound Poisson demand distribution: SKU's with a ratio less than or equal to 1.1 get the Poisson distribution, which is given by λ_i from Equation 4.1 and $p_{ti}(t_0)$ from Equation 4.2. SKU's with a ratio greater than 1.1 get the compound Poisson distribution, which is given by λ_i from Equation 4.3 and $p_{ti}(t_0)$ from Equation 4.4. For convenience, $E(t_i)$ is already calculated in this step, although it is just used in the HSTSP algorithm (Equation 4.8).

Now that the demand parameters are all known (both the input parameters and output parameters of the demand model), we see in Figure 4.16 that the remaining parameters are stored for the HSTSP algorithm. Subsequently, we have implemented the HSTSP algorithm, see the green rectangles in Figure 4.16. The HSTSP implementation follows the HSTSP algorithm of Fung et al. (2001) (see the flowchart in Figure 4.1) closely. However, there are a few differences.

The first difference is that in Figure 4.1, the first step is to "Determine a R_0 ". However, in the implementation R_0 is given and passed on by sub Workflow1 and thus not determined in the HSTSP algorithm.

The next difference between the HSTSP algorithm of Figure 4.1 and the implementation is visible in HSTSP step 1.1: "Find the smallest integer S_i^{new} which satisfies Inequality 4.6 ...". This step is implemented as a Do ... Until loop with Inequality 4.6 as condition. This Do ... Until loop is visible in Figure 4.16 between the two rectangles with "HSTSP step 1.1" in it. The calculation of $p_{xi}(x_0)$, which is not part of the HSTSP algorithm, and HSTSP step 1.1 are done simultaneously in this Do ... Until loop. Hence, the probabilities $p_{xi}(x_0)$ are not calculated in advance, but during the search of S_i^{new} . In this way, no upper boundary for the calculation of the probabilities $p_{xi}(x_0)$ is required (like $t_i \max$ is the upper boundary for the calculation of the probabilities $p_{ti}(t_0)$). That is because on a given moment in the Do ... Until loop, the probability of $p_{xi}(j)$ is calculated and then the probabilities from $x_0 = 0$ up to and including $x_0 = j$ are summed to calculate the service level. If it is not high enough (not equal or higher than $P(i)$), then $p_{xi}(j+1)$ is calculated and added to the sum. Eventually, when the sum is high enough, the process stops and just the right amount of $p_{xi}(x_0)$'s are calculated. Therefore, calculating $p_{xi}(x_0)$ simultaneously with searching for the smallest integer S_i^{new} is very computationally efficient. In fact, the $P(i)$ forms the boundary that determines how many $p_{xi}(x_0)$'s should be calculated.

At the first iteration of the Do ... Until loop, the service level of $S_i = 1$ is already calculated and compared with $P(i)$. This means that the smallest integer S_i which can satisfy Inequality 4.6 is 1. Therefore, there is no possibility of $S_i = 0$ as outcome. However, an order-up-to-level of 0 does not make sense. Besides, the S_i is already initialised as 1 in step 0 of the HSTSP algorithm. The Do ... Until loop then evaluates, one by one, all integers greater than 1. This way of finding the smallest integer S_i which satisfies Inequality 4.6 can therefore be used.

Another (small) difference is in HSTSP step 1.2, where we calculate EC_i^{new} in two steps instead of one. Namely, we first calculate the summation of Equation 4.8 (*expected_ss*) and then use the result in Equation 4.8 to calculate EC_i^{new} .

The last (small) difference is in HSTSP step 2, where we also output EC_i^{old} because it is useful output. Another thing to note is that the outputs of HSTSP step 2 are stored as global variables. In this way, sub Workflow1 can easily retrieve the outputs of sub HSTSP.

We now turn back to sub Workflow1. We can see in the flowchart of Figure 4.15 that the output of sub HSTSP corresponds to the output described in HSTSP step 2 in the flowchart of Figure 4.16. This output is used by sub Workflow1 in the remainder of the HCTSP algorithm (step 2.1, 2.2, 2.3, 3 and 4). Eventually, as we can see in Figure 4.15, sub Workflow1 makes the replenishment schedule based on the output of the HCTSP algorithm. This replenishment schedule starts at the current day and is made for 100 replenishments in advance. The replenishment schedule is presented to the user in a new worksheet, called "Replenish schedule" (Figure 4.17).

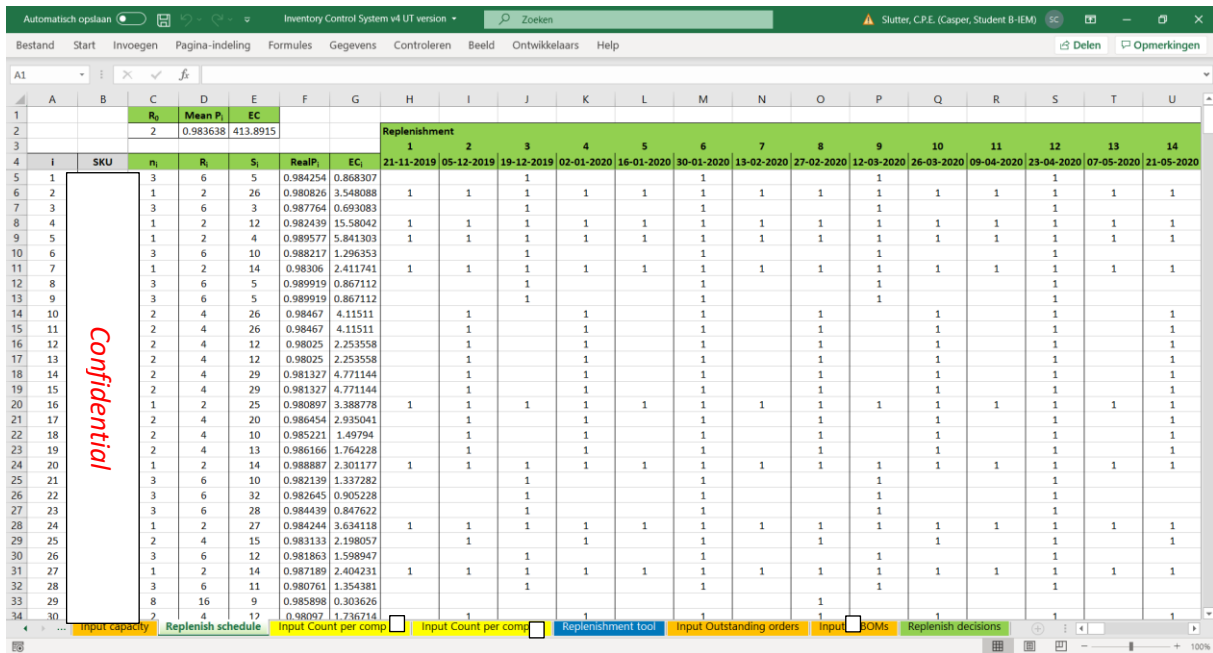


Figure 4.17 Worksheet “Replenish schedule” contains the replenishment schedule, the main output of workflow 1.

We can see the replenishment schedule on the right hand side (starting from column H) in Figure 4.17. Every column corresponds to one replenishment date (the header of the column). If there is a one in a certain cell, it means that the SKU of that row should be reviewed at the replenishment date of that column. As mentioned before, this replenishment schedule is made for 100 future replenishments, which means that the replenishment schedule consists of 100 columns. The other part of the replenishment schedule is the order-up-to-level S_i for every SKU i . These are visible in the fifth column of the worksheet (Figure 4.17). Next to S_i , we also present the values for n_i , R_i , $RealP_i$ and EC_i for every SKU i , because it is also useful SKU specific output of workflow 1. Besides, we give global output about the replenishment schedule (R_0 , $Mean P_i$ and EC) in the second row of the worksheet (Figure 4.17).

We see in Figure 4.15 that sub Workflow1 continues after it has outputted the replenishment schedule. The remainder of sub Workflow1 has to do with the buffer inventory capacity check extension. When we look back at the flowchart of workflow 1 in Figure 4.12, we see that this is indeed the last step of workflow 1. Furthermore, we see in Figure 4.12 that this capacity check is made optional. It is namely not necessary to do the check since it does not affect the replenishment schedule, but it gives useful insights about the feasibility of this schedule in terms of storage capacity.

Regardless of whether the capacity check is chosen or not, sub Workflow1 first deletes possible old output of the capacity check. Subsequently, it checks whether the capacity check should be carried out this time.

First, we need to go back to the user form of Figure 4.14. The user can choose in this user form with a toggle button whether the capacity check should be executed or not, the default setting is “No capacity check” (red toggle button). If the user wants to include the capacity check, (s)he should click the (currently red) toggle button. The toggle button then turns green, indicating that the capacity check will be executed (Figure 4.18).

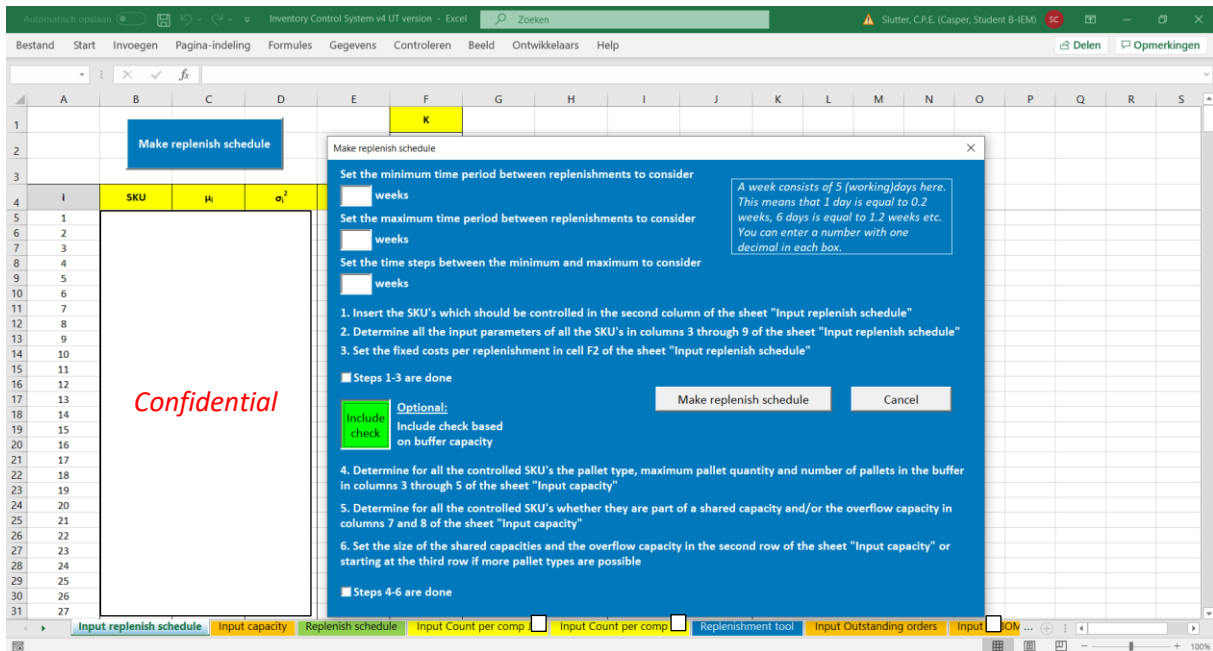


Figure 4.18 The user form of worksheet "Input replenish schedule" with the capacity check selected.

The toggle button assigns either True (include capacity check) or False (no capacity check) to the variable *CapacityCheckYesOrNo*. The user form passes on the value of *CapacityCheckYesOrNo* when it calls sub *Workflow1*. We can see in Figure 4.15 that sub *Workflow1* checks the value of *CapacityCheckYesOrNo* to decide if the capacity check should be carried out.

If the capacity check should not be executed (*CapacityCheckYesOrNo* = False), sub *Workflow1* gives a message on the results of the replenishment schedule to the user, after which sub *Workflow1* and thus workflow 1 is ended. However, if the capacity check should be executed (*CapacityCheckYesOrNo* = True), sub *Workflow1* calls a new sub, sub *CapacityCheck*. This sub contains the complete buffer inventory capacity check extension. It is implemented in a separate sub, so the code of the extension is separated from the code of the basic inventory control system. Besides, a separate sub for the extension is convenient since the extension is optional. In that way, it is easy to simply call the sub if it is required and do nothing if it is not required.

Before we discuss the sub *CapacityCheck*, we will first present a new worksheet, called "Input capacity" (Figure 4.19). This worksheet is only used for the capacity check extension, it contains input that is required by the capacity check. Furthermore, the output of the capacity check is also presented in this worksheet (Figure 4.20). However, worksheet "Input capacity" should only be considered when the capacity check is activated. Nothing is done with the input and no output is given when the capacity check is not activated.

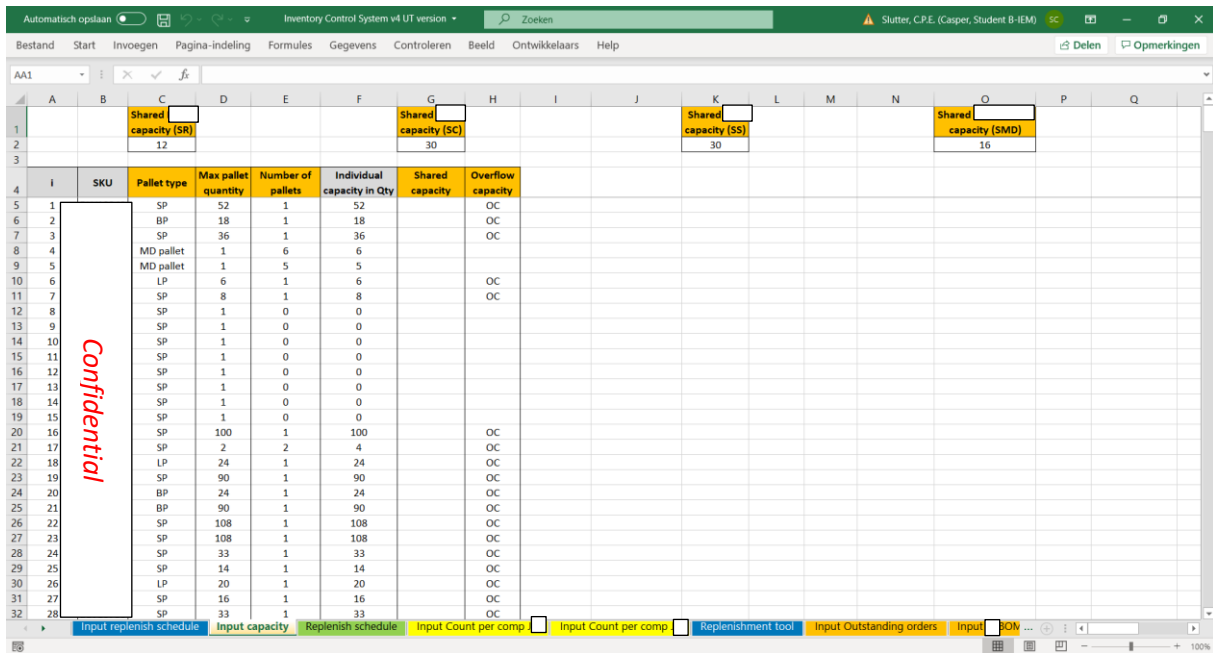


Figure 4.19 Worksheet “Input capacity” contains the input for the optional buffer inventory capacity check extension.

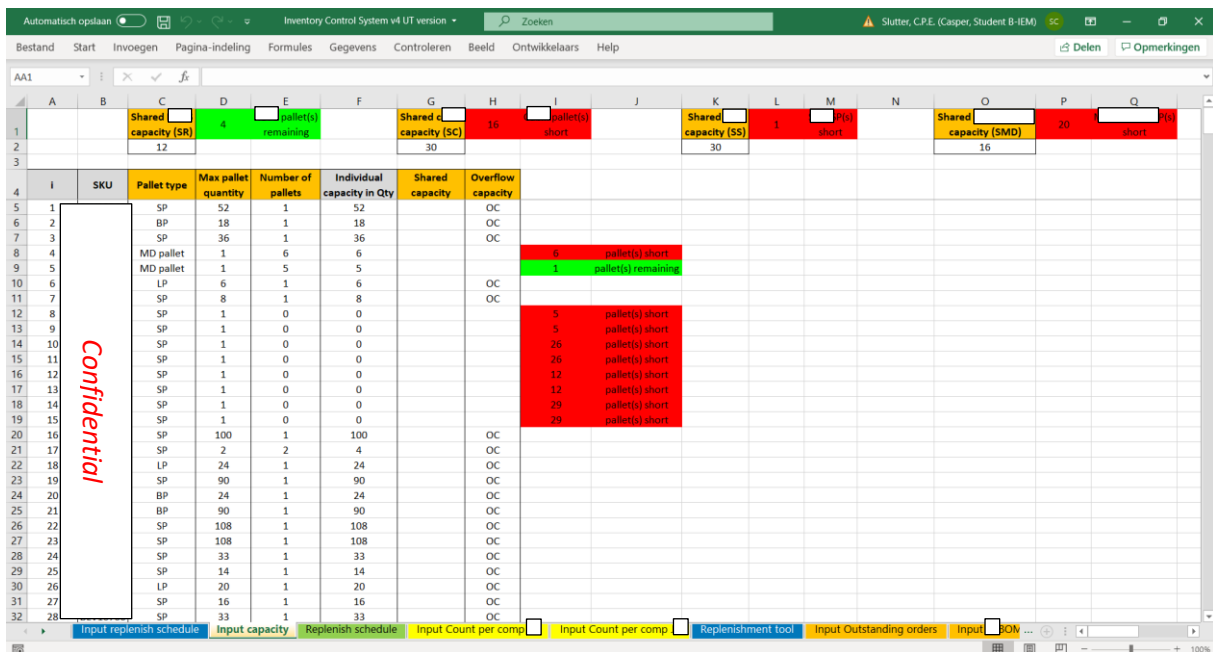


Figure 4.20 The output of the buffer inventory capacity check is also given in worksheet “Input capacity”, but only if the check is activated.

The buffer inventory capacity check extension is implemented according to the methodology in Subsection 4.2.4. We saw there that the buffer inventory can be divided into three capacities: individual capacities (per SKU), shared capacities and an overflow capacity (both per set of SKU’s). You can see the individual capacities in the sixth column of the table with all the SKU’s listed (Figure 4.19). In order to arrive at these sizes of the individual capacities, the user need to input several parameters in column 3 through 5. The individual capacities are then automatically calculated using an Excel formula. The shared capacities and overflow capacity are defined in the first row of the worksheet (Figure 4.19). Their capacity sizes need to be set manually in the second row of the corresponding column. Finally, the user needs to indicate in the seventh and eight column of the

table with all the SKU's, whether the SKU is contained in a shared capacity or the overflow capacity. In the case of a shared capacity, the user also needs to indicate in which shared capacity the SKU is contained.

After the capacity check has been carried out, the output is presented in the same worksheet (Figure 4.20). The results of the capacity check for the individual capacities are shown behind the table with all the SKU's listed (in the ninth and tenth column). When there is/are pallet(s) remaining or pallet(s) short of a SKU, it is shown in the ninth and tenth column of the corresponding row. The results of the capacity check for the shared capacities and the overflow capacity are given in the two columns behind every capacity (in the first row).

We will now look deeper into the sub CapacityCheck. We can see in the flowchart of workflow 1 in Figure 4.12 that sub CapacityCheck contains the implementation of the buffer inventory capacity check extension. In this role, it checks the replenishment schedule and buffer inventory size and decides whether the buffer inventory capacity is large enough. You can find the code of sub CapacityCheck in Appendix E: VBA code of the inventory control system implementation in Excel. The code is graphically represented in the flowchart of Figure 4.21.

