3 May 2017

Improving the overall customer service level

A CASE STUDY AT PHILIPS

JEZUITA, LENA

Contact

Student

Lena Jezuita
E-mail: lenaj137@gmail.com
Mobile: +31 (0) 616011840

Supervisor University of Twente

Matthieu van der Heijden
E-mail: m.c.vanderheijden@utwente.nl
Phone: +31 (0) 53 489 2852

Second Supervisor University of Twente

Leo van der Wegen
E-mail: l.l.m.vanderwegen@utwente.nl
Phone: +31 (0) 53 489 3501

Supervisor Philips

Melissa Kuipers
E-mail: melissa.kuipers@philips.com
Phone: +31 (0) 612505109

Abbreviations

BG – Business Group
CO – Commercial Organization
SC – Supply Center
RDC – Regional Distribution Center
OTTR – On Time to Request
CSL – Customer Service Level
KPI – Key Performance Indicator
OT – On Time
IF – In Full
YTD – Yield to Date
wMAPE – Weighted Mean Absolute Percentage Error
S&OP – Sales and Operations Planning
MAG – Material Article Group
AG – Article Group
PAG – Product Article Group
SKU – Stock-keeping Unit
FCFS – First come first serve
FTL – Full Truck Load
IGM – Integral Gross Margin
COGS – Cost of Goods Sold
WACC – Weighted Average Cost of Capital
KM – Key Module

Contents

Summary

The well-known international company Philips has the mission to improve the life of 3 billion people through their innovations. The driving factor for reaching this goal is to improve customer service while maintaining low inventories and low costs. The target customer service level for 2020 is to deliver 95% of all order lines on time and in full. The order line fill rate of the division of Philips that we study is 89% at the start of this research. Philips is interested in researching how to reach the 95% and still stay profitable. Therefore, the goal of this research is the following:

The goal of this research is to improve a part of Philips' supply chain in order to improve the customer service level in the most profitable way.

The scope of this research is to find ways for customer service level improvement for a certain range of products. We decide to only research how to improve product availability at the warehouses, since low product availability is the main reason for a low customer service level. The order line fill rate for product availability at the warehouse is called CSL-availability. The CSL-availability at the start of this research of the studied division of Philips is 91%. Comparing this to the 89% order line fill rate, we can conclude that only 2% of the order lines are not fulfilled due to other reasons than availability at the warehouse, whereas 9% cannot be fulfilled because of stock-outs at the warehouses.

The scope of this research is limited to two product groups, where one already has a 93% CSL-availability, but can improve the total CSL-availability by 1,5% and thus has a big influence on the total average customer service level. The other product group has a CSL-availability of 87%. This leaves more room for improvement, but due to the relative low amount of order lines for this product type only a 0,5% total CSL improvement can be achieved. Solutions for improvement that we find during this research however can probably also be implemented for other product groups.

The biggest challenge for Philips is to achieve a high service level at low cost, because customers expect highly customized products and fast delivery times. Variability in demand and low forecast accuracies are found to be the main reasons for stock-outs at the warehouses. We researched different ways of coping with demand variability and come to the conclusion that for this scope lead time reduction is the most viable solution. Possible solutions are: changing the mode of transportation from sea to air, late customization or a dual sourcing strategy where a percentage of the products are transported by air and the rest by sea. We built a simulation model to quantify the impact of lead time reduction on the CSL-availability year to date and performed a cost analysis. For the analysis, we chose one SKU per product group, since there is no more data available. Due to the wide variation of SKUs of product group A, the chosen SKU represents only 0,5% of all order lines and is therefore not representative for the whole product group. We recommend Philips to also study other SKUs before implementing a solution for this product group. The SKU of product group B represents 43% of all order lines. Therefore, our conclusions about this product group are much more valuable.

The results and conclusion of the CSL-availability year to date and cost analysis for each solution after running the simulation model are listed below. Note that the recommendations are based on the results for only one SKU per product group. Therefore, we recommend to test the settings for other SKUs of those two product groups before implementation.

- Changing the transportation mode from sea to air reduces the lead time by six weeks for product group A. The improvement in CSL-availability YTD is 3,1% for two weeks of safety stock and 2,6% for one week of safety stock. The profit increase due to less inventory costs, less lost sales and more sales due to less cancellations is 5,3% for two weeks of safety stock and 7,1% for one week of safety stock. We also researched a dual sourcing strategy, where products are partly shipped via sea transport and partly via air transport. However, the most profitable

solution is to ship everything via air. Therefore, we can recommend Philips to implement this solution.

- Building a production line at the production site in North America decreases the lead time and improves the CSL-availability with 3,1%. However, the benefits are clearly not enough to cover the costs of a new production line. Therefore, we recommend Philips not to build a new production line.
- Changing the transportation mode from sea to air reduces the lead time by four weeks for product group B. Here, we found that shipping all orders via air is too costly, because of the high transportation cost. Nevertheless, we also researched a dual sourcing strategy, where products are partly shipped via sea transport and partly via air transport. The best solution is to send 6% for two weeks of safety stock and 25% for one week of safety stock of all shipments via air respectively and the rest via sea. We advise Philips to implement these settings, which will result in a profit increase of 0,5% and 4,1% respectively and an improvement in CSL-availability of 1,6% and 9,2% respectively.
- We develop one feasible idea for late customization for SKU B, which reduces the lead time by six weeks, but adds additional stock in the pipeline. The lead time reduction leads to an average CSL-availability YTD improvement of 1,8% for two weeks of safety stock and 5,3% for one week of safety stock. The additional handling and inventory costs however lead to a profit decrease of 5,1% when having two weeks of safety stock and 1,2% when having one week of safety stock, which makes this solution not desirable for implementation.