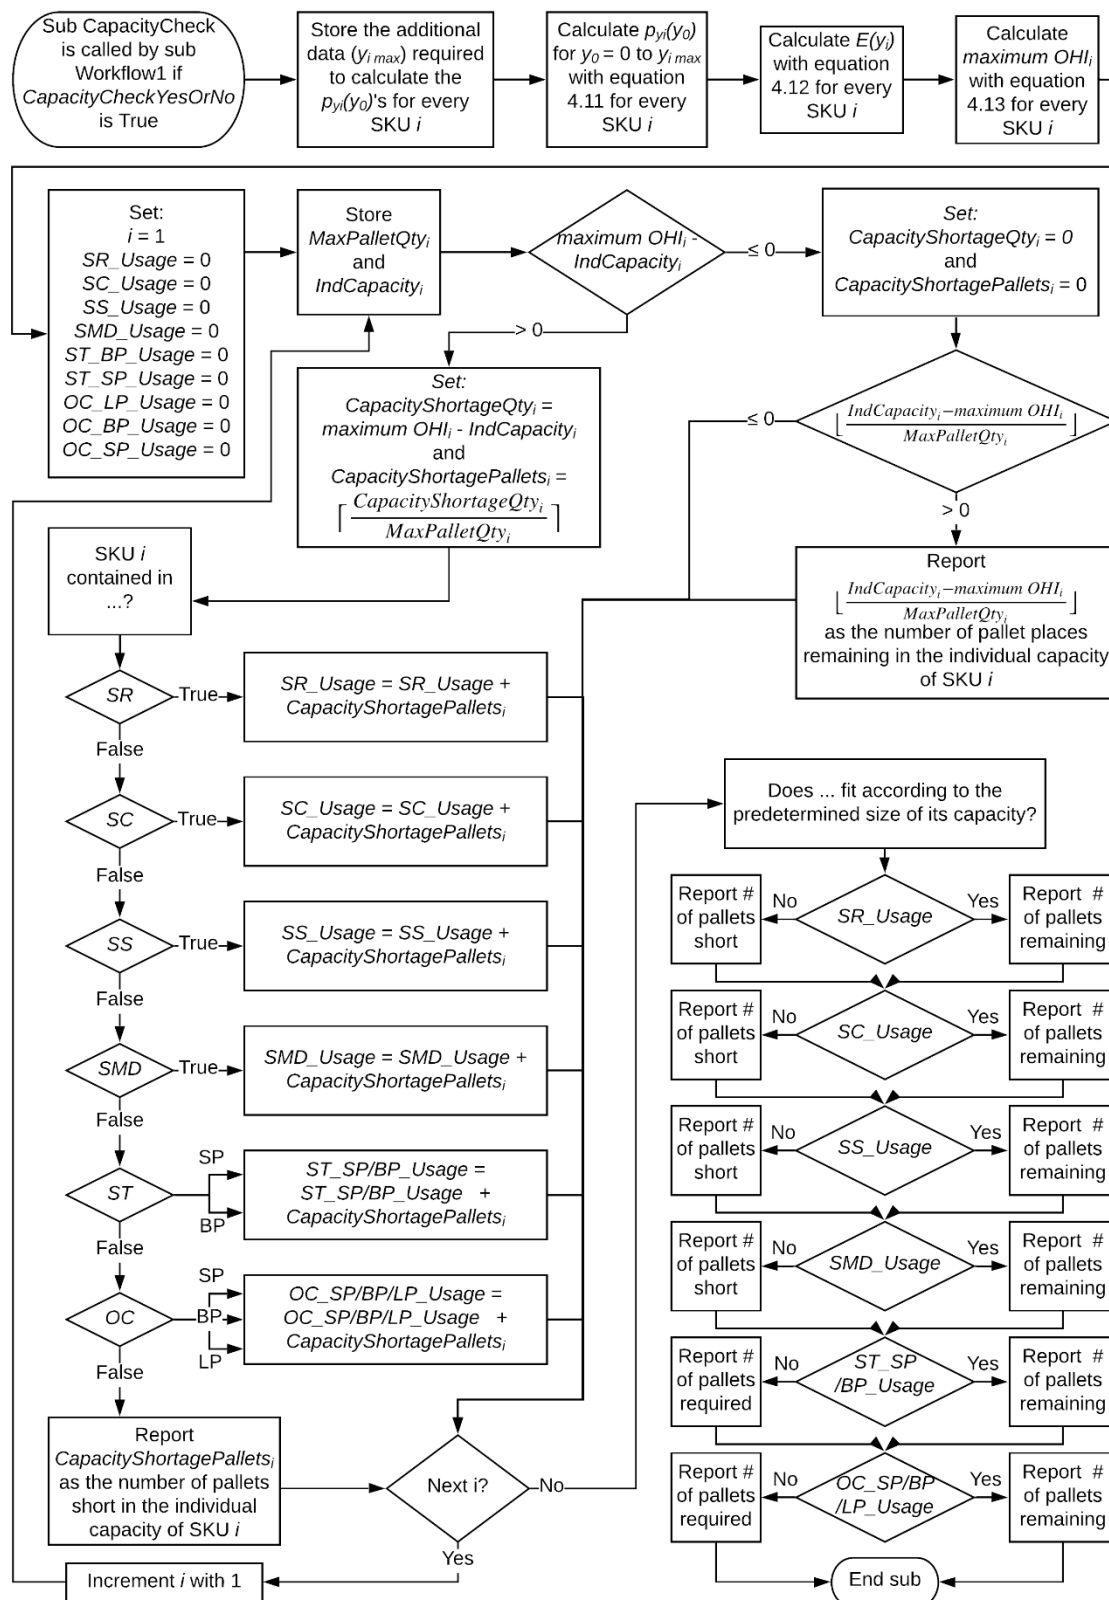


Figure 4.21 The flowchart of sub CapacityCheck

We can see at the start of the flowchart in Figure 4.21 that sub CapacityCheck is called by sub Workflow1 if CapacityCheckYesOrNo is True. The top row of the flowchart consists of calculating

maximum OHI_i for every SKU *i* by using Equation 4.11, 4.12 and 4.13 from the capacity check methodology (Subsection 4.2.4). The rest of the flowchart bears a great resemblance to the flowchart of the capacity checking procedure (Figure 4.7) from methodology Subsection 4.2.4. This is no surprise, since the methodology of the capacity check is tailor-made to the buffer inventory of the A/B series. Therefore, the implementation of the capacity check can follow the methodology closely and so does the flowchart of Figure 4.21.

We want to make a few remarks regarding the implementation of the capacity check. First, we talked about multiple shared capacities in the methodology section of the capacity check. It appears that the buffer inventory can be best divided into five shared capacities: *SR*, *SC*, *SS*, *SMD* and *ST*. This explains why there are six decision diamonds below the rectangle with “SKU *i* contained in ...?” in the flowchart of Figure 4.21, the first five concern the shared capacities and the last one concerns the overflow capacity *OC*. This is in contrast to the two decision diamonds in the flowchart of Figure 4.7, where the first one concerns the shared capacities and the second one the overflow capacity. Likewise, we need to keep track of six capacity usages: *SR_Usage*, *SC_Usage*, *SS_Usage*, *SMD_Usage*, *ST_Usage* and *OC_Usage*. That is why the last step of the flowchart of Figure 4.7 is implemented as six decision diamonds in Figure 4.21, so for every capacity usage it is individually decided whether it is short of pallets or has pallets remaining. The last thing we want to remark is that the usage of shared capacity *ST* and the overflow capacity *OC* are both divided into multiple usages. The reason behind this is that both these capacities can have multiple types of pallets in it. This is in contrast to the other capacities, which can all have one type of pallets in it. Shared capacity *ST* can contain *SP* and *BP* pallets and therefore we calculate *ST_SP_Usage* and *ST_BP_Usage* separately. The overflow capacity *OC* can contain *SP*, *BP* and *LP* pallets and therefore we calculate *OC_SP_Usage*, *OC_BP_Usage* and *OC_LP_Usage* separately. This has also implications for the reporting of the number of pallets short/remaining. Namely, we cannot simply give the number of pallets *short* because there are multiple combinations of *SP/BP/LP* pallets short possible. Instead, we give the number of pallets *required* so the user should determine him/herself what the possible combinations of pallets short are.

After sub CapacityCheck has given all the output of the capacity check (for the individual capacities, shared capacities and overflow capacity), sub CapacityCheck ends. The workflow continues in sub Workflow1. We see in Figure 4.15 that only two steps remain. First, make a message (in the form of a new user form) in which we give some general results of the replenishment schedule and refer to the replenishment schedule in the worksheet (Figure 4.17). This step is done regardless of whether the capacity check is executed or not. The last step is only done if the capacity check is executed. Namely, it is the addition of a sentence to the message in which we refer to the worksheet with the results of the buffer inventory capacity check (Figure 4.20). The resulting message is then shown to the user, see Figure 4.22, and that forms the end of sub Workflow1. This also means that this is the end of workflow 1.

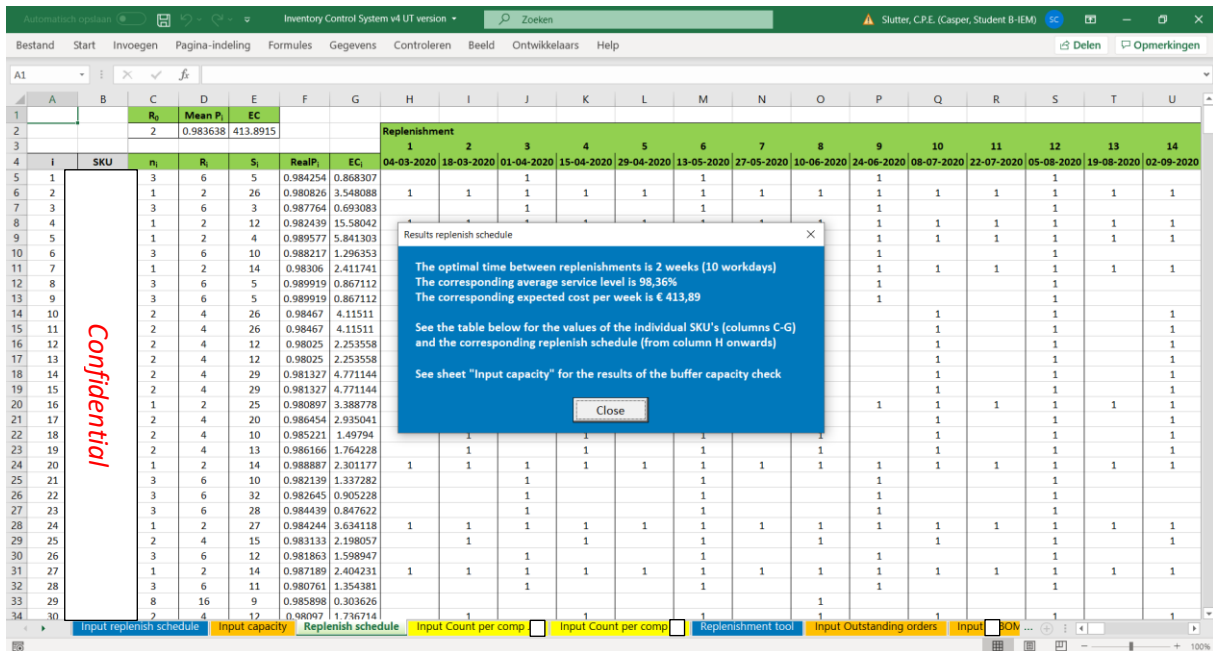


Figure 4.22 The message that the user sees when workflow 1 has ended. It contains general information about the replenishment schedule and it refers to the replenishment schedule and to the results of the capacity check (only when the check is executed).

4.3.3. Implementation of workflow 2

We just described the implementation of the first of two workflows. This workflow is executed only occasionally (only when the demand or costs change significantly or when there is a change in (the number of) critical parts). The result is the *replenishment schedule*. In this schedule you can find which items need to be reviewed on which moments and which order-up-to-level they have. We will now have a closer look at the implementation of the second workflow. This workflow is executed at each review moment (after each base period R_0) and makes use of the *replenishment schedule*, see Figure 4.12. It determines exactly which quantities from which lot(s) should be ordered at a certain review moment (the *replenishment decisions*) to comply with the *replenishment schedule*.

We can see in the flowchart of workflow 2 (Figure 4.12), that the workflow starts with the user form “Make replenish list” and the input of the replenishment date. To this end, worksheet “Replenishment tool” (Figure 4.23) is used.

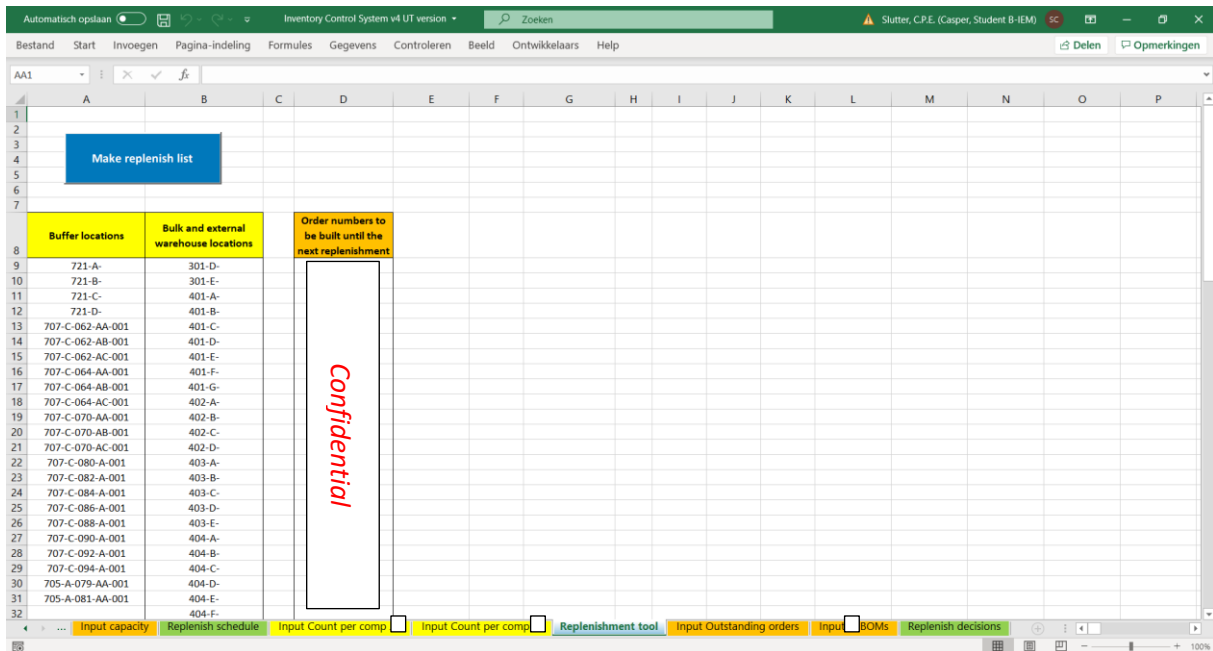


Figure 4.23 Worksheet "Replenishment tool" contains the user form and input for the second workflow.

When the blue button in the worksheet is clicked, the user form shows up (Figure 4.24). The first thing to do in this user form is to select the correct replenishment date from the drop-down list. The drop-down list is programmed in such a way that only relevant replenishment dates are shown: on a replenishment date, only that replenishment date is shown. On a non-replenishment date, both the previous and the first coming replenishment date are shown. This is done because it could be the case that the company is, for example, closed on a replenishment date and wants to do the replenishment in advance or afterwards. In that case, the user can choose the first coming date or the previous date respectively.

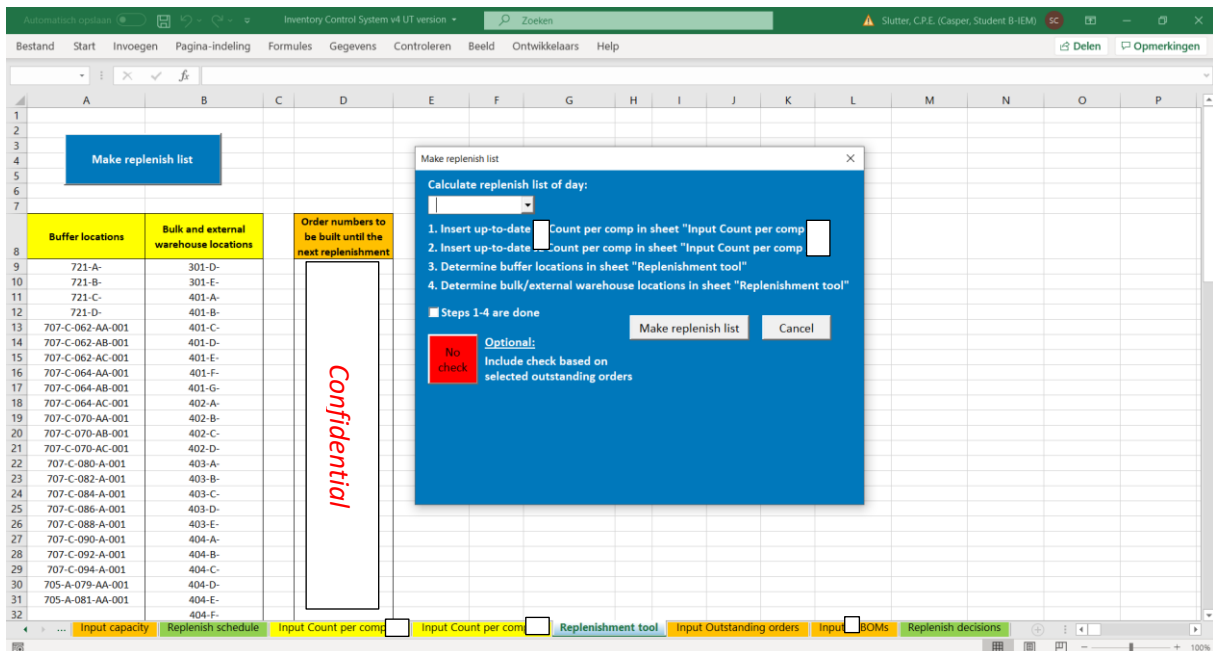


Figure 4.24 The user form of worksheet "Replenishment tool" with which workflow 2 can be controlled.

Furthermore, the user form contains the instructions to insert the required input for the second workflow. This input should be inserted in three places. First, an up-to-date count per comp file for C

parts need to be inserted in a new worksheet, called “Input Count per comp C” (Figure 4.25). The count per comp file is the inventory file of the company, which shows the available lots (including locations and quantities) of every SKU. Secondly, there is another count per comp file, for the L parts, of which an up-to-date version needs to be inserted in another new worksheet, called “Input Count per comp L” (Figure 4.26). Both count per comp files have exactly the same format, so both worksheets look very similar. Thirdly, the user should determine for the locations from the count per comp files, whether they belong to the buffer inventory or to the bulk and external warehouse inventory. The locations of the count per comp files should therefore be split up between the first two columns in the worksheet “Replenishment tool” (Figure 4.23).

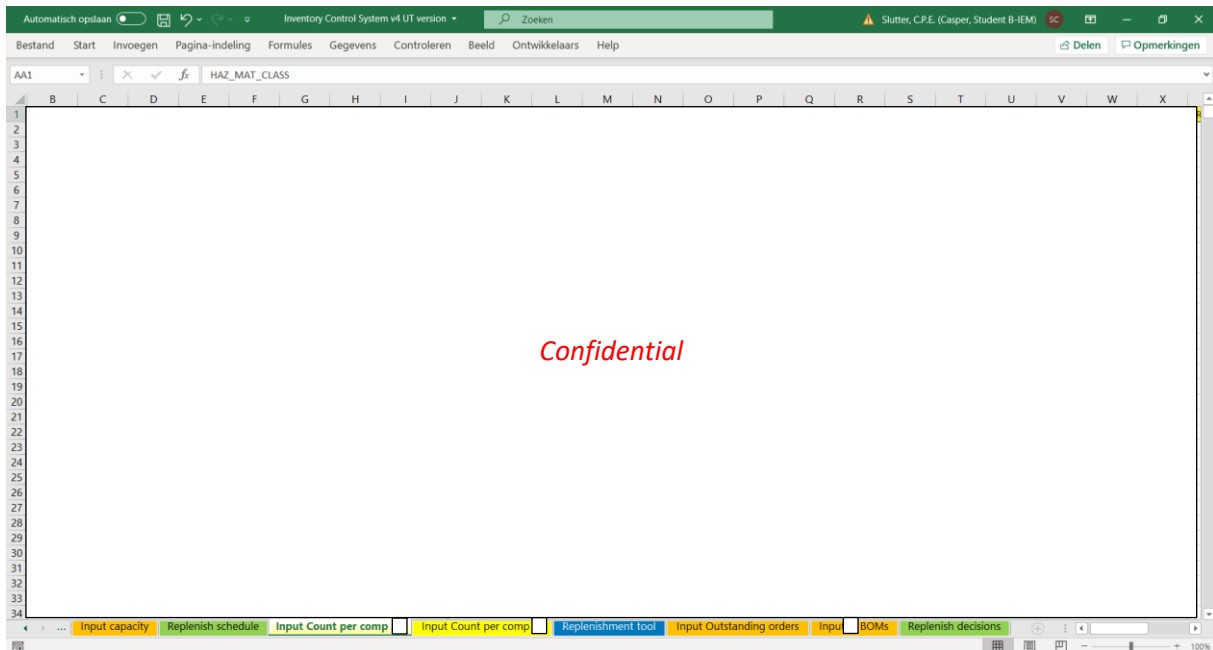


Figure 4.25 Worksheet “Input Count per comp C” contains the available lots of every C SKU.




Figure 4.26 Worksheet “Input Count per comp L” contains the available lots of every L SKU.

Workflow 2 can now be executed (we come back to the optional outstanding orders check later). To this end, the button “Make replenish list” in the user form (Figure 4.24) should be clicked. The user form then calls sub Workflow2, as can also be seen in the flowchart of workflow 2 in Figure 4.12. When calling sub Workflow2, the user form passes on the date for *ReplenishDate* and a True/False for *OutstandingOrderCheckYesOrNo*, which concerns the outstanding orders check (False for now).

Sub Workflow2 is the main sub of workflow 2. We can see in the flowchart of workflow 2 in Figure 4.12 that sub Workflow2 contains the replenishment decisions extension. In this role, it is responsible for making the replenishment decisions. Furthermore, it should call the sub OutstandingOrdersCheck (or not). You can find the code of sub Workflow2 in Appendix E: VBA code of the inventory control system implementation in Excel. The code is graphically represented in the flowchart of Figure 4.27.

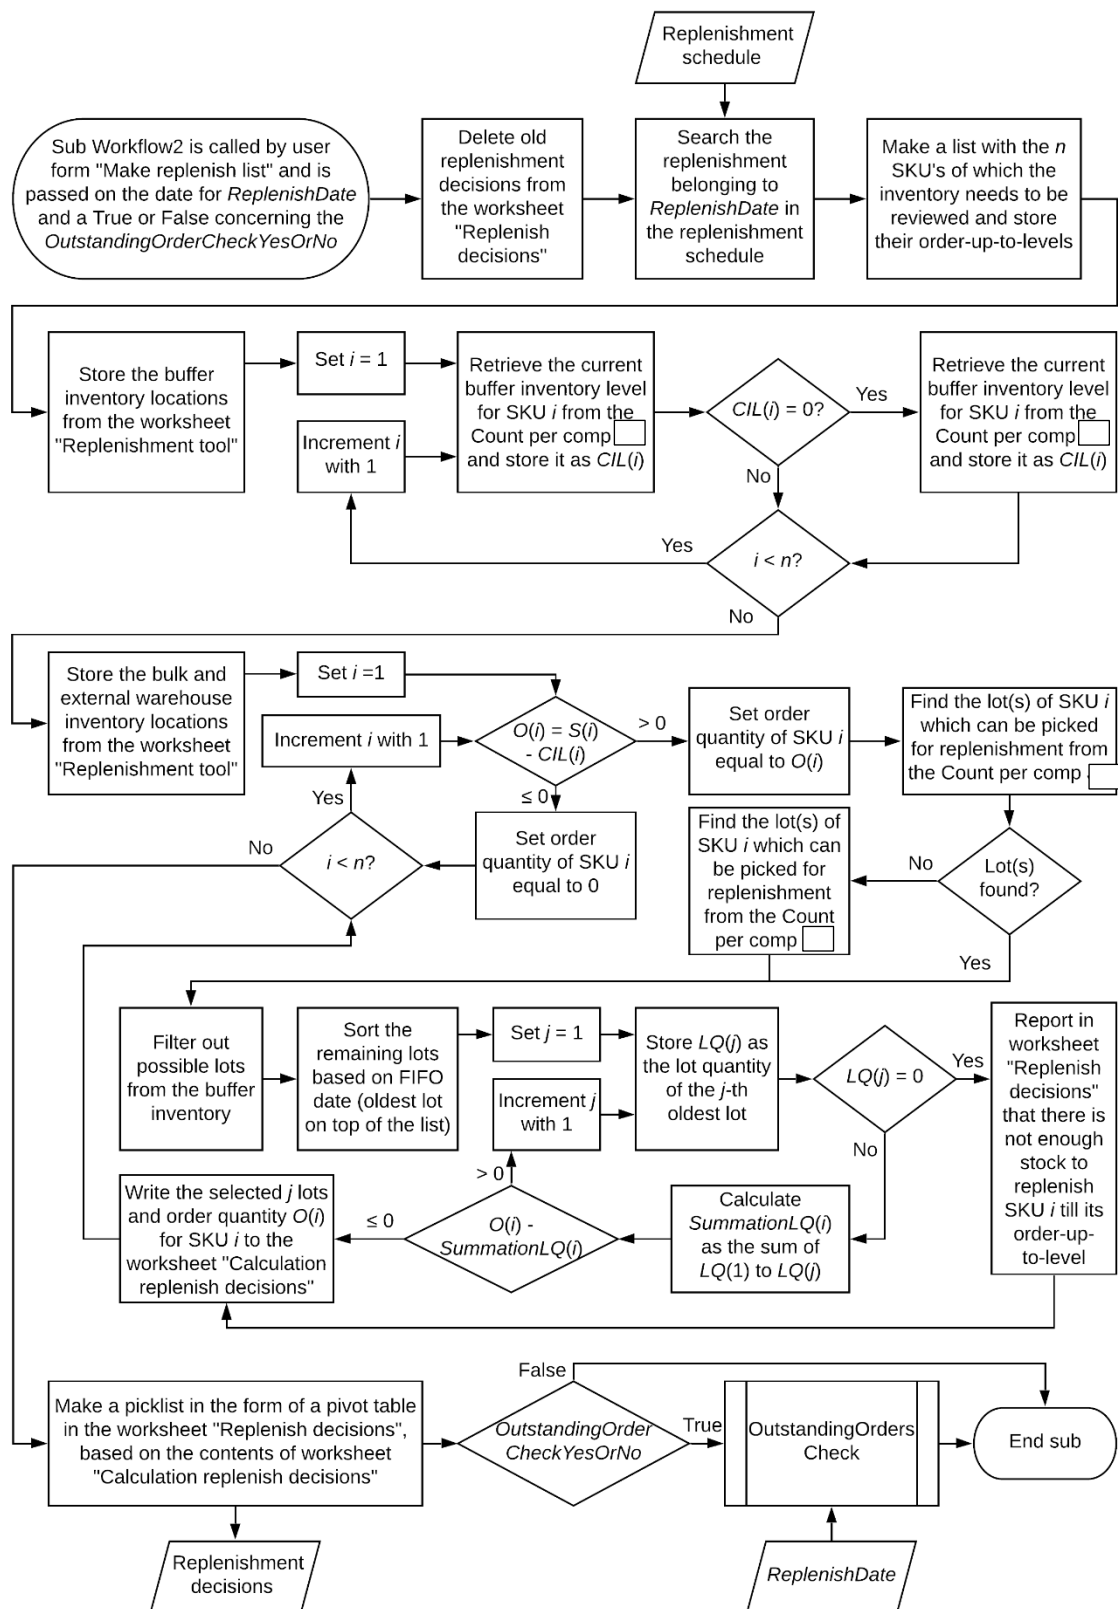


Figure 4.27 The flowchart of sub Workflow2

The first thing that sub Workflow2 does, is deleting the old replenishment decisions. Subsequently, it uses the *ReplenishDate* to find information about the new replenishment from the replenishment

schedule. Therefore, the replenishment schedule (output of the first workflow) is an input in this step of sub Workflow2. Using this information, the sub makes a list with the subset of n SKU's which need to be reviewed with the new replenishment and stores their order-up-to-levels.

We have seen in methodology Subsection 4.2.3 about the replenishment decisions extension that we need the (required) order quantities of these n SKU's as part of the replenishment decisions. We have also seen in this subsection that these order quantities can be calculated as the difference between the order-up-to-level and the current buffer inventory level. Therefore, sub Workflow2 will first calculate the current buffer inventory levels of the n SKU's of the subset. To this end, the buffer inventory locations are stored, after which the first loop of the sub starts. This loop retrieves the current buffer inventory levels from the count per comp C file. If the current buffer inventory level is zero, it might be a L SKU, so the current buffer inventory level is in that case retrieved from the count per comp L file.

Now that the current buffer inventory levels are determined, the sub can determine the (required) order quantities. In addition, the sub advises the lots to satisfy these order quantities. We have seen in Subsection 4.2.3 that this is the second part of the replenishment decisions. Since these lots should come from the bulk inventory or external warehouse inventory, the sub first stores the bulk and external warehouse inventory locations. Subsequently, a new loop starts. The first step of this loop is to determine the order quantity of a SKU by subtracting the current buffer inventory level from the order-up-to-level. If this quantity is less than or equal to zero, then the order quantity is set equal to zero since no replenishment is required for this SKU. The loop then starts again with the next SKU. If the quantity is greater than zero, then the order quantity is set equal to this quantity and the loop will proceed with the advice on the lot(s) to choose. In Subsection 4.2.3, we determined that this advice should be based on the FIFO method. Therefore, the sub will make a list of the lots which can be picked for replenishment and will sort this list from oldest to newest FIFO date. To this end, first the lots which can be picked for replenishment are searched in the count per comp C file. If no lots are found, lots are searched in the count per comp L file, since it is probably a L SKU then. Subsequently, possible lots from the buffer inventory are filtered out the list. The remaining lots are then sorted based on their FIFO dates, so that the lots in the list have an ascending FIFO date.

Next, a new loop starts within the existing loop. The goal of this inner loop is to select the minimum number of oldest lots that cover the entire order quantity, as described in Subsection 4.2.3. This Do ... Until loop works as follows: first, the lot quantity of the oldest lot $LQ(1)$ is stored. If this quantity is smaller than the order quantity, then the lot quantity of the second oldest lot $LQ(2)$ is stored. If the sum of both lot quantities is still smaller than the order quantity, then the loop also stores and adds $LQ(3)$ to the sum. This is done until the sum of the stored lot quantities is greater than or equal to the order quantity, so that the entire order quantity is covered. It can also be the case that there are no lots or not enough lots to cover the entire order quantity. If there are no lots at all, $LQ(1)$ will be equal to zero and the sub will immediately step out of the loop. If there are some lots, but the sum of their lot quantities is not sufficient, the loop will search for the next $LQ(j)$. However, this $LQ(j)$ will be equal to zero since there is no next lot. The sub will then step out of the loop to prevent an infinite loop. In both cases, $LQ(j)$ will be equal to zero. We can see in the flowchart of Figure 4.27 that when this is the case, it is reported in the replenishment decisions that there is not enough stock to replenish that SKU. Besides, the quantity that is short of the SKU is reported. In fact, this forms the third part of the replenishment decisions.

After the (inner) loop ends, regardless of whether the entire order quantity is covered or not, the outer loop continues with writing the selected lots of the SKU to the new worksheet "Calculation replenish decisions", together with the order quantity of that SKU. This worksheet is only used to

store data during the calculation of the replenishment decisions. Therefore, this worksheet contains no input or output of the workflow. For this reason, the worksheet is hidden from the user and is thus not a part of the user interface. Hence, we will not show the worksheet here.

The loop writes the required order quantities and the advised lots for every of the n SKU's one by one to the worksheet "Calculation replenish decisions". This worksheet now contains a list with all the data of the replenishment decisions, but not in a user-friendly way. Therefore, the step after the loop makes a picklist of the replenishment decisions. This picklist is in the form of a pivot table, which is based on the contents of the list in the worksheet "Calculation replenish decisions". This pivot table is given in a new worksheet, called "Replenish decisions" (Figure 4.28). You can see the pivot table in columns A till E. In the first column, you can see the SKU's which need to be replenished. The second column contains the (required) order quantities of these SKU's, the first part of the replenishment decisions. Columns C till E contain the second part of the replenishment decisions: the advice on the lots to satisfy the order quantity of the second column. We have seen in methodology Subsection 4.2.3 that this advice consists of the lot date, lot location and lot quantity of all the advised lots. Therefore, these three lot characteristics are given in this order in columns C till E. We also saw in Subsection 4.2.3 that the advice is based on the FIFO method, so we sort the pivot table in such a way that the advised lots of each SKU have an ascending lot date. In this way, the oldest lots are probably picked first when the user consults the replenishment decisions.

| SKU | Required Order Qty | Advised Lot Date(s) | Advised Lot Location(s) | Each Qty on Location | SKU | Qty short in bulk/ external warehouses |
|-----|--------------------|---------------------|-------------------------|----------------------|-----|--|
| | 20 | 31-12-2019 | 403-A-011-C-0 | 20 | | 25 |
| | 11 | 30-06-2019 | 402-D-048 | 4 | | 25 |
| | | | 707-B-111-A-001 | 1 | | 15 |
| | | | GAC-C-008 | 6 | | 14 |
| | 9 | 28-02-2019 | 705-B-056-C-0 | 16 | | 11 |
| | | | | | | 5 |
| | 26 | 31-12-2017 | 722-F-048-B-0 | 1 | | 19 |
| | 26 | 30-04-2019 | 722-F-104-C-0 | 1 | | |
| | 12 | 31-05-2018 | 722-F-031-H-0 | 1 | | |
| | | | 722-F-032-E-0 | 1 | | |
| | | | 722-F-070-G-0 | 1 | | |
| | | | 722-F-076-C-0 | 1 | | |
| | | | 722-F-098-C-0 | 1 | | |
| | | | 722-F-116-C-0 | 1 | | |
| | | | 722-F-133-G-0 | 1 | | |
| | | 30-06-2018 | 722-F-034-C-0 | 1 | | |
| | | | 722-F-061-G-0 | 1 | | |
| | | | 722-F-076-F-0 | 1 | | |
| | | | 722-F-082-F-0 | 1 | | |
| | | | 722-F-084-C-0 | 1 | | |
| | 12 | 31-05-2018 | 722-F-018-B-0 | 1 | | |
| | | | 722-F-019-G-0 | 1 | | |
| | | | 722-F-050-F-0 | 1 | | |
| | | | 722-F-062-C-0 | 1 | | |
| | | | 722-F-066-D-0 | 1 | | |
| | | | 722-F-100-B-0 | 1 | | |
| | | 30-06-2018 | 722-F-040-C-0 | 1 | | |

Figure 4.28 Worksheet "Replenish decisions" contains the replenishment decisions, the main output of workflow 2.

The third part of the replenishment decisions is the quantity short for the SKU's for which the advice cannot satisfy the order quantity. We saw in Figure 4.27 that these quantities are directly reported to the worksheet "Replenish decisions" during the inner loop of sub Workflow2. Therefore, the quantities short are not given in the picklist (pivot table) in this worksheet, but in a separate table to the right of it, see Figure 4.28. This table contains the concerning SKU's together with the quantities that are short of these SKU's after the advice has been given.

In conclusion, the pivot table and the table with the quantities short together form the replenishment decisions. The user should stick to the required order quantities from the pivot table and is advised to satisfy them by ordering the advised lots from the pivot table. However, as we can

note in Figure 4.28, the sum of the lot quantities of the advised lots is not equal to the required order quantity for every SKU. There are namely three possibilities:

- The sum of the advised lot quantities is equal to the order quantity. This means that a certain number of oldest lots exactly satisfy the required order quantity. The advice is to order exactly these lots (completely) .
- The sum of the advised lot quantities is greater than the order quantity. This means that a certain number of oldest lots together contain a larger quantity than required, whereas with one lot less, the required order quantity would not be satisfied completely. The advice is to either order the lot with the newest lot date (bottom of the list) completely or order only the part of it that is really needed. In other words, order too much or split a lot in two and order exactly enough, respectively. This consideration depends very much on the quantity that will be ordered too much, but it is also very SKU-dependent. Therefore, this choice is left to the company.
- The sum of the advised lot quantities is less than the order quantity. This means that even if all lots from the bulk and external warehouse inventories are ordered, the required order quantity is still not satisfied. In this case, the difference between the sum of the advised lot quantities and the order quantity is given in the table on the right as quantity short. The advice is to order the lots that are advised in any case and try to order the quantity short in another way.

If there is no lot in the buffer/external warehouse inventory available at all, no advice can be given of course. Therefore, a SKU where this applies to, is not given in the picklist at all, it is only given in the table on the right. Here, the quantity short of that SKU will be equal to the complete (required) order quantity. The complete order quantity should therefore be ordered in another way.

We see in Figure 4.27 that sub Workflow2 continues after it has outputted the replenishment decisions. The remainder of sub Workflow2 has to do with the outstanding orders check extension. When we look back at the flowchart of workflow 2 in Figure 4.12, we see that this is indeed the last step of workflow 2. Furthermore, we see in Figure 4.12 that this outstanding orders check is made optional. It is namely not necessary to do the check since it does not affect the replenishment decisions. However, the check will predict whether stockouts will occur when the replenishment decisions are followed and is therefore a useful addition to the replenishment decisions.

We now go back to the user form of Figure 4.24 again. Just as with the capacity check in the first user form, the user can choose in this second user form with a toggle button whether the outstanding orders check should be executed or not. The default setting is “No outstanding orders check” (red toggle button). If the user wants to include the outstanding orders check, (s)he should click the (currently) red toggle button. The toggle button then turns green, indicating that the outstanding orders check will be executed (Figure 4.29).

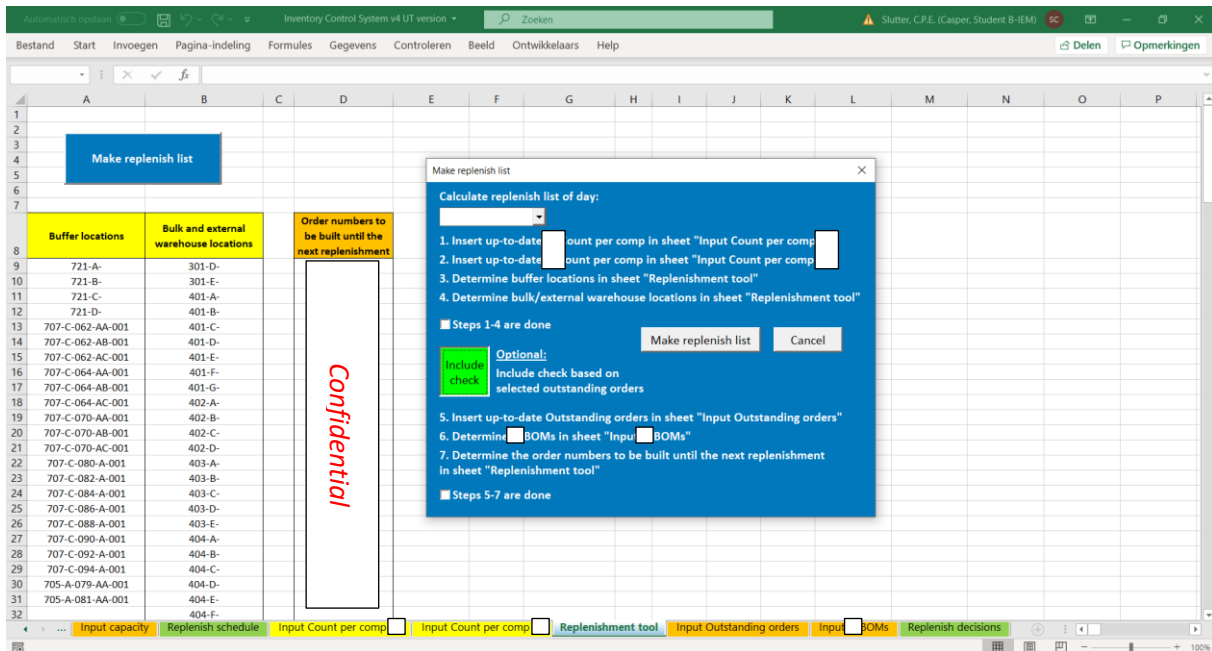


Figure 4.29 The user form of worksheet "Replenishment tool" with the outstanding orders check selected.

The toggle button assigns either True (include outstanding orders check) or False (no outstanding orders check) to the variable *OutstandingOrderCheckYesOrNo*. The user form passes on the value of *OutstandingOrderCheckYesOrNo* when it calls sub *Workflow2*. We can see in Figure 4.27 that sub *Workflow2* checks the value of *OutstandingOrderCheckYesOrNo* to decide if the outstanding orders check should be carried out.

If the outstanding orders check should not be executed (*OutstandingOrderCheckYesOrNo* = False), sub *Workflow2* and thus workflow 2 is ended. However, if the outstanding orders check should be executed (*OutstandingOrderCheckYesOrNo* = True), sub *Workflow2* calls a new sub, sub *OutstandingOrdersCheck*. This sub contains the complete outstanding orders check extension. It is implemented in a separate sub for the same reasons as with the capacity check extension. Namely, the code of the extension is then separated from the rest of the code (the replenishment decisions extension). Furthermore, the extension is optional, so with a separate sub it is easy to either call it or not.

Before we discuss the sub *OutstandingOrdersCheck*, we will first present two new worksheets. These worksheets are only used for the outstanding orders check extension. Both worksheets contain input that is required by the outstanding orders check. When the outstanding orders check is not activated, nothing is done with this input so these worksheets should not be considered in that case.

The first new worksheet is called "Input Outstanding orders" (Figure 4.30). The outstanding orders check will follow the methodology of Subsection 4.2.5. We know from this subsection that the check makes use of the outstanding orders file, an Excel file which can be extracted from the WMS of the company. An up-to-date version of this file should be inserted in this new worksheet. From this list with outstanding orders, we only need the outstanding orders that are planned for the new base period according to the methodology. Therefore, the user should determine these orders and insert the order numbers of the concerning orders in the fourth column of worksheet "Replenishment tool" (Figure 4.23)

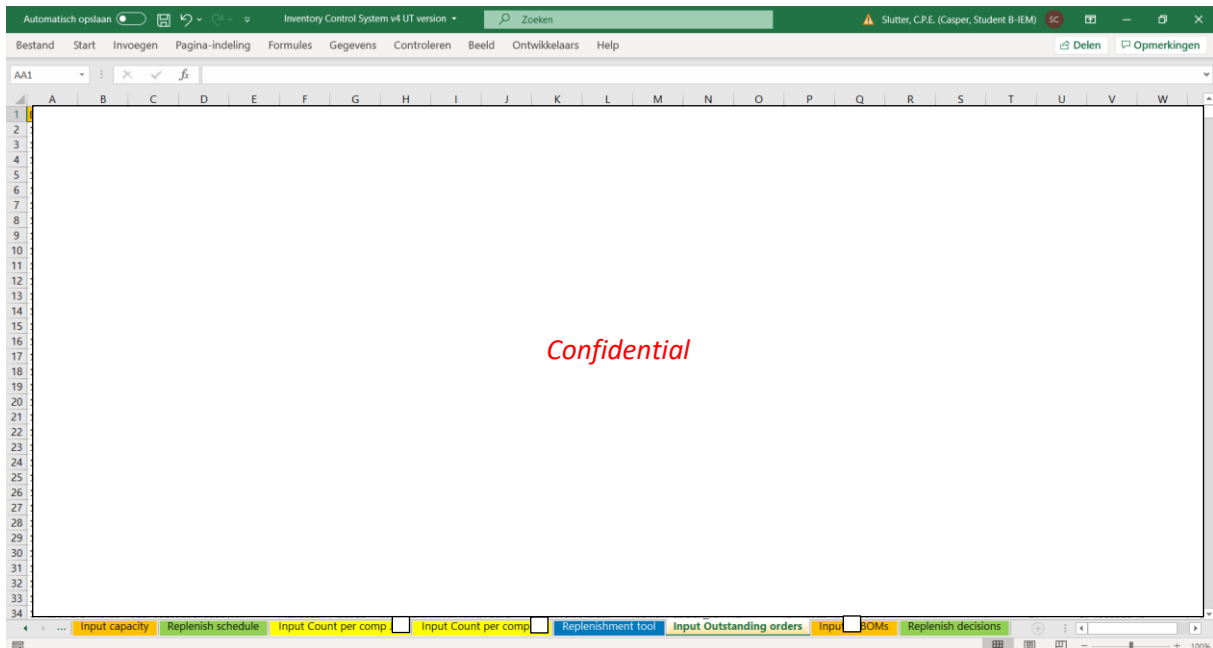


Figure 4.30 Worksheet “Input Outstanding orders” contains the outstanding orders, which is part of the input for the optional outstanding orders check extension.

The outstanding orders file only contains the requirements for the C SKU’s directly, while extra information is needed to derive the requirements for the L SKU’s from this file. This extra information consists of the Bill of Materials (BOM) for every type of the A/B series (every CD number). With the BOM’s, we can calculate the requirements of the L SKU’s based on the requirements of the CD numbers, which in turn can be found in the outstanding orders file. In conclusion, we only need the part with the L SKU’s from the BOM’s of every CD number to find the requirements for the L SKU’s. These BOM’s should be inserted in a new worksheet, called “Input L BOMs” (Figure 4.31). When we look at Figure 4.31, we see that the BOM’s should be inserted in a special way: in the form of a table. All CD numbers of the A/B series that are taken into account by the inventory control system should be inserted in the first row of the table. All the L SKU’s that are controlled by the inventory control system should be listed in the first column. Subsequently, the quantities of the controlled L SKU’s for each CD number should be inserted in the corresponding cell of the table. When the cell is left blank, it means that the SKU is not contained in the type belonging to that CD number.