Our final recommendations for further research to Philips are the following:

- We have seen that big and small customers are supplied from the same stock, which leads to high variations in demand and more complications in forecasting demand accurately. Therefore, we recommend to research if serving big and small customers in a different way or from different stocks leads to higher service levels. For example, bigger customers could be supplied directly from the production site.
- We recommend to investigate if the target setting for the customer service level is cost-effective. Literature suggests that there is an optimal balance between inventories and customer service target settings, see (Jeffery, Butler, & Malone, 2008). Above that point, improvement of customer service level through setting higher safety stocks is too costly. We did not research the optimal target setting for Philips, because finding the right safety stock settings is out of scope. Furthermore, to perform this analysis with the method described in this article, a volume fill rate instead of an order fill rate is needed.
- We advise to reconsider if order line fill rate is the right measurement. Markets can influence the outcome of their CSL by prioritizing small customers just to get a nice KPI. This can result in lost opportunities with big customers. A Market can choose to not fulfill a few big order lines and still have a higher KPI than another Market that prioritizes big order lines over smaller ones, resulting in a higher sales volume but a lower KPI. We suggest Philips to use the volume fill rate instead of the order line fill rate, which makes it easier for production to prioritize orders and gives better insight in the performance of the markets in relation to sales and profitability.

1 Introduction

This research is done within the scope of my graduation project of the Master program Industrial Engineering and Management at the University of Twente. It is held inside a supply chain management team of the global company Philips.

> *A supply chain is a set of facilities, supplies, customers, products and methods of controlling inventory, purchasing, and distribution. The chain links suppliers and customers, beginning with the production of raw material by a supplier, and ending with the consumption of a product by the customer.* (Sabria & Beamon, 2000, p. 581)

The supply chain is a network in which a finished good is produced and distributed to the customer, to fulfill the customers' needs. The objective of supply chain management is to reach a high customer service level in the most cost-efficient way in order to sustain profit and growth. For a global company like Philips, it is substantial to make optimal strategic decisions like "location of facilities (plants and distribution centers), flow of goods throughout the supply chain [...], and assignment of customers to distribution centers" (Sabria & Beamon, 2000). Furthermore, operational optimization which includes determining the safety stocks, production batches, order quantities, production lead times and distribution is significant to reach the objective of a high customer service level. Uncertainties like customer demand have a great impact on the supply chain performance of the company. To forecast the demand accurately is a big challenge for most companies (Beutel & Minner, 2012). Supply chain management aims to find a balance between inventory levels and shortage due to low forecast accuracy.

1.1 Introduction to Philips

Philips is a well-known international company founded in 1891 with its headquarters currently in Amsterdam. Their mission is "Improving people's lives through meaningful innovation". By 2025, Philips wants to help 3 billion people improving their lives through sustainable and healthy innovations. The satisfaction and health of their customers is the driving factor for Philips.
Philips is divided into the sectors Health Systems and Personal Health. The sector Personal Health is organized in three pillars, Business Groups, Markets and a Single-Value Added Layer (which includes HR, Design, Procurement etc.). The sector is divided into five different business groups (BGs). BGs are responsible for strategic review and profit and loss, while being accountable for operations (Manufacturing and Supply Center) and innovation activities. The Markets are responsible for demand planning, in-market activation, e- commerce, consumer care and the Commercial Organizations (CO) and their activities. A BG is divided into several businesses categories, which can be seen as a high level group of products.

This study will focus on a business category inside Philips, further referred to as the case study business. The case study business of Philips is growing and is selling a wide variety of products. The customers are distributors and retailers and not directly the end-consumer. Philips customers expect shorter delivery times and more customization. That is why the case study business needs to be able to react fast and flexible to customer orders to hold its high reputation and constantly gain new customers. The challenge for the Philips business is to deliver reliably to the customers whatever product they want.

The supply chain of the case study business is organized in the following way. The end-customer buys Philips products from the shelf at a retailer, who is not part of Philips. The customers of the Philips business are retailers or distributors who sell to retailers. Customers can order at their Philips Commercial Organization (CO). The COs make forecasts of these orders, which are communicated to one of the Supply Centers (SCs) depending on the region and the type of product. The SCs order according to the forecast at one of the in-house factories or at one of many suppliers. Raw materials, which are necessary for the in-house factories are also ordered at suppliers by the responsible person of the factory. The finished products are stored at Regional Distribution Centers (RDCs), which are linked to the COs and then delivered to the customer or are directly delivered from the factory to a distributor. For a detailed description of the supply chain see Chapter 2.1.

1.2 The customer service level at Philips

The customer service level KPI (CSL), or customer service level of OTTR (on-time to request) orders, measures how many of the order lines were delivered On Time (OT) & In Full (IF) as requested by the customer versus the total number of order lines requested by the customer. To calculate the CSL OTTR, the following formula is used.

$$CSL\ OTTR\% = \frac{\sum Order\ lines\ delivered\ OT\&IF}{\sum Order\ lines\ requested}$$

In Full means that the actual delivered quantity is at least the requested quantity. On time means that the actual delivery date is before or on the requested delivery date.