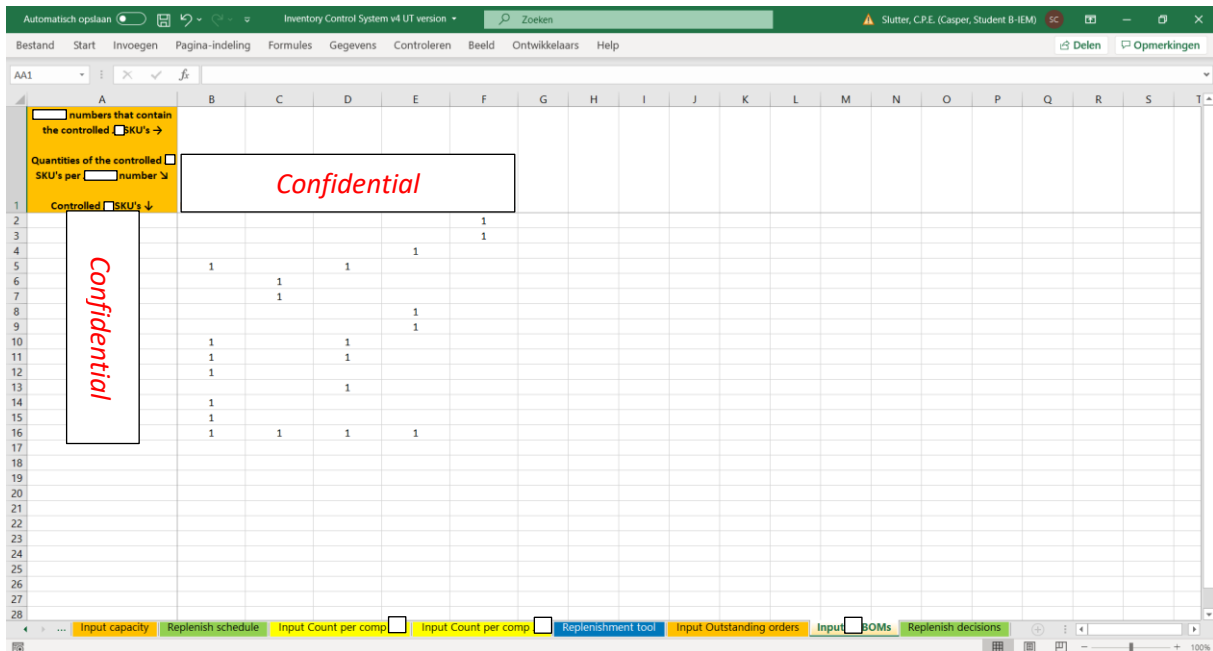


Figure 4.31 Worksheet "Input L BOMs" contains the BOM's of every CD number (only for the controlled L SKU's), which are part of the input for the optional outstanding orders check extension.

We will now discuss the sub OutstandingOrdersCheck in greater detail. We can see in the flowchart of workflow 2 in Figure 4.12 that sub OutstandingOrdersCheck contains the implementation of the outstanding orders check extension. In this role, it checks the replenishment decisions and makes a prediction of the future demand based on selected outstanding orders. Then, it calculates whether the replenishment decisions are sufficient to prevent stockouts. You can find the code of sub OutstandingOrdersCheck in Appendix E: VBA code of the inventory control system implementation in Excel. The code is graphically represented in the flowchart of Figure 4.32.

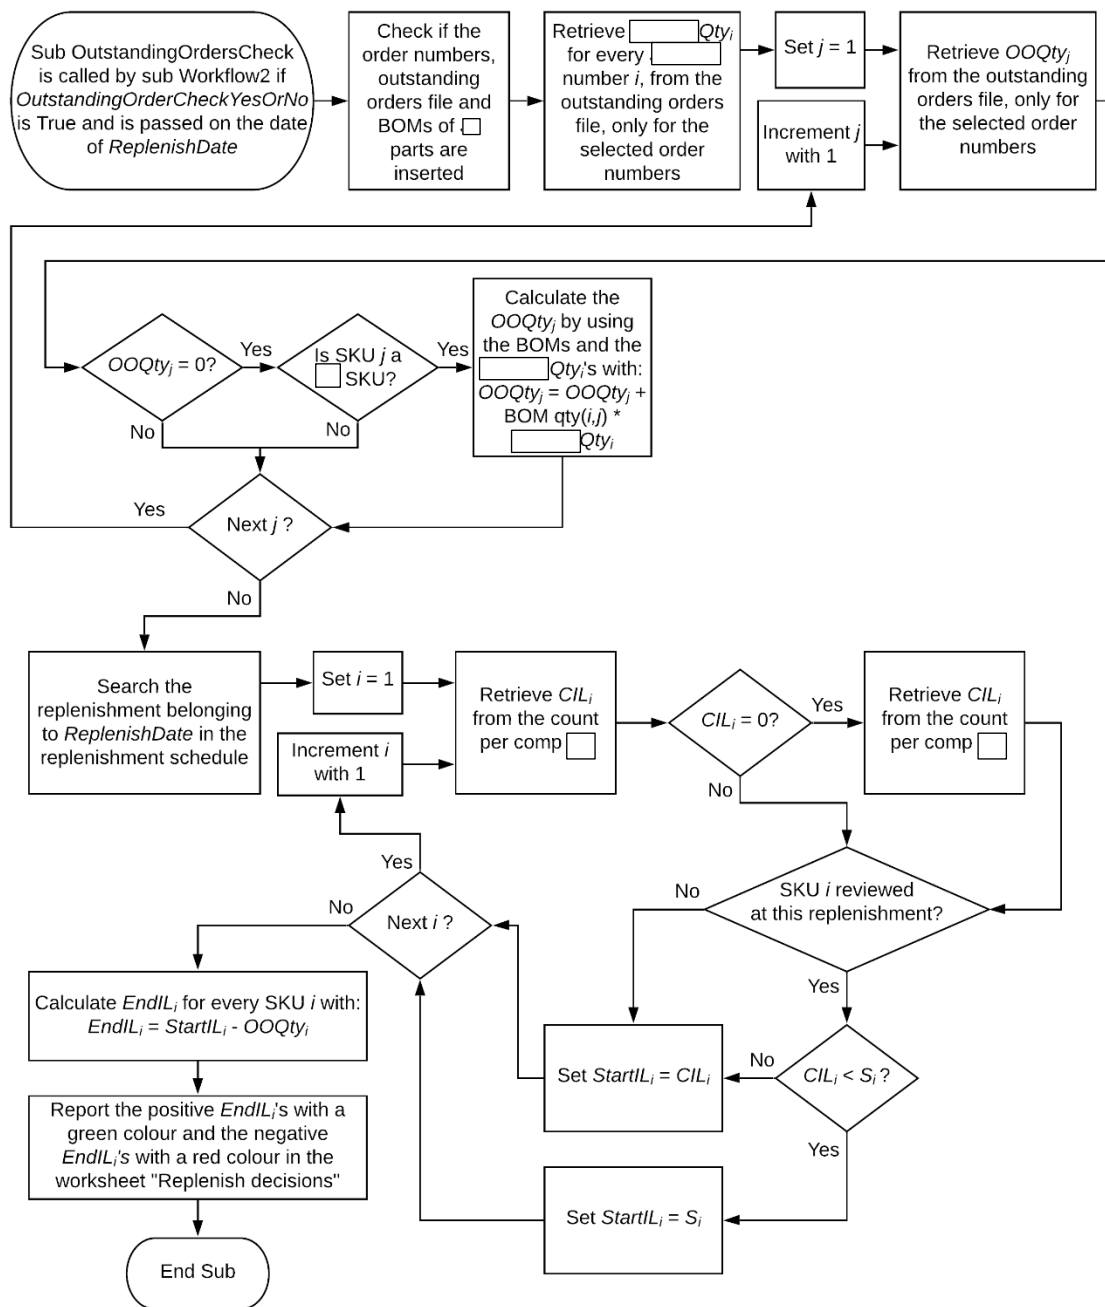


Figure 4.32 The flowchart of sub OutstandingOrdersCheck

We can see in the flowchart of Figure 4.32 that sub OutstandingOrdersCheck starts with checking if all required information for the outstanding orders check is inserted. Subsequently, the $CDQty_i$'s are retrieved from the selected orders of the outstanding orders file for every CD number i . These quantities are used in the next loop. The rest of the flowchart (Figure 4.32) bears a great resemblance to the flowchart of the outstanding orders checking procedure (Figure 4.11) from methodology Subsection 4.2.5. This is logical since the methodology of the outstanding orders check is tailor-made to the way that production happens at the company. Therefore, the implementation of the outstanding orders check can follow the methodology closely and so does the flowchart of Figure 4.32.

We want to comment on the difference between the outstanding orders checking procedure from the methodology (Figure 4.11) and the implementation (Figure 4.32). We see in the methodology that the first step is to retrieve $OOQTY_i$ for every SKU i from the outstanding orders file, only for the selected order numbers. In the implementation, it turns out to be more extensive: it consists of a loop with multiple steps and decisions. The reason for this is that we can only find the $OOQTY_i$ for the C SKU's directly via the outstanding orders file. Therefore, the loop first tries to find $OOQTY_i$ in a direct way, by retrieving them from the outstanding orders file. If $OOQTY_i$ appears to be zero, then SKU i is possibly a L SKU, so the sub checks if the SKU is a L SKU. If so, it calculates the $OOQTY_i$ by looking at the BOMs of worksheet "Input L BOMs" and the $CDQty_i$'s which were retrieved before the loop. This works as follows: if CD number i contains a quantity x of L SKU j according to its BOM, then $OOQTY_j$ is equal to x times $CDQty_i$. If L SKU j is contained in more CD numbers, then the calculation of $OOQTY_j$ consists of the addition of multiple terms of the form $BOM\ qty(i,j) * CDQty_i$. After the loop has determined the $OOQTY_i$'s of every SKU, the sub continues. The next step is, just as in the methodology, to search the replenishment belonging to the $ReplenishDate$ in the replenishment schedule. Subsequently, CIL_i is retrieved for every SKU i . This step is again more extensive in the implementation. Namely, first CIL_i is retrieved from the count per comp C, but when CIL_i appears to be zero, then CIL_i is retrieved from the count per comp L since it might be a L SKU.

You might notice that we already retrieved the current buffer inventory level CIL_i (in the same way) in sub Workflow2 and might wonder why we do the exact same thing again in sub OutstandingOrdersCheck. The reason is that in sub Workflow2, we needed to retrieve CIL_i only for the subset of n SKU's which required inventory review. However, this time we need to retrieve CIL_i for all m SKU's of the inventory control system, since the outstanding orders check is executed for all SKU's and not only the subset of n SKU's. It was chosen not to reuse the CIL_i values of the n SKU's of the subset to prevent confusion and mistakes. Instead, we retrieve CIL_i again for all m SKU's at the same time.

The next few steps of the outstanding orders check implementation are exactly the same as the outstanding orders checking procedure of Figure 4.11 from methodology. Namely, there is a loop where $StartIL_i$ is set for each SKU i , one by one, based on two decisions. The first decision is if SKU i is reviewed at this replenishment and the second decision is if SKU i is replenished at this replenishment (only if $CIL_i < S_i$). When $StartIL_i$ has been determined for every SKU i , the loop ends and the sub continues. The last step of the outstanding orders checking procedure (Figure 4.11) is executed now: the calculation of $EndIL_i$ for every SKU i by using the values for $StartIL_i$ and $OOQty_i$ as determined before.

The last step of sub OutstandingOrdersCheck is to report the outcome of the outstanding orders check: the $EndIL_i$ levels for each SKU i . These levels form the prediction from which the user can see if stockouts for some SKU's are expected. The $EndIL_i$ levels are given in a separate table in the worksheet "Replenish decisions". This table is only given if the check has been activated. You can find the table in that case in the tenth and eleventh column of the worksheet, see Figure 4.33. Positive $EndIL_i$ levels are given in green, since they indicate that the expected inventory just before the next replenishment arrives, will be greater than or equal to zero. In other words, no stockout is expected in the next base period for SKU's with positive $EndIL_i$ levels. Negative $EndIL_i$ levels are given in red, where the amount that the level is negative stands for the amount that is expected to be short in the next base period.

| SKU | Required Order Qty | Advised Lot Date(s) | Advised Lot Location(s) | Each Qty on Location | SKU | Qty short in bulk/external warehouses | SKU | Expected inventory just before the next replenishment arrives |
|-----|--------------------|---------------------|-------------------------|----------------------|-----|---------------------------------------|-----|---|
| | 20 | 31-12-2019 | 403-A-011-C-0 | 20 | | 25 | | 41 |
| | 11 | 30-06-2019 | 402-D-048 | 4 | | 25 | | 26 |
| | | | 707-B-111-A-001 | 1 | | 15 | | 9 |
| | | | GAC-C-008 | 6 | | 15 | | 12 |
| | | | | | | 14 | | 4 |
| | 9 | 28-02-2019 | 705-B-056-C-0 | 16 | | 11 | | 8 |
| | | | | | | 5 | | 11 |
| | | | | | | 10 | | 0 |
| | 26 | 31-12-2017 | 722-F-048-B-0 | 1 | | | | 26 |
| | | | | | | | | 26 |
| | 26 | 30-04-2019 | 722-F-104-C-0 | 1 | | | | 12 |
| | | | | | | | | 12 |
| | 12 | 31-05-2018 | 722-F-031-H-0 | 1 | | | | 26 |
| | | | 722-F-032-E-0 | 1 | | | | 26 |
| | | | 722-F-070-G-0 | 1 | | | | 26 |
| | | | 722-F-076-C-0 | 1 | | | | 25 |
| | | | 722-F-098-C-0 | 1 | | | | 7 |
| | | | 722-F-116-C-0 | 1 | | | | 10 |
| | | | 722-F-133-G-0 | 1 | | | | 20 |
| | | 30-06-2018 | 722-F-034-C-0 | 1 | | | | 42 |
| | | | 722-F-061-G-0 | 1 | | | | 22 |
| | | | 722-F-076-F-0 | 1 | | | | 22 |
| | | | 722-F-082-F-0 | 1 | | | | 27 |
| | | | 722-F-084-C-0 | 1 | | | | 15 |
| | 12 | 31-05-2018 | 722-F-018-B-0 | 1 | | | | 24 |
| | | | 722-F-019-G-0 | 1 | | | | 14 |
| | | | 722-F-050-F-0 | 1 | | | | 21 |
| | | | 722-F-062-C-0 | 1 | | | | 19 |
| | | | 722-F-066-D-0 | 1 | | | | 11 |
| | | | 722-F-100-B-0 | 1 | | | | 10 |
| | | 30-06-2018 | 722-F-040-C-0 | 1 | | | | 14 |

Figure 4.33 The output of the outstanding orders check is given in column J and K of worksheet "Replenish decisions", but only if the check is activated.

After sub OutstandingOrdersCheck has reported all $EndIL_i$ levels in worksheet "Replenish decisions", sub OutstandingOrdersCheck ends. The workflow continues in sub Workflow2. We see in Figure 4.27 that after sub OutstandingOrdersCheck ends, sub Workflow2 also immediately ends. This means that this is the end of workflow 2.

4.3.4. Conclusion on the implementation of the inventory control system

This section described the implementation of the inventory control system in Excel. The user interface of the inventory control system is given by the Excel worksheets and user forms. The model behind the inventory control system is programmed in VBA, the programming language of Excel. The implementation of the model behind the inventory control system in VBA consists of two workflows: workflow 1 and workflow 2. Both workflows have their own user form and worksheets as user interface and both workflows have their own subs where the model is programmed in. Since the model is based on the methodology of Section 4.2, the subs contain implemented versions of the parts of the inventory control system as described in the methodology section. Workflow 1 consists of 3 subs which contain the demand model, inventory model and the buffer inventory capacity check extension, the latter has been made optional. Workflow 2 consists of 2 subs which contain the extension which calculates the replenishment decisions and the outstanding orders check extension, the latter has been made optional.

We have seen in Subsection 4.3.1 that the two workflows have a relation between them. The goal of the first workflow is to calculate the replenishment schedule. This is a procedure that should, in principle, be executed only rarely, since the replenishment schedule is made for the long term. The second workflow, on the other hand, should be executed often, since its goal is to find the replenishment decisions at a replenishment date. In order to find these replenishment decisions, the second workflow needs the replenishment schedule of the first workflow. The replenishment schedule is therefore the link between both workflows, it is an output of the first workflow and an input for the second workflow.

In Subsection 4.3.2, we discussed the implementation of the first workflow. We showed the worksheets that require the input for workflow 1 and the worksheets that give the output of

workflow 1. Besides, we showed the user form that controls workflow 1. Furthermore, we gave flowcharts that graphically depict the code of the subs of workflow 1. An explanation of these flowcharts and the relations with the methodology on which the corresponding code is based, was also given in this subsection. In Subsection 4.3.3, we discussed the implementation of the second workflow in the same way. The complete VBA code of all subs and user forms can be found in Appendix E: VBA code of the inventory control system implementation in Excel.

In this section, we focused on how the inventory control system has been implemented. However, we did not describe in great detail how the implemented inventory control system should be used. In other words, we showed the user interface (worksheets and user forms) but did not describe how the company should use this user interface. This has been explained in a manual specifically intended for the company. This manual is a separate document, included in this report in Appendix D: Manual of the inventory control system.

4.4. Determination of values for inventory control system parameters

We have seen in the previous sections that the inventory control system requires several parameters as input in order to calculate the replenishment decisions. Therefore, values for these parameters need to be found. The quality of the input directly affects the quality of the output. It is thus very important to find representative values for the parameters to get representative values for the expected costs and service level from the system. The parameters can be subdivided in demand parameters, service levels, time parameters and cost parameters. We will now discuss one by one how we determine/estimate these parameters.

Demand parameters

- m the number of SKU's in the inventory control system
We have seen in the current situation analysis (Chapter 2) that there are 90 parts of the A/B series selected as critical parts. Therefore, 90 SKU's need to be controlled by the inventory control system, so m will be equal to 90. The demand information of all these 90 SKU's need to be determined.
- μ_i the mean week demand of SKU i
The first demand parameter that needs to be determined is the mean demand of every SKU. We mentioned before that everything will be expressed in weeks, so we need to find the mean week demand. This will be estimated by taking the average of the week demands of every week of the last year. These week demands have been determined by grouping the day demands of Monday through Friday. The demand of a SKU occurs when the production of a machine that contains this SKU is started. The resulting mean week demands for all SKU's can be found in Table C.1 in Appendix C: Values of inventory control system parameters. The estimation of these mean week demands according to the described procedure has been carried out in a separate Excel file, which is given to the company. The manual of the inventory control system (Appendix D: Manual of the inventory control system) gives a very detailed, step by step explanation on how to repeat this estimation with new data in the future.
- σ_i^2 the variance of the week demand of SKU i
The variance of the week demand is the second demand parameter that needs to be determined. It is estimated in the same way as the mean week demand, but now by taking the variance of the week demands. The variance of the week demand is also found in Table C.1 in Appendix C: Values of inventory control system parameters, for every SKU. The

estimation of these variances has been carried out in the same Excel file as the estimation of the mean week demands. The manual again gives an detailed explanation of this estimation.

- $t_{i\ max}$ the largest possible demand transaction size for SKU i
The largest possible demand transaction size determines for how much values the compounding distribution needs to be evaluated. This parameter is in principle only relevant for the compound Poisson distribution as only in that case demand transaction sizes larger than one are possible. On the one hand, a very large $t_{i\ max}$ means a lot of computational effort and thus time. On the other hand, a very small $t_{i\ max}$ means that the probability distribution of total demand x_i is not modelled correctly. Namely, the calculation of the probabilities $p_{x_i}(x_0)$ is constrained by the value of $t_{i\ max}$ then. The objective is to find a value of $t_{i\ max}$ that is large enough to not constrain the calculation of the probabilities $p_{x_i}(x_0)$, but small enough to limit the computational time. After some testing, it appears that 50 is a good value for the $t_{i\ max}$ of every SKU. Therefore, this value will be used in the inventory control system for every SKU. If errors regarding the calculation of the probability distribution of x_i occur with new demand data, then the values of $t_{i\ max}$ should be raised.
- $y_{i\ max}$ the largest possible demand during the lead time (D_L) for SKU i
We have seen in the methodology of the capacity check extension (Subsection 4.2.4) that the expected demand during lead time $E(y_i)$ of every SKU i is required. A summation until infinity is required for the calculation of $E(y_i)$. It is obviously not possible to implement that in the inventory control system, so instead we use a self-determined maximum of the summation, called $y_{i\ max}$. Just as with $t_{i\ max}$, the computational time becomes large when $y_{i\ max}$ is very large. However, when the $y_{i\ max}$ is very small, the expected value will not be calculated correctly. Namely, high but reasonable values of y_i will not be taken into account by the calculation of the expected value. The objective is, as was the case with $t_{i\ max}$, to find a value of $y_{i\ max}$ that is large enough to find a good expected value of y_i , but small enough to limit the computational time. Again, it appeared that 50 is a good value for the $y_{i\ max}$ of every SKU. Hence, $y_{i\ max}$ will be set equal to 50 in the inventory control system for every SKU. If the demand and/or lead time increase(s) in the future, then the values of $y_{i\ max}$ should possibly be raised.

Service levels

- P_i the service level of SKU i
A service level needs to be defined for each SKU. The service measure that is used is P_1 : the probability that the replenishment cycle ends with no backorders. The service level should be set by the company. We have seen in Chapter 1 that the company wants to prevent production stops and emergency deliveries, so the company does not want backorders to occur. It therefore strives for a high service level. It was agreed with the company to set the service level initially at 98% in the inventory control system for each SKU. However, we will test a range of high service levels in Chapter 5.

Time parameters

- L_i the lead time of SKU i
In Subsection 2.4.3 the lead time has been defined as the time period between the initiation of a replenishment order and the delivery of that order. We also saw that the lead time depends on the location (which external warehouse), but for simplicity we set the lead time equal to one day in the inventory control system for each SKU. One day is the estimate for the maximum lead time, so we are always on the safe side if we assume a lead time of one day. Since we define everything in (work)weeks, we set $L_i = \frac{1}{5}$ week (one day) for each SKU i .

- Range for the base period R_0 (R_0^{min} , R_0^{max} and ΔR_0)
The base period is the interval between replenishments. A SKU can be replenished after each base period or after a multiple of the base period. From the methodology of the inventory model (Subsection 4.2.2) it became clear that the (adjusted) HCTSP algorithm searches the R_0 for which the best solution (minimum EC) is attained. To this end, it evaluates solutions with a R_0 in the range between R_0^{min} and R_0^{max} with steps of ΔR_0 . We need to determine a reasonable range of values for R_0 to consider. It is expected that one replenishment per week is a reasonable number. Therefore, we will choose a range of R_0 values centred around one week. More specifically, we will set $R_0^{min} = \frac{1}{5}$ week (one day), $\Delta R_0 = \frac{1}{5}$ week (one day) and $R_0^{max} = 2$ weeks. This means that we will consider ten values for base period R_0 : one day up to and including ten days (two weeks since a week consists of five working days), in steps of one day.

Cost parameters

- K the major order cost
The first cost parameter that needs to be estimated is the major order cost. This is the fixed cost associated with placing a replenishment order. This cost is independent of the number of SKU's in the order. In this project, the major order cost will be defined as the transportation cost. This is the cost of arranging and driving a truck with trailer from the warehouse(s) to the production location. It does not matter if the trailer is partially or fully loaded, the transportation cost should be the same. Nevertheless, the transportation cost is dependent on the (distance of the) external warehouse and also on the number of external warehouses that is visited with the replenishment. However, the used inventory model requires us to give a single, fixed value for the major order cost. We therefore need to determine a reasonable single estimate: we choose a weighted average transportation cost as estimate for K . For this, we need to determine the transportation costs of all the individual trailer movements. A trailer movement is defined as the transportation of goods with a trailer between two locations of the company. Next, we need to list all the possible combinations of trailer movements and determine the transportation cost of each combination. We thereby assume that a single trailer movement is sufficient for a single external warehouse, i.e. the part of the replenishment order that comes from a certain external warehouse always fits in one trailer. Furthermore, the weight of each combination needs to be determined as the relative frequency of that combination. These relative frequencies can be determined by looking at the historic move files for the occurrence of the different combinations of trailer movements. Eventually, the weighted average transportation cost can be determined by multiplying, for each combination of trailer movements, the transportation cost with the weight and subsequently summing over all these products. This procedure has been executed using the move files from the whole of the last year in a separate Excel file, which is given to the company. The tables from this file regarding the major order cost determination are given as Table C.2, Table C.3 and Table C.4 in Appendix C: Values of inventory control system parameters. The resulting major order cost K can also be found here. The manual of the inventory control system (Appendix D: Manual of the inventory control system) gives a detailed, step by step explanation on how these tables were constructed.
- k_i the minor order cost for including SKU i in the order
The second order cost that should be determined, is the minor order cost. This cost is, in contrast to the major order cost, dependent on the SKU: it is the cost of including the SKU in

the replenishment order. In this project, the minor order cost of a SKU will be defined as the handling cost of a SKU. This is the cost of loading the trailer at the external warehouse(s) and unloading it at the production location for that particular SKU. A lot of a SKU is transported on a pallet and the handling of a SKU therefore consists of putting a pallet in and out of a trailer. The determination of the handling cost of a SKU will thus be based on the handling cost of a pallet and the number of pallets to be handled. The estimation will be as follows. Regarding the handling cost of a pallet, we can determine an average handling cost per pallet. We can do so because we know the handling cost of loading and unloading a full truckload and we know the amount of loading meters that a full truckload can carry. Since we assume that nothing is stacked in the trailer, only the trailer floor will be filled, this amount equals 13.6 loading meters. Furthermore, we can determine the average amount of loading meters that a pallet occupies. These average pallet dimensions can be calculated as a weighted average of the pallet dimensions of every SKU, where the relative frequencies of the SKU's form the weights. Subsequently, we can determine the average number of pallets per (full) truck by dividing 13.6 by the average pallet dimensions. The handling cost per pallet can then be calculated by dividing the handling cost of a full truckload by the average number of pallets per (full) truck.

Regarding the number of pallets to be handled, we should determine the number of pallets of a SKU that is ordered in a replenishment containing that SKU. However, this is an output of the model, but it is already needed as input for a cost parameter of the model. This means that the number of pallets that is going to be ordered at once, should be estimated for every SKU beforehand. Two factors influence this number: the demand of the SKU and the quantity of the SKU on the pallet (lot quantity). We will therefore calculate the ratio D_i/LQ_i for every SKU i , where D_i is the demand per year and LQ_i the lot quantity of SKU i . This ratio denotes the number of pallets required per year. If this ratio is small, it is unlikely that more than one pallet is ordered at a time. However, if the ratio is large, multiple pallets are likely ordered in the same replenishment order. For example, when the ratio is equal to 12, on average one pallet per month is required and it is likely that only one pallet is ordered at a time. We need to determine a boundary for when we assume that more than one pallet is ordered at a time. We set this boundary at 52, which means that on average one pallet is needed per week. For SKU's with a ratio lower than 52, it is assumed that only one pallet is ordered at a time. For SKU's with a ratio higher than 52, we assume that the amount of pallets is ordered so that a week of demand is covered. This amount can be calculated by dividing the ratio by 52 and rounding up to the nearest integer. This assumption is in line with our expectation that one replenishment per week is a reasonable estimate.

Eventually, we need to multiply the estimated number of pallets in a replenishment order by the handling cost per pallet to find the minor order cost k_i . This procedure has been carried out for every SKU in the same Excel file as the estimation of the major order cost K . The tables from this file regarding the minor order cost determination are given as Table C.5 and Table C.6 in Appendix C: Values of inventory control system parameters. The resulting minor order costs k_i for every SKU i can also be found here. The manual of the inventory control system (Appendix D: Manual of the inventory control system) gives a detailed, step by step explanation on how these tables were constructed.

- h_i the holding cost rate of SKU i

The last cost parameter that should be determined is the holding cost rate. The holding cost is also dependent on the SKU, so we need to find the h_i for every SKU i . It is defined as a rate: the holding cost per unit per week. The company already defines holding cost rates per unit per time-unit for all the C SKU's. It is therefore easy to derive the holding cost rates per unit

per week from them for every C SKU. On the other hand, the holding cost rates for the L SKU's are defined per square metre per time-unit at the company. We therefore need to determine the amount of square metres that one unit occupies to find the holding cost rate in the desired form (per unit per time-unit). This amount can be estimated by dividing the pallet dimensions in square metre by the quantity on the pallet. Both the values for the pallet dimensions and the quantity on the pallet are already found when we determined the minor order cost and can be used again here. The determination of the holding cost rates for every SKU has been executed in the same Excel file as the determination of the other cost parameters. The tables from this file regarding the holding cost determination are given as Table C.7, Table C.8, Table C.9 and Table C.10 in Appendix C: Values of inventory control system parameters. The resulting holding cost rates h_i for every SKU i can also be found here. The manual of the inventory control system (Appendix D: Manual of the inventory control system) gives a detailed, step by step explanation on how these tables were constructed.

5. Results of the implemented inventory control system

The inventory control system and its implementation in Excel have been described in Chapter 4. The inventory control system can now be used to calculate the replenishment schedule and subsequently the replenishment decisions that follow from it. This chapter will show and analyse the results of the implemented inventory control system. Section 5.1 contains the raw results, obtained as the output of the inventory control system. In order to determine which parameter influences the output in what degree, sensitivity analyses on different parameters have been carried out in Section 5.2. Finally, Section 5.3 shows the improvement that will be achieved when the inventory control system is used.

5.1. Raw results

The implementation of the inventory control system in Excel is explained in Section 4.3. The Excel file that was mentioned and shown there, will now be used to calculate the raw results of the inventory control system. We also saw there that the inventory control system requires parameters as input for the calculations. Values for all these parameters were determined in Section 4.4. These values will now be used for the calculation of the raw results.

We start with inputting the values of Section 4.4. This comes down to inputting several demand and cost parameters for $m = 90$ SKU's. Besides, for each SKU, the service level P_i is set at 98% and the lead time L_i is set at 0.2 weeks (1 workday). Subsequently, the range for the base period R_0 should be set according to the range given in Section 4.4: $R_0^{min} = 0.2$ working weeks, $R_0^{max} = 2.0$ working weeks and $\Delta R_0 = 0.2$ working weeks. We can now run the first part of the inventory control system: the calculation of the replenishment schedule. The inventory control system now evaluates all replenishment schedules with a base period (R_0) in the range 0.2 - 2 working weeks with steps of 0.2 working weeks. Finally, the inventory control system chooses the best replenishment schedule, which is the schedule with the least total expected costs (EC) per week, while still satisfying the service level P_i of 98% for every SKU. The total expected costs (EC) and mean service levels ($Mean P_i$) of the best evaluated replenishment schedules for each R_0 are given in Table 5.1. The best replenishment schedule, which is chosen by the inventory control system, is highlighted in green in the table.

Table 5.1 The total expected costs (EC) and mean service levels ($Mean P_i$) of the best replenishment schedules in the evaluated R_0 range.

| R_0 (in working weeks) | EC (in euros) | $Mean P_i$ (in %) |
|--------------------------|-----------------|-------------------|
| 0.2 | 712.18 | 98.215 |
| 0.4 | 539.27 | 98.244 |
| 0.6 | 481.30 | 98.235 |
| 0.8 | 455.82 | 98.267 |
| 1.0 | 440.74 | 98.313 |
| 1.2 | 429.21 | 98.312 |
| 1.4 | 419.64 | 98.266 |
| 1.6 | 417.21 | 98.342 |
| 1.8 | 417.63 | 98.358 |
| 2.0 | 413.89 | 98.364 |

It appears that the replenishment schedule with the least total expected costs is the one with a base period R_0 of 2.0 working weeks. This means that it is optimal to do a replenishment once every 2 working weeks (10 working days). The resulting total expected costs will then be equal to €413.89 and the mean service level will be equal to 98.364%. These results are achieved with a specific set of

review periods (R_i), given by Table 5.2. We can see that only 42.22% of the SKU's are replenished after every base period of 2 working weeks. The rest of the SKU's are replenished after a multiple of the base period, such as 4 working weeks (31.11%) and 6 working weeks (17.78%). The largest review period is 16, since one SKU should be best replenished once every 16 working weeks.

Table 5.2 The review periods (R_i) belonging to the best replenishment schedule ($R_0 = 2.0$), together with the number and percentage of SKU's that have a certain R_i .

| R_i (in working weeks) | Number of SKU's with this R_i | Percentage of SKU's with this R_i |
|--------------------------|---------------------------------|-------------------------------------|
| 2 | 38 | 42.22% |
| 4 | 28 | 31.11% |
| 6 | 16 | 17.78% |
| 8 | 4 | 4.44% |
| 10 | 1 | 1.11% |
| 12 | 2 | 2.22% |
| 16 | 1 | 1.11% |

On the basis of these R_i parameters, the replenishment schedule can be constructed. Next to the R_i parameters, the inventory control system gives the S_i parameters for each SKU. These are the order-up-to-levels which determine to which level the SKU's should be ordered with a replenishment. The order-up-to-levels (S_i) which belong to the best replenishment schedule are given in Table 5.3. We see a large deviation in the order-up-to-levels, with the smallest being 1 and the largest being 187. Almost half of the SKU's have an order-up-to-level between 10 and 25. The large differences are logical since the order-up-to-level depends very much on the review period and the demand of the SKU, which can be very different.

Table 5.3 The order-up-to-levels (S_i) belonging to the best replenishment schedule ($R_0 = 2.0$), together with the number and percentage of SKU's that have a certain S_i .

| S_i | Number of SKU's with this S_i | Percentage of SKU's with this S_i | S_i | Number of SKU's with this S_i | Percentage of SKU's with this S_i |
|-------|---------------------------------|-------------------------------------|-------|---------------------------------|-------------------------------------|
| 1 | 1 | 1.11% | 25 | 3 | 3.33% |
| 3 | 2 | 2.22% | 26 | 5 | 5.56% |
| 4 | 3 | 3.33% | 27 | 1 | 1.11% |
| 5 | 3 | 3.33% | 28 | 2 | 2.22% |
| 6 | 2 | 2.22% | 29 | 3 | 3.33% |
| 8 | 3 | 3.33% | 30 | 1 | 1.11% |
| 9 | 1 | 1.11% | 32 | 4 | 4.44% |
| 10 | 8 | 8.89% | 34 | 2 | 2.22% |
| 11 | 3 | 3.33% | 35 | 2 | 2.22% |
| 12 | 13 | 14.44% | 39 | 1 | 1.11% |
| 13 | 3 | 3.33% | 41 | 1 | 1.11% |
| 14 | 4 | 4.44% | 45 | 1 | 1.11% |
| 15 | 2 | 2.22% | 46 | 1 | 1.11% |
| 16 | 1 | 1.11% | 53 | 1 | 1.11% |
| 18 | 1 | 1.11% | 73 | 2 | 2.22% |
| 20 | 4 | 4.44% | 100 | 1 | 1.11% |
| 23 | 1 | 1.11% | 146 | 2 | 2.22% |
| 24 | 1 | 1.11% | 187 | 1 | 1.11% |

The complete set of (R_i, S_i) parameters for each SKU are given in Table F.1 in Appendix F: Detailed results of the best replenishment schedule. Besides, the service levels per SKU are given there as $RealP_i$ for each SKU i . All individual service levels are above 98%, because the 98% service level constraint applies to all individual SKU's. The mean of these $RealP_i$'s is given as $Mean P_i$ in Table 5.1. Finally, the expected costs per SKU are given there as EC_i for each SKU i . The sum of these EC_i 's, together with $K / 2$ (major order cost per week), is given as EC in Table 5.1.

We will now look deeper into the results of Table 5.1. We plotted the results in the graphs of Figure 5.1. The total expected costs EC per week (blue line) and mean service level $Mean P_i$ (orange line) are plotted against the evaluated base period (R_0) values. When we look at the EC line, we can see a clear pattern. Namely, the expected costs decrease with increasing base period length, except for one instance where the expected costs remained approximately the same. Furthermore, the expected costs decrease less quickly with each step of R_0 . There is no clear pattern visible in the $Mean P_i$ line. In general, the mean service level increases with increasing base period length. However, there are drops in mean service level at 0.6, 1.2 and (a large one at) 1.4 working weeks. Nevertheless, the mean service level is always above 98.00%. This is logical since the individual service levels should all be above 98.00% as this is set as a constraint in the inventory control system. In fact, the mean service level is even always above 98.20%. The rest of the $Mean P_i$ line cannot be logically explained, although it is desirable that the mean service level increases with increasing base period. Namely, it means that simultaneously the expected costs go down and the mean service level goes up.




Figure 5.1 The total expected costs EC per week and mean service level $Mean P_i$ from Table 5.1 are plotted against the evaluated base period (R_0) values. The green dot denotes the optimal solution (least EC) as found by the inventory control system.

On the other hand, the pattern of the EC line can be logically explained when we recall that the total expected costs consist of the (minor and major) order cost and the holding cost. When R_0 is very small, i.e. one or two working days, a lot of replenishments with small quantities will occur. This means that the order cost will be extremely large and the holding cost will be extremely small. However, the expected cost can be greatly reduced if the order cost will be lower at the expense of

higher holding cost. This can be explained by the basic EOQ model that we gave in Section 3.4.1. We saw there that the total costs consist of two cost terms: a reciprocal ordering cost term and a linear holding cost term. Thus, in the beginning the order cost will decrease a lot faster than the holding cost will increase, so the total expected cost will decrease a lot. As the base period becomes larger, the order cost will decrease only a bit faster than the holding cost will increase, so the total expected cost will decrease only a bit. The decrease in total expected cost will become less and less, until the decrease stops and the increase in total expected costs start. The point where this happens is the optimum, the total expected costs are the lowest here.

We can see a similar pattern in the *EC* line of Figure 5.1. From $R_0 = 0.2$ to approximately $R_0 = 1.4$ the *EC* decreases significantly, first very fast and then increasingly slower. From $R_0 = 1.4$ to $R_0 = 2.0$, the *EC* does not change very much. This is in accordance with the EOQ model, where the *EC* does not change much around the optimum. However, it is striking to see that the *EC* line is not convex. Namely, the *EC* rises for the first time from $R_0 = 1.6$ to $R_0 = 1.8$, but it decreases again from $R_0 = 1.8$ to $R_0 = 2.0$. We already noted this non-convex behaviour when we made the inventory control system, see Section 4.2.2. If we had used the original HCTSP algorithm in our inventory control system, then we would incorrectly have chosen $R_0 = 1.6$ as the optimal base period. Since we adjusted the HCTSP algorithm and implemented the adjusted version in the inventory control system, it correctly chooses the optimal base period of 2.0 working weeks.

There is one more aspect about the *EC* line in Figure 5.1 that we are going to analyse. The optimal replenishment schedule happens to be at $R_0 = 2.0$ working weeks. However, this is also the end of the evaluated R_0 range. The question now is whether the optimal solution really occurs at $R_0 = 2.0$ or at a higher R_0 which is not evaluated because the inventory control system stops at 2.0. The latter is entirely possible, especially since the *EC* decreased from $R_0 = 1.8$ to $R_0 = 2.0$, so it could decrease further from $R_0 = 2.0$ to $R_0 = 2.2$. Since it is an interesting question, we will analyse it by means of a new figure: Figure 5.2. This figure shows again the total expected costs *EC* per week plotted against the base period R_0 . However, this time the *EC*'s for $R_0 = 2.2$ through 4.0 are also evaluated and plotted.




Figure 5.2 The total expected costs (*EC*) per week is plotted against the base period (R_0) on the range 0.2 through 4.0. The green dot denotes the optimal solution (least *EC*) as found by the inventory control system.

The first part of the graph ($R_0 = 0.2$ through 2.0) is exactly the same as in Figure 5.1. The graph has been extended with the *EC* values for $R_0 = 2.2$ through 4.0 in steps of 0.2. We can see that the *EC*

immediately increases after $R_0 = 2.0$ and it keeps increasing. This means that the optimal solution indeed occurs at $R_0 = 2.0$. Thus, the optimal solution coincidentally occurred at the end of the original evaluated R_0 range.

When we look at the shape of the complete EC graph, we notice the same behaviour as with the basic EOQ model. First, it approximately takes on the shape of a reciprocal function graph (coming from the reciprocal ordering cost term from the EOQ model), since the order costs play a much larger role than the holding cost at small R_0 's. Then, the graph is relatively flat around the optimum, just like with the EOQ model. Finally, it approximately takes on the shape of a linear function graph (coming from the linear holding cost term from the EOQ model), since the holding costs play a much larger role than the order costs at large R_0 's.

We have now discussed the results of the replenishment schedule, but we did not include the buffer inventory capacity check yet. To find the results of the capacity check, we first need to input the capacity information in the inventory control system. Subsequently, we run the inventory control system again with the same settings as before. Since the capacity check is not a constraint in the inventory control system but merely a check on the results of the replenishment schedule, it does not influence the replenishment schedule. The (results of the) replenishment schedule will therefore be the same as presented above. We will only give the additional capacity check results now.

Recall that we divided the buffer inventory capacity in three parts in Subsection 4.2.4:

- 4. Individual capacities, for (almost) every SKU individually
- 5. Shared capacities, for certain sets of SKU's
- 6. Overflow capacity, one large capacity for one large set of SKU's

SKU's are stored on pallets in these capacities (the quantity of a SKU on a pallet depends on the SKU). The number of pallets that a capacity can hold is determined in the capacity information input. The capacity check calculates the number of pallets that a capacity should hold. There are three cases possible for each capacity: it has a certain number of pallets short, it has a certain number of pallets remaining or it has just the right amount of pallets. The results of the individual capacities are summarised in Table 5.4. The detailed capacity check results of the individual capacities per SKU are given in Table F.2.

Table 5.4 Summary of the capacity check results of the individual capacities.

| Case | Number of individual capacities with this case | Total number of pallets short/remaining |
|----------------------------------|--|---|
| Pallet(s) short | 14 (6) ⁵ | 156 (12) ⁵ |
| Pallet(s) remaining | 2 | 2 |
| Just the right amount of pallets | 74 | - |

We can see that the majority of the individual capacities (74 of the 90) have just the right size, i.e. no pallets short or remaining. However, there are also some (14) individual capacities with a too small size. As discussed in the footnote, 8 of these 14 individual capacities have a size of 0 (not part of the buffer capacity), which explains the very high number of pallets short. If we do not take these into

⁵ There are 8 SKU's which are not kept in the buffer inventory. Therefore, the individual capacities of these 8 SKU's are equal to 0 and all the pallets that should be kept in the buffer inventory for these SKU's according to the replenishment schedule are counted towards the "Total number of pallets short". This comes down to 144 pallets. It is not completely fair to include these 144 pallets since (a part of) these pallets is/are already stored outside the buffer inventory. If we do not take these SKU's into account, we get the values between brackets, which give a more realistic view of the pallet shortage in the current buffer inventory.

account, there are only 12 pallets short.

Two missing pallet places can be compensated with the two pallets places that remain unused due to the two individual capacities with a total of two pallets remaining. This is on the condition that the unused pallet places are of the same size or larger than the missing pallet places. This appears to be the case if we look at the detailed results in Table F.2. The net result is then a shortage of 154 pallets regarding the individual capacities part of the buffer inventory capacity.

The next part of the buffer inventory capacity is the shared capacities part. We have defined five shared capacities in Section 4.3.2: SR, SC, SS, SMD and ST. The last part of the buffer inventory capacity is the overflow capacity part. The capacity check results of both parts are given in Table 5.5.