A CSL OTTR failure occurs when a customer receives an order not in full or not on the requested delivery date; or both. In Philips terminology, this failure is called a CSL OTTR hit and is categorized in six buckets; namely customer, availability, Sales Office, warehouse, distribution and other. Every hit can only be placed in one of the buckets. It can happen that a CSL OTTR hit occurs due to multiple reasons. Therefore the hit needs to be reviewed in the right order. For every bucket it is checked if the hit falls in that bucket. The first bucket is called *customer* bucket and includes all failures which are caused by wrong expectations of the customer, for example when he orders a product which is already phased out. The next bucket is the *availability* bucket. A hit occurs if products are not available at the RDC due to production or transport issues or if the forecast was not accurate. Next, the hit can be assigned to the *Sales Office* bucket if there is some financial issue like a credit block issue, which means that a customer has exceeded his credit limit and still places an order. Since the system does not allow to process orders above the credit limit, this order will not be fulfilled. The next bucket is the *Warehouse* bucket. Mostly late goods dispatch causes a hit in this bucket. A hit falls in the next bucket, the *Distribution* bucket, if there are problems due to the distribution from the RDC to the customer. Mainly these problems occur due to a failure of the carrier, e.g. traffic delay or planning error. The last bucket is for all other reasons like system errors, which could not be placed in one of the other buckets. Due to the responsibilities of the business group, this research will focus on the improvement of the CSL OTTR of the availability bucket, also called CSL availability.

1.3 Goal

Philips strives for high quality products and a high customer service level. Their goal is to reach a customer service level of on-time to request orders of 95% by 2020. The goal for 2016 is to reach a CSL of 87% for all products in all BGs. The case study business reached a CSL OTTR Year to date (YTD) of 84% in 2015.

1.4 Problem description

Because of the growing targets inside the Philips business, it is a prerequisite to deliver all receiving orders on time and in full when requested by the customer. So the CSL OTTR is a key enabler to support the strategy to grow. The operational committee has the vision to grow to a 95% CSL OTTR in 2020. The Supply Chain Management team would like to investigate what is needed to be done differently to reach this target.

When receiving an order, the Philips business wants to deliver the product on-time and in full, regardless of orders already promised. The demand forecasts are made by the Markets. For the

Markets, it is difficult to reach a high forecast accuracy. The company measures the forecast accuracy using the weighted mean absolute percentage error (wMAPE). At the start of this research, the wMAPE of the forecast of two month before the sale is 47% on average for the first 3 month of 2016. Therefore, the supply chain management team needs to setup their supply chain flexibly to cope with low forecast accuracy.

Another challenge the Philips business faces is the fact that different international customers have different requirements regarding the products; for example the language on the packaging. In the current set-up, late customization is not possible for all products and therefore the amount of products produced for a specific region has to be decided in an early stage. That is why the finished goods are not flexible, which means that they cannot be used for different markets, which leads to wasted resources due to high inventories or high customer order lead-times due to unavailability of products.

1.5 Research questions

1.5.1 Main research statement

From the problem description and the goal of the company, the following main research statement can be formulated:

The goal of this research is to improve a part of Philips' supply chain in order to improve the customer service level in the most profitable way.

1.5.2 Sub-questions

To reach the research goal, a number of sub-questions can be asked. First we need to take a look at the current state of the supply chain and gather all the data needed for the project. Then the data can be analyzed and possible causes for the problem can be found. After the analysis, different solutions to the given problem can be gathered from given models and/or brainstorming. The best possible solutions should be compared and finally an advice for the best solution can be given. Due to time limitations, we will not research how to implement the solution. This process leads to the following sub-questions divided into three phases.

Definition phase:

1. *How is the current supply chain of the case study business structured?*
   In Chapter 2, we describe the current structure of the supply chain in general for all products of the case study business. We look at the order process and the product flow.

2. *How is the customer service level measured? What are the results? What are the root causes for low customer service levels?*
   We describe how the customer service level is measured in Chapter 2. We gather data about customer service levels for different product groups and show the results. We find product groups which are the main drivers for low customer service level and look at the root causes for the low performance.

Analysis phase:

3. *Which factors influence the service levels and how do they relate to each other? What solutions/ interventions to improve the service level have already been found by researchers?*
   Literature provides a lot of information on customer service levels. We perform a literature study about the factors that influence the customer service level and possible solution models in Chapter 3.

4. *Which of those solutions/ interventions could be suitable for Philips and should be analyzed in further detail?*
   To answer this research question, we talk to stakeholders and discuss different types of solution models. Chapter 4 gives an overview of different solution designs and answers whether or not a solution is suitable for Philips. We decide on solution models that we want to investigate in this chapter.

Solutions phase:

5. *What are the expected costs and impacts of each solution / intervention on the customer service levels? What set of solutions / interventions should Philips implement?*
   In order to generate results, we first develop a simulation model in Chapter 5. Then in Chapter 6, we report the costs and improvement in customer service level of each solution. With this information, we evaluate whether a solution should be implemented or not. The implementation of the solution is not be researched due to time limitations.

1.6 Scope

For this research, the following scope is defined:

- The research will be done in the case study business.
- The data used for this research will be limited to the data from the last year, January 2015 until August 2016.
- The Regional Distribution Centers (RDCs) have their own team working on the planning. Therefore the forecasts and lead-times between the RDC and the customer will be taken as input.
- The analysis of the forecasts and forecast accuracy will not be done in this research since the supply chain should be able to react fast to forecast changes.
- The research will only focus on CSL failures due to availability reasons at the RDC, like unavailability of a product due to production, transportation or forecast issues.

1.7 Research Approach

In Chapter 1 we have described the background and defined the main objective of the research and the corresponding research questions. The main goal of this research is to find ways to improve the customer service level of the case study business of Philips.

In Chapter 2, we give an overview of the current supply chain of the case study business. We define our scope to be CSL OTTR failures due to unavailability of products. We gather data to find the product groups which are the main drivers for low CSL OTTR and their measured root causes. In Chapter 3, we perform a literature study about which factors influence the customer service level and what types of solutions are known to improve the customer service level. After the literature study, we talk to stakeholders to discuss which of the solutions that are described in literature are suitable in practice and for this type of research in Chapter 4.

To model our solutions and generate results, we use discrete-event simulation described in (Law, 2007). We develop a simulation model in Chapter 5. In Chapter 6, we design our experiments and state the results. We compare the costs and improvement in service level of each solution and give advice on which solution Philips should implement. Finally, in Chapter 7, we list further recommendations and our overall conclusions of this research.

2 Context analysis

The first part of the research has the aim to gather all information needed for this research to be able to find the core root causes to the problem. To answer sub-question 1, the supply chain structure will be described in further detail to gain a better understanding of the processes inside Philips. In the second part of this section, the customer service level of different product groups will be evaluated in order to answer sub-question 2. This will be the basis to make a decision on which problem to tackle.

2.1 The supply chain

The supply chain of Philips is quite big and complex due to the vast amount of suppliers, factories, markets and products which are spread out globally. This chapter will give a general overview of the order flow and the physical product flow.

2.1.1 The products

The products are divided into different groups on different levels. The highest level is the Material Article Group (MAG), which is the most aggregated level. The MAG consist of multiple Article Groups (AGs). One step down in the hierarchy is the Product Article Group. This is a code which gives information about the kind of product and information on the factory it comes from and is only used by the factories. The next level is a so called 12NC code which contains all specific information of the product that is needed for production, e.g. color, size, type and country sold in.

2.1.2 The customer

The customers of Philips are retailers like e.g. Amazon or distributors, who sell to retailers. The retailers and distributors are not a part of Philips. They are selling the final product to the consumer.

Philips' customers are spread all over the world and therefore split up into different markets. Philips has several Commercial Organizations (COs), which are responsible for customer orders from one or several specific countries. For example, one CO is responsible for the market DACH, which includes customers from the countries Germany, Austria and Switzerland.

2.1.3 Order process

The demand planners at the COs forecast the total volumes of the customer orders. They are in close contact with their customers and have information about special discounts or promotions. Forecasts of monthly sales volumes are made for the next 12 months and need to be updated weekly. The 6-12 months forecasts are needed for strategy planning and making strategic decisions on, for example, capacity. The 3-6 months forecasts include information on new product introductions or on products which are phasing out. For the next 0-3 months, a weekly volume forecast is required for the production. In these three months, there is a frozen window in which the production planning is fixed. Therefore, forecast changes during that frozen period cannot be taken into account in the production planning anymore. However, most changes are made in the frozen windows, which can lead to stock-outs in the warehouses and therefore a lower customer service level. When a demand planner needs more products in a week of the frozen time window, then the factories can deliver this earliest in the first week after the frozen window. However, if there is extra capacity left, then the factories can deliver earlier or airfreight can be used to shorten the delivery time.
The forecast that is created by the demand planner is referred to by the unconstrained forecast within Philips. It is the forecasted demand of the markets without any capacity restrictions. For this forecast, a statistical method is used to calculate a baseline for the forecast. This baseline is then evaluated and sometimes changed by the demand planner based on experience and additional information like promotions. The supply planners at the SCs get the so called unconstrained forecast of all COs. Together with the factory, either one of the in-house factories or one of the suppliers or both, a production schedule is made. Due to capacity restrictions, the factories cannot always supply the

forecasted quantities. Then the capacity is allocated using allocation rules. After that step, a confirmed supply planning is sent out to the demand planners. The quantity that the factory confirmed to deliver is referred to by the constrained forecast within Philips. The raw materials, which are necessary for the in-house factories, are also ordered at suppliers by the responsible person of the factory.

2.1.4 Physical product flow

The products are produced at one of the in-house factories or at the suppliers. The finished goods are delivered to the warehouses of the COs, which are called Regional Distribution Centers (RDCs). From the RDCs, the products are delivered to distributors and the distributors send the products further to the retailers or directly to the consumer, see Figure 2.1. In some cases, finished goods will be sent directly to the distributor and not through the RDC or the finished goods will be sent directly from the RDC to a retailer, see Appendix A.1. The transportation is done by ship, train, truck or air freight depending on the urgency and the location. Each CO is linked to one RDC. In Europe, there are three RDCs, the 3DCs, which store products for different European countries. In that way, if one country has a shortage, the products can be re-allocated to another country taking into consideration packaging and language requirements. Other countries around the world have, with some exceptions, one RDC per country. Sometimes an RDC can also have a packaging function, which means that finished goods are combined and are sold as a set of products. These sets can be unique per country. Since the RDCs, factories and customers are spread all over the world, the replenishment lead times and the customer lead times vary for different products and different markets.