Table 5.5 The capacity check results of the shared capacities and the overflow capacity.

| Type of capacity | Capacity check result |
|-------------------------------|--|
| Shared RB part capacity (SR) | 4 RB part pallets remaining |
| Shared CB part capacity (SC) | 16 CB part pallets short |
| Shared SE part capacity (SS) | 1 SE part SP short |
| Shared MD part capacity (SMD) | 20 MD part LP's short |
| Shared TR parts capacity (ST) | 2 SP's short OR 2 BP's short and 1 SP remaining ⁶ |
| Overflow capacity (OC) | 3 LP's and 1 BP remaining ⁶ |

We see very different results per capacity: some capacities have many pallets short while some have only a few pallets short. There are also capacities which have a few pallets remaining. Therefore, it is possible to increase the size of some capacities at the expense of the size of other capacities. However, the net result will still be a capacity shortage regarding the shared capacities part of the buffer inventory capacity.

We have now given the raw results of the replenishment schedule and the buffer inventory capacity check. However, this is only part of the output of the inventory control system, since it also gives the replenishment decisions as output. Nevertheless, we cannot give raw results of the replenishment decisions, since they completely depend on the replenishment day. Namely, the (number of) SKU's that are replenished, are dependent on the replenishment day and the quantity that should be ordered of each SKU is also dependent on the inventory levels at that day. Therefore, the replenishment decisions of two replenishment days will likely never be the same. Hence, we cannot give one set of raw results of the replenishment decisions. The same goes for the outstanding orders check.

5.2. Sensitivity analyses

This section is about sensitivity analyses on different parameters. With a sensitivity analysis we analyse how uncertainty in the output of the inventory control system is related to uncertainty in one of its input parameters. We determined estimates for the input parameters in Section 4.4. However, there is uncertainty if these estimates really reflect the real values of these parameters and we are uncertain how deviations affect the output of the inventory control system. Therefore, we will recalculate the output with alternative values for parameters. We will vary only one parameter at once and thus keep all other parameters fixed. In this way, we can understand the relationship between the parameter that is changed and the output of the inventory control system. Furthermore, we can test the robustness of the optimal replenishment schedule when changing that

⁶ Please note that shared capacity ST and the overflow capacity can hold a combination of different pallet types. Therefore, the capacity check result is also a combination of different pallet types.

parameter. We will change several parameters: In Subsection 5.2.1 we will change the cost parameters (K, k_i, h_i) one by one and compare the results. Subsequently, a sensitivity analysis with changing service level (P_i) will be carried out in Subsection 5.2.2. Finally, we will perform a sensitivity analysis with changing buffer inventory capacity in Subsection 5.2.3.

5.2.1. Sensitivity analysis with changing cost parameters (K, k_i, h_i)

We will start with varying only the value of the major order cost K . In Section 4.4 we determined an estimate for the major order cost and defined it as K . However, this estimate is possibly not close to the real value of the major order cost. Since the deviation could be large in both the positive and negative direction, we will test a large range of values around K : $0.25 \cdot K, 0.5 \cdot K, 0.75 \cdot K, 1.25 \cdot K, 1.5 \cdot K, 1.75 \cdot K$ and $2 \cdot K$. All other parameters will be kept fixed at the values that were used to find the optimal replenishment schedule of Section 5.1. We are interested in the influence of major order cost K on the expected costs EC of the optimal replenishment schedule. The question that we ask ourselves is: “What will happen to the EC if the optimal replenishment schedule of Section 5.1 will be used while K would be x times lower/higher than estimated?” We will only evaluate EC at $R_0 = 2.0$ weeks, since this is the R_0 that belongs to the optimal replenishment schedule. First, we will calculate the EC value with $0.25 \cdot K$ at $R_0 = 2.0$. To this end, we multiply the value for K with 0.25 in the input of the inventory control system. Subsequently, we run the inventory control system and only record the EC value at $R_0 = 2.0$. We do the same for every value around K described above. We have plotted the resulting EC values against the factor that K has been multiplied with, see the blue points in the scatterplot of Figure 5.3.

In the same manner, we varied the minor order cost k_i and the holding cost h_i . We multiplied the costs with the same factors as with the major order cost K . However, we needed to change k_i and h_i 90 times, since there are 90 k_i and h_i values (one k_i and h_i for every controlled SKU) as opposed to just one K value. After the inventory control system had run with all values of k_i and h_i (one after the other) and all corresponding EC 's were recorded, the EC values have been added to the scatterplot of Figure 5.3. The orange points correspond to the EC values with changed minor order costs and the grey points to the EC values with changed holding costs.

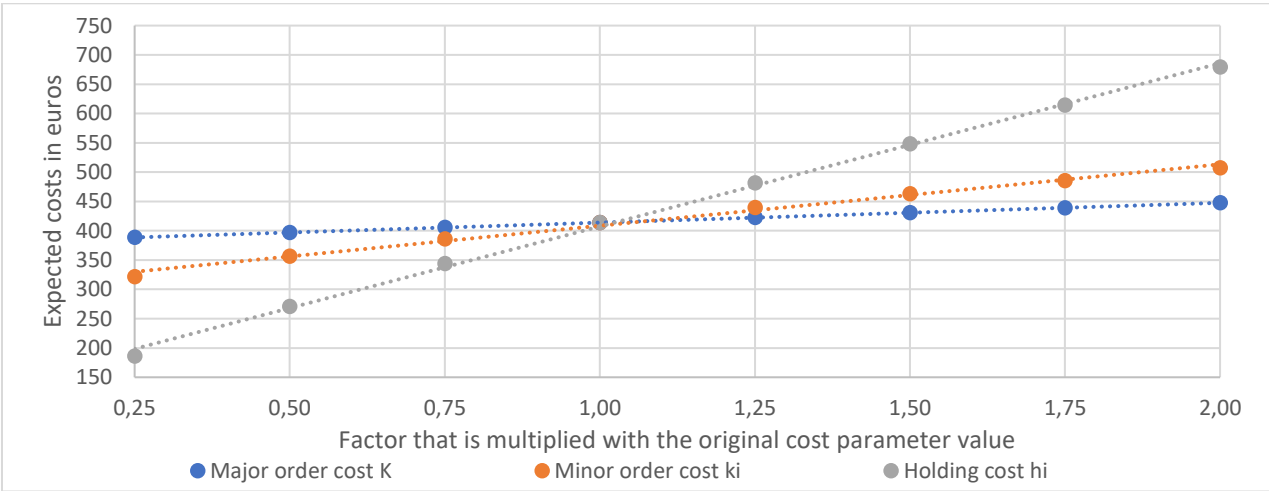


Figure 5.3 This scatterplot shows the relationship between the expected costs EC and values for the three cost parameters K, k_i and h_i . The values of the three cost parameters are given as a range of factors that the cost parameters have been multiplied with (1 corresponds to the original value for K, k_i and h_i as determined in Section 4.4).

We also added three linear trend lines to the scatterplot of Figure 5.3, one for each set of EC values. Some conclusions can be drawn when we look at the points and trend lines. Firstly, for each set of EC values, all the points approximately lie on the corresponding linear trend

line with a positive slope. This means that both K and k_i and h_i are almost perfectly positively correlated with EC . This can be explained by looking at the objective function of mathematical programming Problem 4.9. We can see there that K , k_i and h_i only occur as constants in different terms that are added (K and k_i are numerators in fractions and h_i is a factor of a product). Hence, a (almost) total positive linear correlation between $K/k_i/h_i$ and EC was to be expected. Secondly, the slope of the three trend lines differs significantly. Since the EC values follow the trend lines, it means that the influence of the three cost parameters on EC differs significantly. The cost parameter corresponding to the trend line with the steepest slope has the most influence on the EC of the (optimal) replenishment schedule. It turns out that a change in the holding cost parameter value has the largest influence on EC , as the grey trend line is the steepest. To be more precise, the grey trend line has a slope of 278.46. This means that the EC approximately increase (decrease) with €278.46 when h_i increases (decreases) with one time the original value of the holding cost parameter. Regarding the minor order cost k_i , the increase/decrease in EC is only around €104.75 (slope of the orange trend line). The major order cost K has the least influence on EC since its (blue) trend line has a slope of only 33.74. Thus, the EC only increase (decrease) with approximately €33.74 when the value of K increases (decreases) with the value of the original estimate of K .

To conclude, we can now answer the question that we asked ourselves: “What will happen to the EC if the optimal replenishment schedule of Section 5.1 will be used while $K/k_i/h_i$ would be x times lower/higher than estimated?” Namely, K , k_i and h_i all show a positive correlation with EC : the EC will increase when $K/k_i/h_i$ increases and the other way around. In fact, these correlations are all almost perfectly positively linear. However, the degree in which the EC increase/decrease differs per cost parameter. A change in holding cost parameter values results in the largest change in EC value. Therefore, the inventory control system is most sensitive to a misestimation of the holding cost. If we misestimate the major or minor order cost, the EC of the optimal replenishment schedule will be less affected.

We just carried out a sensitivity analysis to see the influence on EC given (the base period of) the optimal replenishment schedule ($R_0 = 2.0$). However, the optimal replenishment schedule might no longer occur at $R_0 = 2.0$ and the EC will be different in that case. Therefore, it is interesting to study the influence that a change in a cost parameter has on the optimal base period. Since we saw that the EC is the most sensitive to a change in holding cost at a given base period, we expect that the holding cost parameter has the most influence on a change in optimal base period as well. We will therefore only do a sensitivity analysis with changing holding cost parameter values. Thus, the question that we ask ourselves here is: “What would be the optimal base period when the holding cost would be x times lower/higher than estimated?” We will only test the following values around h_i : $0.5 * h_i$, h_i , $1.5 * h_i$ and $2 * h_i$. Obviously, we will change the holding cost parameter values for each of the 90 SKU’s. We run the inventory control system for every holding cost parameter value described above. We record the EC values for the base periods $R_0 = 0.2$ through 4.0 weeks each time. We used the resulting EC values to plot EC lines for every holding cost parameter value in Figure 5.4. The expected costs EC are plotted against the base period R_0 in the figure.




Figure 5.4 This figure shows 4 expected costs lines over the R_0 range 0.2 through 4.0 working weeks. Each line corresponds to a certain set of values for the holding cost parameter. The green dots denote the optimal solution (least EC) for every holding cost parameter value.

We can see in Figure 5.4 that the shape of the EC line changes quite a bit when the value for the holding cost parameter changes. It becomes clear that the minimum EC is reached earlier when the holding cost is larger and the other way around. Hence, the optimum is reached at a shorter base period R_0 when the holding cost increases and at a longer base period R_0 when the holding cost decreases. This can be logically explained when we recall the basic EOQ model again. We saw a trade-off between holding cost and order cost there. This trade-off also applies here: when the holding cost becomes more important (increases), we will order more often (short base period) but in smaller quantities, so that less inventory is on hand. On the other side, if the holding cost becomes less important (decreases), the order cost becomes more important and we will order less often (longer base period). This corresponds with the shifts in the optima in Figure 5.4. Due to the trade-off, we expect shifts in the opposite direction when performing a similar sensitivity analysis with changing (major or minor) order cost.

We can now also answer the question that we asked ourselves: “What would be the optimal base period when the holding cost would be x times lower/higher than estimated?” We already know that the optimal base period R_0 will be equal to 2.0 working weeks with the original estimate of the holding cost. If the holding cost would be half the original estimate, the optimal R_0 will be equal to 3.0 weeks (increase of 1.0 weeks). The EC will drop to €264.80, which is a decrease of 2.12% compared to the EC of €270.53 corresponding to $R_0 = 2.0$. On the other side, if the holding cost would be 1.5 times the original value, R_0 will be equal to 1.4 weeks (decrease of 0.6 weeks) and the EC will decrease with 1.45% from €547.95 to €540.00. Therefore, an overestimation of the holding cost (real holding cost lower than estimate) will have more impact on the optimum than an underestimation. If the holding cost would be double the original estimate, the optimal R_0 will also be equal to 1.4 weeks (decrease of 0.6 weeks), which is a further indication that an underestimation of the holding cost has less impact on the optimum.

5.2.2. Sensitivity analysis with changing service level (P_i)

We are now going to perform a sensitivity analysis with changing another parameter, namely the service level P_i , which should be set for every SKU i included in the inventory control system. The company should determine the height of the service level. Obviously, the company strives for a high service level, but this comes at a cost: a high service level means keeping much inventory on hand which leads to high holding cost. A trade-off should therefore be made by the company. In Section 4.4, it was decided to set the service level for each SKU at 98.0% in this project. Hence, this value has been used to find the optimal replenishment schedule of Section 5.1. Nevertheless, it is interesting to study the effect of other service levels on the expected costs. It will offer an insight into the relationship between service level and expected costs and thus helps the company in making the trade-off. We will therefore carry out a sensitivity analysis with changing P_i . We will take the same approach as in Subsection 5.2.1, which means we are first going to look at the influence of P_i on the EC of the optimal replenishment schedule, so only at the given optimal base period $R_0 = 2.0$. We will test the following service levels, which will be set for each SKU: 95.0%, 95.5%, 96.0%, 96.5%, 97.0%, 97.5%, 98.5%, 99.0%, 99.5% and 99.9%. These values were chosen because they are all high service levels and they are approximately centred around 98.0%. Since a service level of 100% can never be guaranteed, we will choose 99.9% instead. All other parameters will be kept fixed at the values that were used to find the optimal replenishment schedule of Section 5.1. The inventory control system was run for every service level described above and the EC value at $R_0 = 2.0$ had been recorded every time. The resulting EC values are plotted against the service level in Figure 5.5. The corresponding linear trend line is also plotted in the figure.

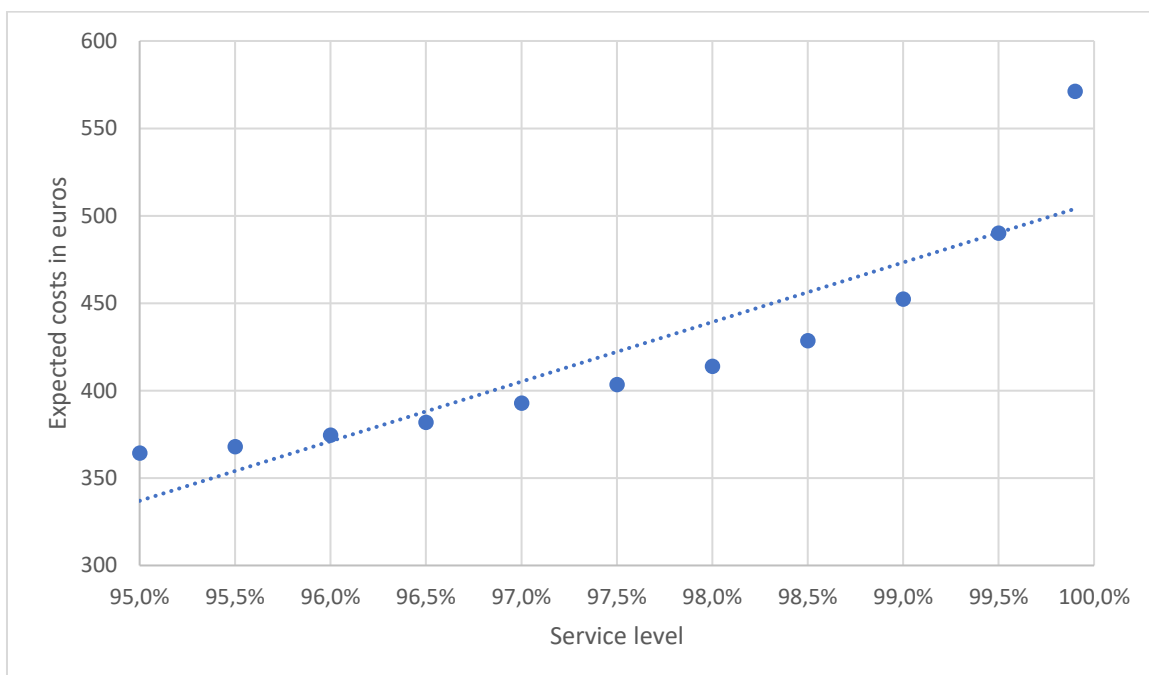


Figure 5.5 This scatterplot shows the correlation between service level and expected costs, where all points are evaluated at $R_0 = 2.0$. The service level has been varied in steps of 0.5% between 95.0% and 99.9% (since 100% is unattainable).

When we look at Figure 5.5, we can clearly see a positive correlation between the service level and expected costs, as the points show an upward trend. We already expected that because of the trade-off between service level and (holding) cost. However, the positive correlation is not linear, since the points do not follow the linear trend line. Nevertheless, we can observe a certain relationship: the expected costs increase more quickly with each equal step higher in service level. Thus, a per cent gain in service level becomes increasingly more expensive. This becomes especially extreme when

the service level is approaching 100.0%. It can be logically explained if we recall that we assumed future demand which is stochastic and unbounded. This means that there is always a (very) small probability of a very high demand. With a very high service level, even this very small probability should be covered by choosing a very high order-up-to-level. However, most of the time the demand will not be equal to (or not even close to) such a very high demand. Therefore, a large part of the inventory will remain on hand in most replenishment cycles, which results in very high holding cost. That explains the explosive increasing costs for very high service levels, although it also applies to the progressive increase in costs at lower service levels.

Just as in Subsection 5.2.1, we are now going to carry out a sensitivity analysis to analyse the influence that a change in service level has on the optimal base period. We will only test the following values for the service level: 95.0%, 96.0%, 97.0%, 98.0%, 99.0% and 99.9%. We will always set the service levels of all SKU's to the same value. The inventory control system was run for every service level described above and the EC values for the base periods $R_0 = 0.2$ through 4.0 weeks had been recorded each time. The resulting EC values were used to plot EC lines for every service level in Figure 5.6. The expected costs EC are plotted against the base period R_0 in the figure.




Figure 5.6 This figure shows 6 expected costs lines over the R_0 range 0.2 through 4.0 working weeks. Each line corresponds to a certain service level. The green dots denote the optimal solution (least EC) for every service level.

Some observations can be made when we look at Figure 5.6. Firstly, the distance between the EC lines increases with increasing service level. This is in accordance with the sensitivity analysis of Figure 5.5, where we observed this relationship at $R_0 = 2.0$. Now we can confirm that this also holds for the other base periods in the range $R_0 = 0.2$ through 4.0. Secondly, we can see that the shapes of the EC lines are generally the same. Consequently, the optimal base periods do not differ much for different service levels. We can see that the optimal base period fluctuates in the range $R_0 = 1.6$ through 2.2 working weeks. Therefore, we can conclude that the service level has not much influence on the optimal base period.

5.2.3. Sensitivity analysis with changing buffer inventory capacity size

The last sensitivity analysis that we will perform is a sensitivity analysis with changing buffer inventory capacity size. This sensitivity analysis will be different compared to the previous sensitivity analyses. Namely, this time it is not possible to observe the influence on *EC*, since the capacity is not an input for the inventory model. Therefore, the capacity does not influence the outcome (replenishment schedule) of the inventory control system. Thus, the value of *EC* will not be affected by the size of the capacity. The size of the capacity is not a constraint, it is only used in a check after the replenishment schedule has been made. With this check we predict if there will be pallets short or remaining in the buffer inventory, given the capacity of this inventory. Therefore, changing the size of the capacity does affect the number of pallets short/remaining. Hence, a sensitivity analysis on the number of pallets short/remaining with changing buffer inventory capacity size will be done.

Recall that we divided the buffer inventory capacity in three parts: individual capacities, shared capacities and an overflow capacity. We want to study the effect of changing one of these capacity sizes at a time. We will multiply the capacity size with the following factors: 0.8, 0.9, 1.1 and 1.2. These relatively small changes were chosen because these changes are still realistic. We start with changing the sizes of the individual capacities. To this end, we multiply the original capacity sizes of all the individual capacities with 0.8 in the input of the capacity check. Subsequently, we run the inventory control system with the capacity check activated. Afterwards, we sum up all the pallets short to find the total numbers of pallets short and we do the same for the total number of pallets remaining. Then, we subtract the total number of pallets remaining from the total number of pallets short to find the net number of pallets short. Note that we do not take the different pallet sizes into account, we just assume that every pallet is the same. This will be less precise, but it simplifies the sensitivity analysis as we only need to calculate and analyse one value. We calculate the net number of pallets short for every factor described above by running the inventory control system several times with the different factors. We have plotted the resulting net number of pallets short against the factor that the original individual capacities sizes have been multiplied with, see the blue points in the scatterplot of Figure 5.7. The same approach has been taken when changing the sizes of the shared capacities and the overflow capacity. The results are also plotted in Figure 5.7: the orange points correspond to the shared capacities and the grey points to the overflow capacity. Linear trend lines, one for each set of points, are also plotted in the figure.

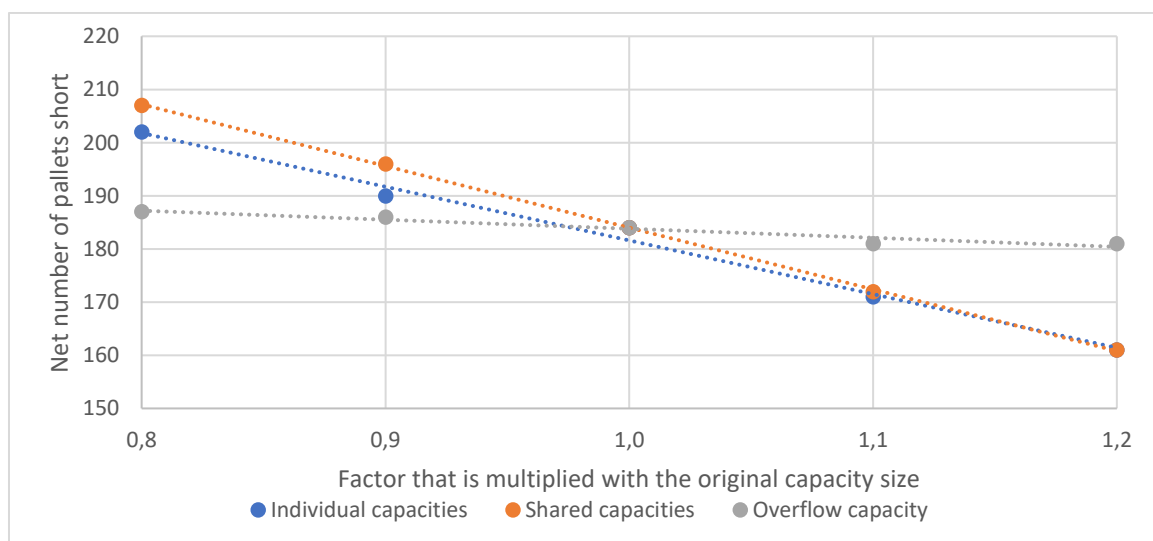


Figure 5.7 This scatterplot shows the relationship between the net number of pallets short and the capacity size for the three capacity parts. The values of the capacity size are given as a range of factors that the capacity size has been multiplied with (1 corresponds to the original capacity size as determined in Section 4.4).