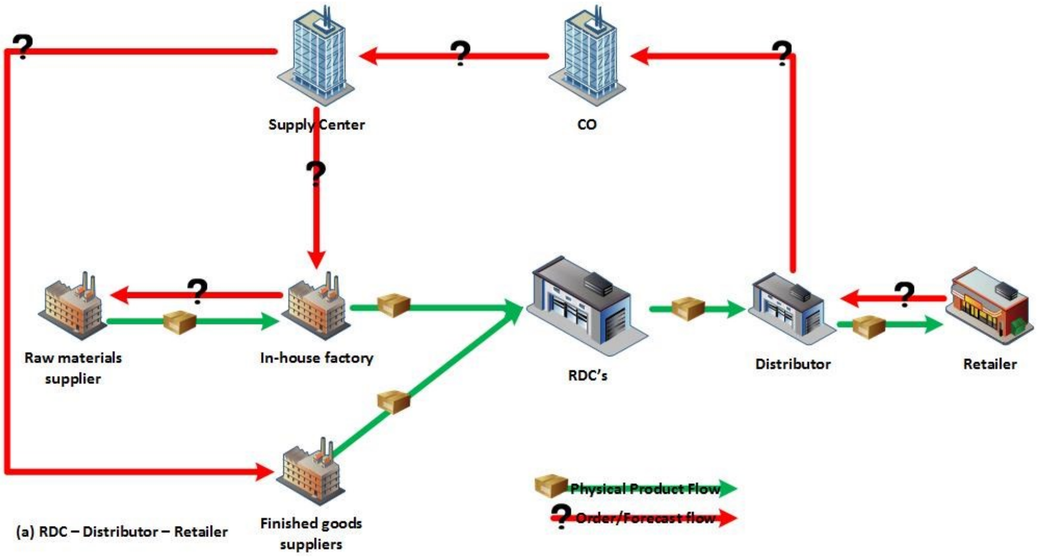

Figure 2.1 Supply chain structure

2.1.5 Inventories

Philips divides inventories into commercial inventory and industrial inventory. Commercial inventory is the inventory of finished goods at the RDC. Markets are responsible for holding the right amount of stock. High inventory levels can lead to high inventory holding costs and have as a risk that the inventory does not get consumed. On the other hand, if inventory is too low there is a risk that Philips cannot deliver the right products to its customers and thereby decreases customer service levels. For every product there is a certain level of safety stock, which differs depending on the size of the product, the desired service level and the demand for the product.

Industrial inventory consists of all component and key module inventory in the supply chain before it is a market specific finished good. The BG is responsible for handling the industrial inventories. The availability of components influences whether factories can manufacture the finished good, which is ordered by the CO in the market. For this inventory, there is the same trade-off between inventory costs and availability.

2.2 Customer service level

Philips reviews its CSL OTTR on a weekly and monthly basis. Every week the CSL OTTR is reported for the current week and compared to the performance of the weeks before. At the end of the month, the performance of the whole month is compared to previous months. Also the CSL OTTR YTD (yield to date) is reported, which is the average CSL OTTR of all weeks of the year. The goal is to reach a CSL OTTR YTD of 95% by 2020. At the start of this project (week 20 of 2016), the CSL OTTR YTD of the case study business is 89%. So the CSL OTTR needs to improve by 6% within the next 3 years to reach to target of 95% for that business.

The CSL OTTR failures are categorized in six buckets; customer, availability, sales office, warehouse, distribution and other. Due to CSL OTTR reports of the Customer Collaboration Team of Philips, 78% of all CSL OTTR failures are assigned to the availability bucket. Since this is the biggest bucket and due to the responsibilities of the supply chain management team and the business group, the focus of this research will be limited on how to decrease the amount of hits, which fall into the availability bucket. For the rest of this research, it is assumed that the other the buckets will not change. In the following, we call the CSL OTTR YTD for product availability at the RDC CSL-availability YTD.

The CSL-availability YTD at the start of this research (2016 week 20) of the case study business is 91%. Due to the calculation of the CSL-availability, if the CSL-availability score is increased with 1%, the total CSL OTTR will increase also with 1% assuming that the other buckets do not change. Since the current CSL OTTR YTD of the case study business is 89%, the CSL-availability YTD needs to increase by 6% to reach the goal of a CSL OTTR YTD of 95%. Therefore, the target for the CSL-availability YTD is 97%.

2.2.1 The availability bucket – root causes

Inside the CSL-availability bucket, there are three main root cause categories namely production, mainstream and forecast and other CO related issues. Every availability hit is categorized in one of the three root causes. We first explain the meaning of the three categories and then quantify how many failures fall in each category in Table 2.1. A hit can be identified as either production or mainstream and if it does not fit in one of these root causes, it automatically gets the forecast and other CO related issues root cause assigned. When a hit's root cause is production, then there are any kind of production issues like capacity shortage, quality issues or machine downtime. The mainstream root cause is about all kinds of transportation issues from the factory to the RDC, e.g. traffic delays or carrier capacity issues. Finally the forecast and other CO related issues root cause is about everything which does not fit to the production or mainstream root cause and relates to a low forecast accuracy of the markets. Root causes for this category can be e.g. long lead times to the markets or customer collaboration issues. Other CO related issues are for example low safety stock settings or selling new products to the customer which is not in the warehouse yet. Due to one of the three previous reasons, the requested product is not available at the RDC and therefore the CSL OTTR failure is categorized into the CSL-availability bucket. For a more detailed root cause analysis, see Appendix A.2.

To be able to measure the CSL OTTR availability, all order lines requested by the customer and all orders lines that caused an availability hit are gathered in a database and linked to the specific root cause (production, mainstream or forecast and other CO related issues). Every order can be linked to a certain business, market, supply center and product. Also the order quantities and the quantities that caused the hit are stored in the database. With this information, the CSL OTTR availability score for each product, market, business, supply center and root cause can be calculated. Philips reports the results in a weekly CSL-availability dashboard. In the dashboard, the scores can be evaluated on different levels. For example a product can be reviewed on Main Article Group (MAG) or Article Group

(AG) level down to the specific product. Markets can be reviewed per region, cluster or specific country. This gives the business the possibility to analyze which product groups or products have the biggest impact on CSL in what market and from which supply center they are ordered. The impact of a specific product can be calculated through dividing the amount of hits of that product by the total amount of orders of all products.

Table 2.1 shows results of our root cause analysis given the data from the CSL-availability dashboard from 2015 and the first half of 2016. 84% of the availability hits of the case study business of 2016 have forecast and other CO related issues as a root cause. In 2015 it was 74%. The rest of the hits in 2016 is split evenly through production and mainstream. In 2016, of the 9% of the order lines with a CSL-availability hit, 7,5% were due to forecasting and other CO related issues and only 1,4% are due to production and mainstream issues.

Table 2.1 Root cause analysis

	2015		2016	
Root cause	% of hits	Impact on CSL-availability	% of hits	Impact on CSL-availability
Forecasting & Other CO related Issues	74%	7,70%	84%	7,50%
Mainstream delay issue	19%	1,90%	8%	0,70%
Production issue	8%	0,80%	8%	0,70%
Total		**10,40%**		**8,90%**
CSL-availability		**89,60%**		**91,10%**

2.2.2 Product group focus

In this section, the CSL OTTR of different product groups is analyzed to decide on which product groups the focus should lie during this research. The data used for this research is from the year 2015 and the first half of 2016. The years 2015 and 2016 will be compared to analyze developments and distinguish current issues from structural issues. Stakeholders are also asked to evaluate if possible differences in product group performance are caused by incidental issues.

Recall that the customer service level is calculated by dividing the amount of all OTIF order lines by the total order lines during a period. Therefore every hit has the same impact on the CSL OTTR, independent of the quantity ordered or the size or importance of the customer. If the number of hits is reduced, then the CSL OTTR improves and therefore the goal is to minimize the number of hits. Also it holds that the more hits, the more impact on the CSL OTTR. Logically, a product group with a big amount of order lines has a higher chance to have a lot of impact than a product group with a small amount of orders. If we look at the amount of hits within all orders of a specific product group, then a product group with a big amount of order lines can perform better than a product group with a small amount of order lines, but still have a greater impact on the total CSL availability OTTR, when the number of hits is bigger.

When the focus lies on a small product group which performs badly on its own, a reduction of the amount of hits might be easy to achieve. However, it will have a relatively small impact on the CSL OTTR, because it has a relatively small number of hits. Product groups with a really small amount of hits or no hits can be left out, because they have almost no impact on the CSL OTTR. Therefore it would be preferred to look at a big product group in terms of order lines with a bad performance on its own. If a big product group performs really well on its own compared to other product groups, still a big improvement on the CSL OTTR can be achieved, but it will probably be harder to really improve the performance. So, a balance needs to be found between the impact of the product group and the possibility of improvement.

To find the right product group to focus on, first the CSL-availability data from the Philips CSL-availability report was analyzed for different AGs. Table 2.2 shows the results of this analysis. The column "Impact on total CSL" shows with how much percentage the CSL OTTR of the case study business rises if the CSL availability of that AG improves to 100%.
Then we interviewed the S&OP manager and the corresponding Supply Planners, see Appendix A.3, to find out the root causes for low CSL performance. In that way, low performance due to current issues, like reduced safety stock due to maintenance in the factory or new product introductions, could be filtered out. The focus can then be on AG's with a structural issue and a high impact on the CSL availability.

Table 2.2 Impact on CSL per product group, sorted on highest to lowest impact on the total CSL in 2016

AG		**2015**		**2016**		
	% of all order lines	Impact on tot. CSL	CSL-availability	Impact on tot. CSL	CSL-availability	Comments
A	**21%**	**1,59%**	**92%**	**1,51%**	**93%**	**Structural and biggest impact**
C	7%	0,86%	88%	1,11%	83%	Temporary problem in production
D	13%	1,20%	91%	1,09%	92%	
E	21%	1,76%	92%	1,08%	95%	
F	11%	1,76%	86%	0,88%	92%	Supplier issue in 2015, now solved
G	3%	0,34%	86%	0,62%	76%	Temporary problem with supplier
H	4%	0,39%	91%	0,58%	86%	Temporary problem in production
B	**4%**	**0,88%**	**78%**	**0,53%**	**87%**	**Structural**
I	1%	0,12%	87%	0,22%	85%	Temporary problem with supplier
J	3%	0,22%	92%	0,21%	93%	
K	2%	0,24%	89%	0,20%	91%	
L	1%	0,11%	93%	0,17%	88%	Temporary problem in production
M	1%	0,14%	86%	0,17%	84%	Structural
N	2%	0,24%	91%	0,15%	94%	

For this research, we choose to focus on product group A from the in-house production site A and product group B which is produced at a supplier.

Product group A accounts for 21% of all order lines. It has a CSL-availability of 92% in 2015, which leads to a reduction of the total CSL-availability of the case study business of 1,6%, see Table 2.2. Although the CSL-availability is already reasonably high, the impact on the overall CSL is the biggest. Therefore a small improvement for this product group can result in a big improvement for the overall CSL.

The products of product group A are produced in an in-house factory in Europe and are sold globally. The order process for the products is visualized in Figure 2.2. In week 1, the markets fill in the forecasts on Monday into a planning system. The demand planners make a high level production plan manually on Wednesday taking into consideration the production capacity and the prioritization wishes of the markets. The factory planner will confirm or adjust the plan on Thursday and the confirmed plan will be sent out to the markets on Friday. So the total planning lead time is one week. The next week products will be produced and are planned to be ready at the latest Saturday at 6:00 am. So in general the production lead time is one week, however prioritization is possible to shorten the lead time. Products which are produced in the beginning of the week can be shipped earlier and therefore arrive earlier at the RDC. As soon as the products are ready, they are booked into the Warehouse at the factory and can be shipped as soon as there is a full truck load available for shipping. The distribution time to the countries varies depending on the distance and shipping method.