When we look at Figure 5.7, we can draw some conclusions. Firstly, there exists a negative correlation between the net number of pallets short and the capacity size for each set of points (individual capacities, shared capacities and overflow capacity). This is logical since we expect that the pallet shortage decreases when the capacity size increases. Secondly, we see that the negative correlation is quite linear, as all points lie on or close to their corresponding linear trend lines. This is also not strange as it makes sense that every 10% added to or subtracted from the capacity size results in the same decrease or increase in number of pallets short. Thirdly, the slopes of the three trend lines are different. This means that the influence that each capacity part has on the net number of pallets short differs. We see that a change in size of the overflow capacity has the least effect on the net number of pallets short. The other two capacity parts have approximately the same effect on the net number of pallets short, although the shared capacities part has somewhat more influence when decreasing its size starting from the original capacity size. The orange points are namely clearly above the blue points at 0.8 and 0.9 in Figure 5.7.

In conclusion, the size of a capacity part, whether it is the individual capacities part or the shared capacities part or the overflow capacity part, is almost perfectly negatively linear correlated with the net number of pallets short. However, the three capacity parts have a different effect on the net number of pallets short: the individual capacities part and the shared capacities part have almost the same effect whereas the overflow capacity part has a clearly smaller effect.

5.3. Improvement compared to the current situation

The idea of implementing the inventory control system is to arrive at a new situation which is an improvement over the current situation. This section gives the improvements which are expected when the inventory control system will be implemented. Since the goal of the inventory model is to minimise the expected costs given a certain service level, our KPI is the expected costs. We will therefore measure the improvement in expected costs in Subsection 5.3.1. However, there are also other improvements expected. These improvements are not measurable but will certainly have an impact. We will list these improvements in Subsection 5.3.2.

5.3.1. Improvement in expected costs

In this subsection we are going to determine the improvement (if there is any) in expected costs (*EC*) when the devised inventory control system will be used. To this end, we will determine the *EC* of the current situation and compare it to the *EC* of the new situation. The *EC* of the current situation will be defined as the total expected costs over the past year. The *EC* of the new situation will be defined as the total expected costs per year given by the optimal replenishment schedule of the inventory control system. We will now first explain how these two *EC* values have been determined and then we will look at and compare the results.

Regarding the calculation of the *EC* of the current situation, we first need to recall that the total expected costs consist of three components: minor order cost, major order cost and holding cost. Therefore, we are going to determine the minor order cost, major order cost and holding cost from the past year separately and add them up afterwards to find the *EC* value.

The major order cost from the past year has been determined by counting all the trailer movements (containing the controlled SKU's) from the past year and multiplying them with the corresponding cost of trailer movement. However, this also includes replenishments to the bulk inventory whereas we only consider replenishments to the buffer inventory in our project. Furthermore, these trailer movements also contained SKU's from different products that are not controlled by the inventory control system. Thus, the cost of these trailer movements should actually be shared with these products. For these two reasons, the determined value for the major order cost per year is somewhat overestimated. The determination of the major order cost can be found in Table G.1 and

Table G.2 in Appendix G: Determination of the costs of the current situation.

The minor order cost from the past year has been determined by counting the number of pallets that were replenished in the past year and multiplying that number with the handling cost per pallet as found in Section 4.4. This has been done for all controlled SKU's and the resulting values were added up to find the minor order cost from the past year. However, this does only include the replenishments from the external warehouses and not the replenishments from the bulk inventory as there is no data recorded about the latter. Consequently, we only have data for 48 of the 90 controlled SKU's since the other SKU's were apparently only replenished from the bulk inventory. For this reason, the determined value for the minor order cost per year is quite lower than in reality and is thus underestimated. The determination of the minor order cost can be found in Table G.3 in Appendix G: Determination of the costs of the current situation.

The holding cost from the past year has been determined by adding up the holding cost for all the individual capacities and shared capacities in the buffer inventory. The holding cost for each individual capacity has been determined by multiplying the average inventory level with the holding cost per unit per week as found in Section 4.4. The result is subsequently multiplied with 52 (weeks) to find the holding cost per year. The average inventory level has been defined as half the capacity size as we assume that the inventory goes from full capacity to zero capacity. The holding cost for each shared capacity has been determined in a similar way by multiplying the average inventory level with the pallet dimensions in m^2 and with the holding cost rate per m^2 per week. The result is then multiplied with 52 (weeks) to find the holding cost per year. Note that we did not include the holding cost of the overflow capacity in the calculation of the holding cost from the past year. Since we have no data on the usage and contents of the overflow capacity, it is hard to determine the holding cost of the past year for the overflow capacity. Therefore, the overflow capacity has been omitted from the calculation. This means that the determined value for the holding cost per year will be lower than in reality, so we underestimate the holding cost. The determination of the holding cost can be found in Table G.4 and Table G.5 in Appendix G: Determination of the costs of the current situation.

Regarding the calculation of the EC of the new situation, we can immediately derive the total expected costs per year from Section 5.1. We saw there that the total expected costs of the optimal replenishment schedule (which will be used in the new situation) are equal to €413.89 per week. We can simply multiply this value with 52 (weeks) to find the total expected costs per year. However, we want to split the total expected costs per year in minor order cost, major order cost and holding cost per year, just like in the current situation. To this end, we need to look back at the objective function of mathematical programming Problem 4.9. This objective function is minimised by the inventory control system as it gives the value for EC . The objective function consists of a major order cost term, a minor order cost term and a holding cost term. We can see that the major order cost term is equal to $\frac{K}{R_0}$. Since we know the value for K (input to the inventory model) and the value for R_0 (2.0, since that is the R_0 belonging to the optimal replenishment schedule), we can easily calculate the major order cost term. We only need to multiply the result by 52 (weeks) to find the major order cost per year.

The minor order cost term is equal to $\sum_{i=1}^m \frac{k_i}{R_i}$. Since we know the k_i for every of the controlled m SKU's (input to the inventory model) and the value for R_i for every of the controlled m SKU's (output of the optimal replenishment schedule), we can easily calculate this summation to find the value for the minor order cost term. Again, we need to multiply this value by 52 (weeks) to find the minor order cost per year.

The holding cost term can now be calculated quite easily since the remainder of the objective

function corresponds to the holding cost term. We can simply subtract the major and minor order cost per year from the total expected costs per year to find the holding cost per year.

The resulting total expected costs per year for both the current situation and the new situation are given in Table 5.6. The total expected costs are in both cases split in minor order cost, major order cost and holding cost per year. Furthermore, the percentages increase or decrease in costs of the new situation with respect to the current situation are given in the table. Additionally, pie charts of the current situation and the new situation are given in Figure 5.8. These pie charts give a graphical representation of the share of each cost component in the total expected costs.

Table 5.6 This table shows the minor order cost, major order cost, holding cost and total expected costs per year for both the current situation and the new situation. Besides, the percentage increase or decrease in cost of the new situation with respect to the current situation has been given for each cost.

| | Current situation | New situation | Percentage in/decrease |
|--------------------------------------|-------------------|-------------------|------------------------|
| Minor order cost per year | € 4274.00 | € 5621.57 | 31.53% |
| Major order cost per year | € 9080.00 | € 1754.22 | -80.68% |
| Holding cost per year | € 10138.57 | € 14146.57 | 39.53% |
| Total expected costs per year | € 23492.57 | € 21522.36 | -8.39% |




Figure 5.8 These pie charts show the shares of the minor order cost, major order cost and holding cost in the total expected costs for both the current situation and the new situation.

The first and most important thing that becomes clear from Table 5.6 is that there will be an decrease in total expected costs per year if the optimal schedule from the inventory control system will be used. This means that the situation which will arise then, will be an improvement over the current situation. According to our calculations, the decrease in total expected costs per year will be equal to 8.39%. However, there were some factors that affected the accuracy of our determination of the total expected costs of the current situation. Some of the factors overestimated the expected costs and some of the factors underestimated the expected costs. Nevertheless, we have a feeling that the factors that underestimated the expected costs outweigh the factors that overestimated the expected costs. The real expected costs of the current situation are therefore probably higher than our determined value. This means that the gap with the expected costs of the new situation will probably be higher. The decrease in expected costs will therefore be at least 8.39%. All in all we can say that the shown improvement in expected costs is a conservative estimate.

Furthermore, we can reason how the improvement in expected costs will be realised by looking at Table 5.6 and Figure 5.8. We see that the major order cost decreases sharply with more than 80% and its share in the total expected costs becomes only 8%. This comes at the expense of higher minor order cost and holding cost, which increase with 31.53% and 39.53% respectively. However, the decrease in major order cost is larger than the increase in both minor order and holding cost together, which means that the total expected costs decrease. The increase in holding cost and the decrease in major order cost can be explained if we consider the number of replenishment orders that will be placed. We know from the current situation that multiple replenishments take place per week. We can also confirm this by looking at the replenishment data of the past year, where we counted 116 replenishments which contained SKU's which are studied in this project (2-3 replenishments per week). However, in the new situation, a replenishment order will be placed only once in two weeks according to the optimal replenishment schedule. This means only $52 / 2 = 26$ replenishments per year. This large decrease in number of replenishments per year causes the large decrease in major order cost per year, as major order cost is incurred at every replenishment order. It also explains why the holding cost per year has increased so much: you order less often but when you order, you will order more. This means that you keep more inventory on hand which causes higher holding cost. However, the increase in minor order cost per year cannot be explained by the decrease in number of replenishments per year. We do not have an explanation for the higher minor order cost per year in the new situation.

Finally, we want to make a few remarks regarding the distribution of holding cost and order cost (as a whole, so major + minor order cost) over the total expected costs. When we compare the pie charts in Figure 5.8, we see that the distribution holding cost/order cost is 43%/57% in the current situation and 66%/34% in the new situation. A reason why the share of the holding cost becomes so large compared with the share of the order cost for the new situation has been given in the previous paragraph. Namely, the number of replenishments per year decreases, so less orders are placed and more inventory is held. However, we can give another reason: our new situation is focused on satisfying a high service level. Achieving a high service level requires holding much inventory on hand, so that backorders will occur only rarely. Therefore, the holding cost will play a major role and it will have the largest share in the total expected costs. This is due to the stochastic nature of the demand in our project. If demand would have been deterministic, service level would not play a role since demand satisfaction would be guaranteed. The holding cost and order cost would in that case have a perfect 50%/50% balance. In our case, we hold a lot of inventory to aim at a high service level (complete demand satisfaction cannot be guaranteed). Consequently, there will not be a perfect balance: there will be more holding cost, 66% in our case.

5.3.2. Other improvements

Next to the improvement in expected costs, we expect other improvements when the inventory control system will be implemented. However, these improvements are not measurable, like our KPI expected costs. Nevertheless, we believe that these improvements will actually exist. We will now list and explain these other improvements:

- Achieving desired service levels

In the current situation, the company does not know what its current service levels are and are therefore not sure if the desired service levels are achieved. However, it is safe to assume that the desired service levels are not achieved currently as it is a problem in the problem cluster of Figure 1.1 (Required part(s) not available in buffer). Since the service levels are constraints in the inventory model, the inventory control system will take the desired service levels into account when creating the optimal replenishment schedule. The optimal

replenishment schedule will tell the company what to do to achieve the desired service levels and shows the real service levels that will be achieved in that case. Hence, the company gets insight into the service levels and can be more sure that they are achieved. This can be regarded as an improvement over the current situation.

- Setting different service levels

In addition to the previous point, the company has the possibility to set service levels in the inventory control system. Currently, the company does not set service levels and it is therefore difficult to strive for high service levels. The inventory control system will facilitate this. It is also very flexible since the company can change service levels in the future, if they change their strategy for example. The company can even set different service levels for different SKU's, if it appears that some SKU's give more problems than others.

- Less emergency deliveries

Currently, emergency deliveries exist when there is a stockout of a SKU and this SKU is needed immediately. These emergency deliveries cost time and money, but also cause frustration among employees. Following the optimal replenishment schedule created by the inventory control system will reduce the number of stockouts. Therefore, the number of emergency deliveries will be reduced, which is an improvement.

- Standardisation of the replenishment process

Using the inventory control system in the replenishment process will give a lot more structure. Currently, the replenishment process is not structured: there is no systematic way of inventory control, the replenishments are based on the feeling and experience of employees. The inventory control system takes away these issues since it is a systematic way of inventory control. It does not rely on the feeling and experience of employees anymore. This is an improvement over the current situation. Namely, anyone from the company can do the replenishments now by using the inventory control system in the user-friendly Excel file and the manual that accompanies it. The advantage is thus that the company is less dependent on the employee who is responsible for the replenishments. The replenishments can still happen without problems when the employee is sick, on holiday or leaving the company. Furthermore, more structure probably means that the employee will make less errors and will also make sure that multiple employees carry out the replenishments in the same way. Therefore, standardisation of the replenishment process will certainly be an improvement.

- Extending the inventory control system to different product series

The inventory control system of this project has been made for one assembly line (2 product series) only. However, it is possible to extend the inventory control system to assembly lines of different product series. These assembly lines should be set up in the same way as the assembly line in this project. Furthermore, all input data that is required for the controlled SKU's in this project should also be gathered for the SKU's of the other product series. The improvement when extending the inventory control system to different product series is obvious: the improvements that were discussed in this section will probably also apply to those product series. Besides, the replenishment orders for the different product series can be grouped together when the inventory control system will be used for multiple product series at the same time. This is possibly an improvement over placing replenishment orders for the different product series separately.

6. Conclusions and Recommendations

The final chapter of this project covers the conclusions and recommendations that we made based on the previous chapters. In Section 6.1, we will draw conclusions on what we did and how we did our research. Furthermore, we will list the findings and the managerial insights in this section. Recommendations to the company based on our research will be given in Section 6.2. This section will also contain topics for future research. Finally, we will describe the contribution that we made with our research in Section 6.3. This section is subdivided in Contribution to theory and Contribution to practice.

6.1. Conclusions

This project was started since the company indicated that the replenishment of parts to the buffer inventories of the assembly lines is currently not efficient. Given the limited time that is available for this project, it was decided to look at the replenishment of one buffer inventory only. The main problems that the company faces regarding the replenishment process are that the costs are high and required parts are not always available for assembly. However, when looking for the cause of these problems, we found out that the core problem is the lack of a systematic way of inventory control. An inventory control system is needed for a systematic way of inventory control. Therefore, we formulated the following action problem: "There is no inventory control system, whereas there should be an appropriate inventory control system to base replenishment decisions on." To solve this action problem, the following central research question should be answered: *What is an appropriate inventory control system for the buffer of the A/B series assembly line?*

The first thing that we did was analysing the current situation in Chapter 2, which is important to understand the context in which the inventory control system is going to work. Furthermore, the scope of the inventory control system can be determined further and the part characteristics that the inventory control system should take into account can be investigated.

The current situation is analysed in this chapter by answering two knowledge questions. The first knowledge question that was answered is: *What is the current situation regarding the planning and control of the A/B series assembly line and corresponding buffer inventory?* We found out that the planning takes place according to a week planning in which they plan which orders they will build in that week. We also found out that the parts used in the assembly line, are stored at external warehouses and at the production location. The inventory at the production location can be subdivided into bulk inventory (used to replenish the buffer inventory) and buffer inventory (stored close to the assembly line and used to supply it). The replenishments which we study take place from an external warehouse inventory or bulk inventory to the buffer inventory. These replenishments are made by employees of the assembly line. They should in principle follow a FIFO policy, where they use the lot with the oldest lot date for replenishment.

The second knowledge question that was answered is: *Which parts of the A/B series should be included in the inventory control system and what are their characteristics?* We determined which parts to include in the inventory control system (critical parts) and which not. It appeared that a basic distinction could be made between C parts and L parts. The C parts could be further subdivided in different part families. Some part families are considered relevant and all the parts from these families are chosen as critical parts. An ABC inventory classification method has been used to determine the critical L parts. Only the L parts classified as class A and B will be critical parts. In the end, we determined a total of 90 critical parts. For these 90 parts, we determined values for the following characteristics: part commonality, storage locations, lead times, lot sizes and buffer capacity.

The next step towards answering our research question is conducting literature study. This has been done in Chapter 3. We learned that an inventory control system uses an inventory model for determining replenishment decisions. Besides, a demand model is needed as input to the inventory model. Therefore, our literature study focused on both demand models and inventory models. We first conducted literature study on different demand models. The aim was to answer the following knowledge question: *“What demand models are proposed in inventory management literature?”*. It appeared that a basic distinction could be made between deterministic and stochastic demand models. For deterministic demand models, a further distinction could be made between constant demand and time-varying demand. For stochastic demand models, we only looked at demand models that assume approximately constant average demand. Stochastic demand models are commonly modelled by a probability distribution. We looked at both discrete and continuous demand distributions and at a decision rule to decide between the two. The discrete distributions that we discussed were the Poisson and Compound Poisson distribution. We also discussed a decision rule to decide between them. Finally, we discussed the continuous Normal and Gamma distribution and a decision rule to decide between them.

Subsequently, we conducted literature study on different inventory models. The knowledge question that was answered is: *What inventory models are used in inventory control systems according to literature?* We only considered inventory models for single-echelon inventory control systems. These inventory models were divided in models with individual control of single items and models with coordinated control of multi items. Since the demand model is an important input to the inventory model, the classification of the demand models has also been used to classify the inventory models. This means that both the single item inventory models and the coordinated multi-item inventory models were subdivided in deterministic and stochastic models. The deterministic inventory models were further subdivided in models that assume constant demand and time-varying demand. The stochastic inventory models were further subdivided in continuous review and periodic review models. For each class of this classification, inventory models were found in literature, which are explained in Chapter 3.

With the knowledge we gained from the current situation analysis of Chapter 2 and the literature study of Chapter 3, we could answer our central research question. This is done in Chapter 4. In the first section, the foundation of an appropriate inventory control system has been chosen by choosing the most appropriate demand model and inventory model.

It is chosen to use a stochastic demand model which uses a discrete probability distribution. From decision rules from literature, it appeared that the logarithmic compound Poisson distribution could best be used to model the demand of 90% of the critical parts. The demand of the remaining 10% of the critical parts should be modelled by a Poisson distribution. The demand model of the inventory control system should therefore consist of the Poisson distribution and the logarithmic compound Poisson distribution.

Subsequently, the most appropriate inventory model is chosen. We consider only stochastic coordinated multi-item inventory models, since we chose a stochastic demand model. Furthermore, we only consider periodic review models, since periodic review of the inventory is more practical for the company than continuous review and the studied periodic review models generally perform equally well or even better. We can thus conclude that the most appropriate inventory model is a periodic review stochastic coordinated multi-item inventory model. Three of these inventory models were discussed in the literature study, we chose for the inventory model proposed by Fung et al. (2001). This is a periodic review coordinated inventory model that makes use of order-up-to-level policies and a service measure. The model assumes compound Poisson demand. However, the model is not implemented as it is in our inventory control system, we needed to make changes to their

model. Firstly, we have incorporated the “normal” Poisson distribution from our demand model in their inventory model. Secondly, it turned out that one of their algorithms did not give the global minimum of a certain range of solutions, but merely a local minimum. Therefore, we adjusted that algorithm so that it finds the global minimum of a specified range and used the adjusted version in our inventory control system.

The developed inventory control system basically works as follows: several demand parameters (such as the mean and variance), cost parameters (major order cost, minor order cost, holding cost), lead times and service levels are inputted to the inventory control system. Besides, a range in which the base period R_0 (time between consecutive replenishments) should fall, should be given. We needed to determine values for these parameters, where some parameter values could be easily derived while some parameter values required more calculations. Based on the demand parameters, the demand model models the demand of the critical parts as Poisson demand or logarithmic compound Poisson demand. Using this input, together with the lead times, service levels, cost parameters and the range for the base period, the inventory model finds the values for the decision variables R_0 , n_i ($R_i = n_i * R_0$) and S_i that minimise the total expected costs while satisfying the service levels. Here, the R_i denotes the time between replenishments containing SKU i and S_i denotes the order-up-to-level of SKU i . Since the inventory model is a non-linear mixed-integer programming problem, it is too difficult to solve it exactly. That is why two heuristic algorithms are used to solve the problem. One of these algorithms has been adjusted, so that it always finds the global minimum of a given range of the base period. Please check Section 4.2 for the details.

At this point, the (R_i, S_i) parameters of every critical part are found. However, the concrete replenishment decisions are still not determined with merely these parameters. With replenishment decisions we mean *when* (date) and *how much* (quantity) should be ordered, for every critical part. Therefore, we made an extension to the inventory control system that automates the derivation of the replenishment decisions from the (R_i, S_i) parameters using the current inventory levels. Furthermore, an extension that does a capacity check on the buffer inventory is added to the inventory control system. This extension is developed since the inventory model does not take capacity into account, although the buffer inventory has a limited capacity. It is therefore a good idea to check if the buffer inventory capacity is large enough to store the order-up-to-levels that result from the inventory model.

Finally, an extension that checks if the replenishment decisions lead to stockouts, is developed and added to the inventory control system. The inventory control system uses stochastic demand to find the replenishment decisions. However, there is short term deterministic demand information available from outstanding orders. This extension uses this information to predict if and where stockouts will occur if the replenishment decisions are followed.

We made the implementation of our proposed inventory control system in Excel. The models behind the inventory control system are programmed in VBA. The worksheets act as a user interface, used for giving input and receiving output. Besides, a manual has been made for the company, which gives a very detailed explanation on how to use the user interface.