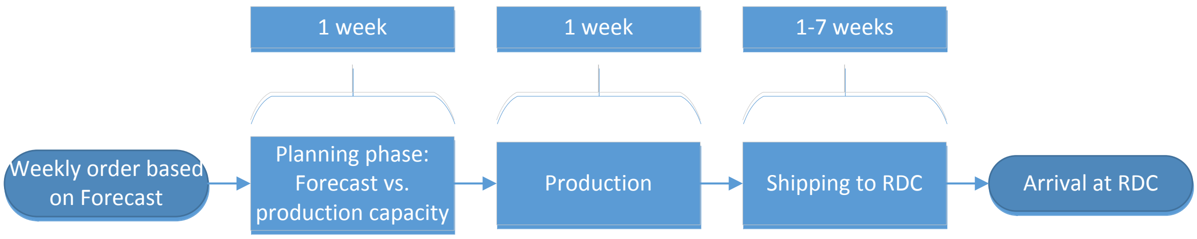

Figure 2.2 Product group A order lead time

The total lead time from the point of placing an order until the order arrives thus varies from 3 to 9 weeks depending on the location of the market.

Product group B accounts for 4% of all order lines. It had a CSL-availability of 78% in 2015, which lead to a reduction of the total CSL-availability of the case study business of 0,88%, see Table 2.2. The impact on the overall CSL is lower than that of product group A, but the CSL-availability is much lower, which leaves much room for improvement. Therefore we expect the effect on the CSL to be quite high.

The products of product group B are ordered at a supplier in Asia and shipped to Europe. Therefore the scope for this project will be all the products, which are sent to Europe. The order process of product group B is similar to that of product group A. The order confirmation takes one week, the production takes three weeks and the delivery of the products from Asia to Europe takes five to six weeks, see Figure 2.3. Therefore the total lead time is nine weeks.

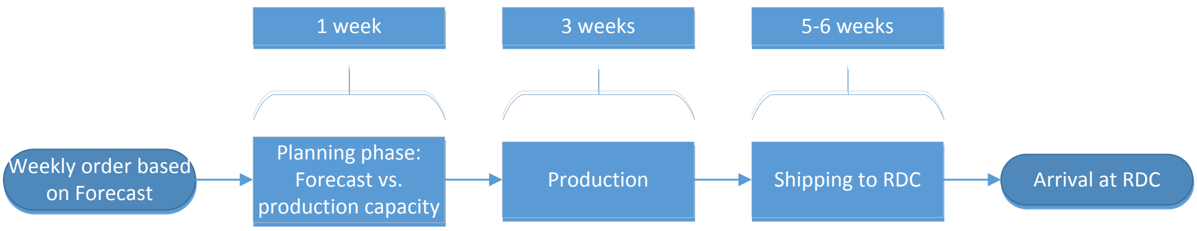

Figure 2.3 Product group B order lead time

2.2.3 Root causes product groups

The following table shows the root cause split for product groups A and B. We can see that also within these product groups, the biggest root cause for a CSL-availability hit are forecast errors.

		Root cause split		
AG	CSL-availability 2016	Forecasting & Other	Mainstream	Production
A	93%	6%	1%	0%
B	87%	10%	2%	1%
TOTAL	91%	7%	1%	1%

Since forecasting is the biggest issue for the case study business and since it has the biggest impact on the customer service level, the focus of this research lies on understanding more in detail the origin of this root cause and finding a way to resolve these issues. To understand better why a hit occurs, we brainstorm with several stakeholders listed in Appendix A.3 amongst which the Supply Chain Manager Customer Collaboration about potential root causes for forecasting errors, transport problems and production issues, see Appendix A.2. From that, we can conclude that demand variability is one of the main factors that influence the forecast accuracy. Note, that we excluded improving forecast methods from our scope in Chapter 1.6. We are interested in understanding better how demand variability influences the customer service level and in how the customer service level can be improved. Therefore, we conduct a literature study in Chapter 3.

2.3 Conclusions

In this chapter, we gave an overview of the current supply chain, described how the customer service level is measured at Philips and analyzed the main drivers and root causes for low customer service level.

Due to CSL OTTR reports of the Customer Collaboration Team of Philips, 78% of all CSL OTTR failures are assigned to the availability bucket. We focus on improving product availability due to this high percentage and due to the responsibilities of the supply chain management team. The CSL-availability YTD, which is the CSL OTTR YTD of product availability, at the start of this research (2016 week 20) of the case study business is 91%.

Most of the CSL-availability hits, 84% in the first half of 2016 and 74% in 2015, are due to forecast errors. Due to the responsibilities of the supply chain management team, we do not want to research the effect of different forecasting methods, but we want to find ways of being more independent of bad forecasts.

The focus of this research will lie on product group A, because of its big impact on the CSL-availability. Although, it has a quite high CSL-availability YTD of 93%, small improvements in this product group have a big impact, up to 1,5%, on the overall CSL-availability of the case study business.
The focus of this research will also lie on product group B, because of its structurally low CSL-availability YTD of 87%. This CSL-availability YTD leaves much room for improvement.