The Excel tool determines a long-term replenishment strategy once: a list with replenishment dates and the SKU's which should be ordered at these dates, together with their order-up-to-levels. Furthermore, it derives the concrete replenishment decisions at every replenishment date. Please check Section 4.3 for the details.

Our findings are presented in Chapter 5 including sensitivity analysis and the improvement if the proposed inventory control system is used. It appeared that the optimal replenishment schedule is the one with a base period of 2 working weeks. This means that it is optimal to do a replenishment

once every 2 working weeks (10 working days). The (R_i, S_i) parameters belonging to this optimal replenishment schedule were also returned as output, with all R_i 's being integer multiples of 2 working weeks to ensure coordinated multi-item replenishments. When we replenish according to these (R_i, S_i) parameters for all critical parts, the total expected costs per week will be equal to €413.89 and a mean service level of 98.364% will be achieved. All individual SKU's will achieve a service level above 98%, which is in line with our 98% service level constraint. Regarding the results of the buffer inventory capacity check, we found that for some capacity parts there are pallet places remaining and for some capacity parts there are pallet places short. However, only a part of the capacity shortage can be compensated by shifting the sizes of the capacity parts, as there are far more pallet places short than remaining.

Subsequently, we performed sensitivity analyses on different parameters to determine which parameter influences the output of the inventory control system in what degree. Firstly, we varied the values of the major order cost, minor order cost and holding cost one by one and compared the results. We concluded that all three costs are almost perfectly positively correlated with the expected costs. However, they do not have the same influence on the expected costs: a change in holding cost has the largest influence. Therefore, the inventory control system is the most sensitive to a misestimation of the holding cost. We also looked at the influence of the holding cost on the optimal base period. It turned out that the optimum is reached at a shorter base period when the holding cost increases and at a longer base period when the holding cost decreases. Secondly, we performed a sensitivity analysis with changing service level. We found out that the service level and expected costs are positively correlated, but the correlation is not linear: the expected costs increase more quickly with each equal step higher in service level. We also looked at the influence of the service level on the optimal base period, but it appeared that this influence was only small. Finally, a sensitivity analysis with changing buffer inventory capacity size was performed. The buffer inventory capacity was divided in three parts: individual capacities, shared capacities and an overflow capacity. We analysed the effect on the net number of pallets short by changing one of these capacity parts at a time. It appeared that there exists a quite linear negative correlation between the net number of pallets short and the capacity size, for each capacity part. However, the influence that each capacity part has on the net number of pallets short is different: a change in size of the overflow capacity has the least influence whereas the other two capacity parts have a larger (but approximately the same) influence.

The goal of our inventory model is to minimise the expected costs subject to a service level constraint. Our KPI is therefore the expected costs. Hence, we have focused on the improvement in expected costs first. This improvement was measured by calculating the expected costs of the current situation (costs of the past year) and of the new situation (costs with the inventory control system) and comparing them. It became clear that the new situation will be an improvement over the current situation. Namely, the decrease in expected costs will be at least 8.39% according to our calculations. This is probably a conservative estimate, since the calculations underestimated part of the expected cost of the current situation. The decrease in expected costs has been realised by a sharp decrease in major order cost. The minor order cost and holding cost increase, but the decrease in major order cost is larger than both increases together. We concluded that the decrease in major order cost and increase in holding cost are due to a large decrease in number of replenishments per year. In the new situation there will only be 26 replenishments per year, whereas we counted 116 replenishments containing SKU's of this project in the past year. As a result, the subdivision of the total expected costs in holding cost and order cost changed from 43% and 57% in the current situation to 66% and 34% in the new situation, respectively.

Next to the improvement in our KPI expected costs, we discussed other improvements which are

expected when the inventory control system will be implemented. These other improvements are not measured but it is believed that they will exist. First of all, the company gets insight into the service levels and can be more sure that they are achieved. Secondly, the company can set different service levels for different SKU's and can change them in the future. Furthermore, the number of emergency deliveries will be reduced, which saves time, money and frustration. Besides, the inventory control system gives a lot more structure to the replenishment process. By standardising the replenishment process, the company becomes less dependent on the (experience of) employees who are responsible for the replenishments. Finally, it is possible to extend the inventory control system to assembly lines of different product series. All improvements that we discussed here will probably also apply to those assembly lines. Further improvement could possibly be achieved by grouping the replenishment orders of all assembly lines together instead of placing these orders at separate times.

Findings

Here our main findings are listed below:

- Base period, so the company will know how often replenishment orders should be placed
- Review periods for all critical parts, so the company will know *when* to order *which* SKU's
- Order-up-to-levels for all critical parts, so the company can determine *how much* of the SKU's should be ordered
- Service levels for all critical parts, so the company will know which service levels will be achieved
- Major order cost per week, minor order cost per week, holding cost per week and the total of the three: the expected costs of inventory control per week and the improvement over the current costs
- Current capacity of the buffer inventory and the increase in capacity that is needed to store the found order-up-to-levels

Managerial insights

The gained managerial insights are listed below:

- The total costs of buffer inventory control are expected to be at least 8.39% lower when the proposed inventory control system is used to control the buffer inventory of the A/B series
- This decrease in costs is mainly achieved by reducing the number of replenishments from multiple times per week to once every two weeks
- Some critical parts are replenished every two weeks whereas the other critical parts are replenished after a multiple of two weeks (e.g. every four, six or eight weeks)
- From the buffer inventory capacity check we learn that the capacity of the buffer inventory is not large enough to store the order-up-to-levels of all critical parts at once
- From sensitivity analysis it becomes clear that a change in holding cost has a larger influence on the expected costs than a change in major or minor order cost
- From sensitivity analysis it also becomes clear that the expected costs increase more quickly with each equal step higher in service level
- It is possible to apply the inventory control system to more assembly lines, in order to solve the same problems and probably also cut inventory control cost there

6.2. Recommendations and Future research

During the project we came up with advice for the company on certain issues. Furthermore, we have made recommendations to the company based on the results of our research. These

recommendations are listed and explained below. Some of the recommendations lead to ideas for future research, but we also came up with other topics for future research. These topics for future research are also listed and explained below.

Recommendations

- Use the Excel tool with the implemented inventory control system to make the long term replenishment schedule. Follow this replenishment schedule: place a replenishment order on the indicated replenishment dates.
- Use the Excel tool with the implemented inventory control system to calculate the replenishment decisions on every replenishment day given by the replenishment schedule. Follow the replenishment decisions: the given required order quantities of each SKU should be ordered. Preferably, order the advised lots to satisfy the required order quantities. In this way, the FIFO policy will be respected.
- Make a buffer inventory location for every part of the A/B series, such that the inventory control system can calculate buffer inventory levels for all parts. Currently, the TR parts of the B Large series do not have a buffer inventory location assigned to them. Therefore, the inventory control system assumes that their buffer inventory level is always zero.
- Update the estimation of the demand and cost parameters regularly and change them in the inventory control system. Better input to the inventory control system will give better output.
- The capacity check of the inventory control system indicates that the buffer inventory capacity should be enlarged. However, the question is in which degree this really is the case in practice. Namely, the capacity check assumes that the maximum possible on hand inventory of every SKU is present at the same time in the buffer inventory. This will likely never be the case in reality. We therefore recommend to first observe in practice what will happen with the buffer inventory utilization when the inventory control system is used. If it turns out that the capacity is too small, then (partly) enlarge the capacity as indicated by the capacity check results.
- Do not store SKU's on pallets anymore on pick locations in the buffer inventory. Instead, place a decking on the pallet racks and place the SKU's thereupon. This has two advantages. Firstly, there will be more space available on the pallet racking, since the area increases (there will not be any gaps between pallets anymore). Secondly, it is no longer required that precisely one or two pallets of a SKU are stored in the pallet racking, but there could also be the contents of e.g. half a pallet or one and a half pallet stored in the racking. The pick locations become more flexible in this way.
- Extend the inventory control system with a calculation of the amount of truck loading meters that a replenishment requires. In this way, we can already predict how many truck trailers are needed for a replenishment and schedule the right number of trucks.
- Extend the inventory control system to multiple assembly lines. To make those assembly lines compatible with the inventory control system, they should be set up the same as the assembly line of the A/B series. This means that they should have a production file and a buffer inventory. There are two possibilities: the first possibility is to use the inventory control system for every assembly line separately. You will then get an optimal replenishment schedule for every assembly line. Some replenishments for different assembly lines might overlap and can in that case be combined into one replenishment. The other possibility is to use the inventory control system for every assembly line at the same time. You will then get only one optimal replenishment schedule for all assembly lines. Every

replenishment will then consist of a mix of SKU's from all assembly lines. Research is needed to find out which possibility will be better, especially in terms of expected costs.

- There is a problem when a new SKU is introduced and the inventory of this SKU should be controlled by the inventory control system. The problem is that this SKU has no historic demand, so the demand parameters cannot be determined. We recommend to do one of the following two things to cope with this problem. The first option is to request historic demand information from the manufacturer/supplier and determine the demand parameters based on that information. The second option is not including the SKU in the inventory control system in the first place and replenishing it manually (like is done currently with all SKU's). Keep track of the demand of the SKU and as soon as some weeks with demand have past, reasonable estimates of the average and variance of the week demand can be made. These values should be updated more frequently than the values of the SKU's with a long demand history. Another way how the inventory of the new part is controlled is via the outstanding order check. If the part demand of the outstanding orders is larger than the estimate based on the average and variance, then the check will indicate that more units should be ordered than just the replenishment schedule says. However, to make this check work, it is important to at least include initial demand parameters of the part in the inventory control system.

Future research

- Research on production scheduling can be carried out. During the identification of the core problem, we determined that the lack of a production schedule is a problem. However, it was not chosen as a core problem since a production schedule was not needed to devise an inventory control system. Nevertheless, we made an extension (outstanding orders check) that predicts at a replenishment moment whether stock outs will occur before the next replenishment moment. This extension requires a production schedule of the outstanding orders for the coming two weeks, since that is the time between replenishments. Therefore, we recommend to look into the production scheduling of outstanding orders for (at least) two weeks in advance.
- As described in the recommendations, future research could focus on the extension of the inventory control system to multiple assembly lines. Both possibilities described there should be developed and the best one should be chosen.
- Look into the forecasting of demand. Currently, the demand is modelled based on parameters derived from historic demand. However, the future demand might change and then the used demand parameters will not be accurate enough. It is therefore useful to forecast this future demand, so better values for the demand parameters will be found. Another advantage is that we can take the seasonality of demand into account when forecasting the demand.
- If a production schedule for a long time period can be made in the future, it might be worth considering an inventory control system solely based on deterministic demand. This will probably decrease the risk of stockouts even more. We propose to investigate inventory models using a rolling horizon with frozen period (see Chapter 3) in that case.
- If the company also wants to have an inventory control system for the bulk inventory and/or external warehouses in the future, it can be a good idea to design a multi echelon inventory control system. This will be preferred to designing multiple single echelon inventory control systems, as it is generally better to minimise the inventory control costs of the whole system (across all echelons) at once instead of just minimising the costs of all levels separately.

Namely, the total inventory costs of the whole system are expected to be lower under a multi echelon inventory control system.

6.3. Contribution to theory and practice

Our research has contributions to both theory and practice. A contribution to theory is a contribution that is made to the scientific literature. We used an inventory model from literature in our project, but we did not implement it in the inventory control system exactly in the way it was described in literature. Namely, we changed and added certain aspects of the inventory model to make it a better fit to our situation. Furthermore, the inventory control system included specifically made extensions in addition to the inventory model. These contributions to theory are listed and explained below. Next to the contributions to theory, our research has contributions to practice. A contribution to practice is a contribution that is made to the company. The outcome of our inventory control system is valuable information to the company, which they can use to improve their replenishments. Furthermore, the usage of our inventory control system contributes to other aspects of the replenishment process. These contributions to practice are also listed and explained below.

Contribution to theory

- Let the algorithm always choose the global minimum expected costs of the given range
This is a change made to the used inventory model from literature. This model contains a HCTSP algorithm which stops searching for the minimum expected costs when the expected costs rise for the first time. However, when the graph of the expected costs is non-convex (what appeared to be the case in our research), the global minimum of the given range might not be found. We adjusted the search procedure of the HCTSP algorithm so that it always finds the global minimum of the given range: our algorithm does a full enumeration of the solutions in the specified range, whereas the original algorithm has a stopping criteria, even when a range is specified. However, this adjustment is at the expense of computational effort, so it takes more time to find the solution, especially if the specified range is large. We cope with this by assuming only reasonable, practical values as a specified range, e.g. 1, 2, 3 days and not values like 1.06, 2.32, 3.78 days.
- Addition of Poisson demand next to the compound Poisson demand
Here we combined the theory from demand models and inventory models. From the theory about demand models we learned that we need “normal” (unit-sized) Poisson demand for some SKU’s and compound Poisson demand for the other SKU’s. However, the inventory model only assumes the latter demand. Therefore, we incorporated the former demand in the model, which was easy since Poisson demand is a special case of compound Poisson demand. As a result, we do not have one distribution as input for the inventory model anymore, but a mix of two distributions and a decision rule to decide between them for every SKU.
- Extension which does a capacity check on the buffer inventory
This extension has been made to include a capacity dimension, since the inventory model does not take the capacity of the buffer inventory into account. It was a conscious choice to include the capacity dimension as a check and not as a constraint. This has been done since we did not want the capacity to be the bottleneck, it should not influence the results of the inventory model. The size of the buffer inventory capacity should be adjusted to the results of the inventory model and not the other way round. However, it should also be possible to include a capacity dimension as a constraint, when capacity is a real bottleneck.

Contribution to practice

- Concrete replenishment decisions
We made an extension to the inventory control system that automates the derivation of the replenishment decisions from the (R_i, S_i) parameters using the current inventory levels. The company can use these concrete replenishment decisions immediately to place a replenishment order.
- Achieving desired service levels at minimum costs
By replenishing according to the replenishment decisions, the company will replenish in the best way possible. This means that they will achieve the least costs of inventory control possible, while still satisfying the desired service levels for every controlled SKU.
- Structured and standardised replenishment process
The implementation of the inventory control system will lead to more structure and a standardisation of the replenishment process. Therefore, every employee will be able to place replenishments by using the inventory control system.
- Adhering to the FIFO policy when replenishing the advised lots
When the company orders the advised lots to satisfy the required order quantities of the replenishment decisions, it will adhere to the FIFO policy that should in principle be followed when replenishing. This is the case since the advised lots are the lots with the oldest lot dates.
- Future-proof system for new warehouses and/or changing parts
The (externa) warehouses where inventory of the controlled SKU's is stored, can change over time. The same goes for the SKU's themselves: when a product series is updated or when a new product is released, new SKU's need to be controlled or existing SKU's might be replaced. These changes can all be handled by the inventory control system. Hence, the inventory control system provides a future-proof system for inventory control.
- Extendable to different product series
The inventory control system is made for three product series from the same assembly line. However, it can be extended to different product series of other assembly lines, if they are set up in the same way. Therefore, the inventory control system provides in principle a complete solution for inventory control to the company.
- Extendable with product scheduling
There is information on outstanding orders available, but this is currently not used to make a production schedule. However, the inventory control system contains an extension in which a production schedule of outstanding orders can be used to predict stockouts. Therefore, the inventory control system provides a use for a production schedule of outstanding orders, if the company decides to make such a schedule in the future.

References

- Atkins, D. R., & Iyogun, P. O. (1988). Periodic versus 'Can-Order' Policies for Coordinated Multi-Item Inventory Systems. *Management Science*, 34(6), 791–796.
- Axsäter, S. (2006). *Inventory Control* (Second). Springer.
- Boctor, F. F., Laporte, G., & Renaud, J. (2004). Models and algorithms for the dynamic-demand joint replenishment problem. *International Journal of Production Research*, 42(13), 2667–2678. <https://doi.org/10.1080/00207540410001671660>
- Fung, R. Y. K., Ma, X., & Lau, H. C. W. (2001). (T,S) policy for coordinated inventory replenishment systems under compound Poisson demands. *Production Planning & Control*, 12(6), 575–583. <https://doi.org/10.1080/095372801750397699>
- He, J., Rexford, J., & Chiang, M. (2010). Design for Optimizability: Traffic Management for a Future Internet. In *Algorithms for Next Generation Networks* (pp. 3–18).
- Heerkens, H., & Van Winden, A. (2012). *Geen Probleem: Een aanpak voor alle bedrijfskundige vragen en mysteries*. Van Winden Communicatie Nieuwegein.
- Jensen, P. A., & Bard, J. F. (2002). *Operations Research Models and Methods*.
- Praveen, M. P., Simha, J. B., & Venkataram, R. (2016). Techniques for Inventory Classification: A Review. *International Journal for Research in Applied Science & Engineering Technology*, Vol. 4(X), 508–518. Retrieved from www.ijraset.com
- Sahin, F., Narayanan, A., & Robinson, E. P. (2013). Rolling horizon planning in supply chains : review , implications and directions for future research. *International Journal of Production Research*, 7543(May). <https://doi.org/10.1080/00207543.2013.775523>
- Silver, E. A., Pyke, D. F., & Thomas, D. J. (2017). *Inventory and Production Management in Supply Chains* (Fourth). CRC Press.
- Slack, N., Brandon-Jones, A., & Johnston, R. (2016). *Operations Management* (Eighth). Pearson.
- Viswanathan, S. (1997). Periodic Review (s, S) Policies for Joint Replenishment Inventory Systems. *Management Science*, 43(10), 1447–1454.

Appendix A: Critical parts and values of their relevant characteristics

Table A.1 CM SKU's with their descriptions and the type of machine which uses the TR parts

Confidential

Table A.2 CL SKU's with their descriptions and the quantities in which they were used last year

Confidential

Table A.3 CD numbers with their descriptions

Confidential

Table A.4 L SKU's classified as class A parts

Confidential

Table A.5 L SKU's classified as class B parts

Confidential

Table A.6 Values of some relevant characteristics of the critical parts

Confidential

Table A.7 Buffer capacity split in individual capacities, shared capacities and overflow capacity

Confidential

Appendix B: Goodness of fit tests

Confidential

Figure B.1 Goodness of fit test of first SKU

Confidential

Figure B.2 Goodness of fit test of second SKU

Confidential

Figure B.3 Goodness of fit test of third SKU

Confidential

Figure B.4 Goodness of fit test of fourth SKU

Confidential

Figure B.5 Goodness of fit test of fifth SKU

Appendix C: Values of inventory control system parameters

Mean and variance of the week demand

Table C.1 This table shows the estimates of the mean and variance of the week demand for every SKU.

Confidential

Major order cost determination

Table C.2 This table shows the possible combinations of trailer movements together with their transportation costs and relative frequencies. The major order cost K is then calculated as the weighted average transportation cost with the relative frequencies as weights.

Confidential

Table C.3 This table shows the transportation costs of all trailer movements. This information is used in Table C.2 to determine the transportation costs of the trailer movements combinations.

Confidential

Table C.4 This table shows which combination of trailer movements occurred on which day from the past year. This information is used in Table C.2 to determine the relative frequencies of the trailer movements combinations.

Confidential

Minor order cost determination

Table C.5 This table shows the estimation of the handling cost per pallet. On the one hand, information from Table C.6 is used here to determine the average pallet dimensions. On the other hand, the handling cost per pallet is used in Table C.6 to determine the minor order cost k_i of every SKU i .

Confidential

Table C.6 This table shows the values that are directly or indirectly required for the determination of the minor order cost k_i of every SKU i . The values for k_i are eventually given in the last column for every SKU i .

Confidential

Holding cost determination

Table C.7 This table shows the holding cost rate per m^2 per year for C SKU's as defined by the company and the conversion to the desired form (holding cost rate per m^2 per week). The latter is used in Table C.8 to determine the holding cost h_i of every C SKU.

Confidential

Table C.8 This table shows the determination of the amount of square metres that one unit of every C SKU occupies, as defined by the company. Using this information and the C rate in Table C.7, the holding cost rate h_i of every C SKU is calculated in the last column.

Confidential

Table C.9 This table shows the holding cost rate per EUR-pallet per month for L SKU's, as defined by the company, and the conversion to the desired form (holding cost rate per m^2 per week). The latter is used in Table C.10 to determine the holding cost h_i of every L SKU.

Confidential

Table C.10 This table shows the determination of the amount of square metres that one unit of every L SKU occupies. Using this information and the L rate in Table C.9, the holding cost rate h_i of every L SKU is calculated in the last column.

Confidential

Appendix D: Manual of the inventory control system

Confidential

Appendix E: VBA code of the inventory control system implementation in Excel

Confidential

Appendix F: Detailed results of the best replenishment schedule

Table F.1 This table shows, for each controlled SKU, the review period (R_i), order-up-to-level (S_i), service level ($RealP_i$) and expected costs (EC_i) belonging to the best replenishment schedule ($R_0 = 2.0$).

Confidential

Table F.2 The capacity check results of the individual capacities per SKU. The 74 controlled SKU's that are not mentioned in this table all have just the right amount of pallets in their individual capacities.

Confidential

Appendix G: Determination of the costs of the current situation

Major order cost of the past year

Table G.1 This table shows the possible combinations of trailer movements together with their transportation costs and number of occurrences in the past year. The number of occurrences is determined by counting the corresponding combination of trailer movements in Table G.2. The major order cost of the past year is then calculated as the sum of the products of cost and number of occurrences.

Confidential

Table G.2 This table shows which combination of trailer movements, containing the controlled SKU's, occurred on which day from the past year. This information is used in Table G.1 to determine the number of occurrences of each combination of trailer movements.

Confidential

Minor order cost of the past year

Table G.3 This table shows the determination of the minor order cost of the past year. It is calculated as the sum of the handling costs of the past year for every SKU. These handling costs are the product of the number of pallets handled in the past year and the handling cost per pallet.

Confidential

Holding cost of the past year

Table G.4 This table shows the determination of the holding cost of the past year. It is calculated as the sum of the holding costs per year of all the individual SKU's and the holding costs per year of all the shared capacities, which are calculated in Table G.5.

Confidential

Table G.5 This table shows the calculation of the holding cost per year for every shared capacity. This information is used in Table G.4 to find the holding cost of the past year.

Confidential