3 Literature study: Customer service level improvement

Before choosing a strategy on improving the service level, a deep understanding of the customer service level and the factors which influence the customer service level is needed. This chapter focusses on answering sub-questions 3 and 4. First, a literature study is done to gain insight in factors which influence the customer service level and their relation, see sections 3.1 and 3.2. Then, we choose a focus area for this research and study solution methods which are found in literature, see section 3.3.

3.1 Customer service level in inventory management

The customer service level is a measurement which is used in companies to measure their performance. Performance measurements which are customer service related are "stock service level, delivery precision, delivery reliability, delivery lead time and flexibility" (Jonsson & Mattsson, 2009, p. 47). The CSL-availability measure, used by Philips, is a stock service level, which measures the amount of order lines which can be delivered in full directly from stock. In literature this sort of measure is often called order fill rate or line item fill rate. Larsen & Thorstenson (2008) describe the order fill rate "as the fraction of complete orders that can be filled directly from inventory" (p. 798). Therefore the problem of improving the customer service level can be seen as an inventory management problem.

3.1.1 Characteristics of the order fill rate

Let us again look at the formula of the order line fill rate to discuss the characteristics of the performance measure:

$$CSL\ OTTR\% = \frac{\sum Order\ lines\ delivered\ OT\&IF}{\sum Order\ lines\ requested}$$

Imagine a set of order lines with some having a great amount of products on order and some only a small amount of products on order. To reach a higher customer service level, a company should then prioritize the order lines with a small amount of products on order to minimize the amount of failures due to unavailability of products. Logically the customer service level is lower when prioritizing the other way around. In practice, customers with big order volumes have a higher chance to have a higher prioritization, because they also drive sales. To reach a trade-off between customer service level and sales, companies can split the order of a big customer to prevent too many order line hits.

Moreover when delivering to several markets, replenishment decisions in situations of shortages at the production site influence the total customer service level. The amount of order lines and volumes per order line can differ between markets. Therefore prioritizing a market with customers who frequently order small order lines opposite to markets with customers who order a large bulk of products in one order line on a monthly basis can lead to a higher customer service level, because of simply fulfilling more small order lines instead of fulfilling bigger order lines. We will evaluate in the discussion in Chapter 7 what kind of effect this characteristic can have on the company.

3.1.2 Demand variability in relation to the customer service level

Many researchers ((Gupta & Maranas, 2003); (Jeffery, Butler, & Malone, 2008); (Towill, 1996)) state that a big challenge for companies nowadays is to achieve a high customer service level at the lowest cost possible, because customers expect highly customized products and fast delivery times. Companies need to react quickly and flexible to changes in demand to avoid excess stock or unavailability of stock.

Variability in demand is described a lot in literature. Lee et al. (1997) describe that order variation throughout the supply chain, also called the Bullwhip effect, can lead to "excessive inventory, poor customer service due to unavailable products or long backlogs, uncertain production planning (i.e., excessive revisions), and high costs for corrections, such as for expedited shipments and overtime" (p. 93). The authors suggest fighting the bullwhip effect and therefore improving inventory availability through "information sharing, channel alignment, and operational efficiency" (Lee, Padmanabhan, & Whang, 1997, p. 98).

Jeffery, Butler & Malone (2008) found that "the [minimum cost volume fill rate] service level decreases as forecast error and demand variability increase" (p. 231). "The higher the demand uncertainty, the more difficult it is to generate accurate forecasts" (Jonsson & Mattsson, 2009, p. 105). So high demand variability leads to inaccurate forecasts and can therefore lead to a poor customer service due to unavailability of stock (Beutel & Minner, 2012).

3.2 Coping with demand variability

In this section, we evaluate methods from literature to cope with demand uncertainty and decide which methods to look at regarding the scope of the case study.

3.2.1 Forecast accuracy

One way to cope with demand variability is to invest in better forecast accuracy (Li, Erlebacher, & Kropp, 1997). In this research, we will not focus on finding a better forecasting method, but assume the current method as given and will use the forecast data as input. A way to reach a better forecast accuracy is by looking at the length of the forecast horizon.

> *The longer the horizon, meaning the farther in the future that must be forecast, the more difficult it will be to avoid forecast errors. By reducing throughput times to allow a shorter forecast horizon, measures to cut throughput times in material flows are effective methods of improving potential forecast precision.* (Jonsson & Mattsson, 2009, p. 111)

Towill (1996) states that researchers have found as a rule of thumb that "reducing the lead time by 50 per cent will reduce the forecast error by 50 per cent" (p. 17). These statements let us assume that lead time reduction will have a positive effect on the forecast accuracy and therefore improve the customer service level.

3.2.2 Safety Stocks

To buffer against demand uncertainty and forecast errors and to reach a certain target service level, companies add safety stocks to their inventories. Beutel and Minner (2012) describe that "inaccuracy of forecasts leads to overstocks and respective markdowns or shortages and unsatisfied customers" (p. 637). The authors state that safety stocks are used to secure against forecast errors.

A standard way of calculating the size of the safety stock for a periodic review system with an order size of a multiple of the fixed order quantity is described in (Beutel & Minner, 2012):

$$SS = k \cdot \sqrt{LT \cdot {\sigma_d}^2 + {\mu_d}^2 \cdot {\sigma_{LT}}^2}$$

where SS is the safety stock in volume, k is the service level factor, LT is the replenishment lead time plus review period, σ_d is the standard deviation of demand, μ_d is the mean demand per period and σ_{LT} is the standard deviation of lead time. From the method, we can see that if we keep all other factors constant, the service level will increase when the lead time decreases. Because the forecast error will decrease with shorter lead times, the effect on the service level should in theory be even bigger.

3.2.3 Lead time

In the last two sections, we have seen that lead time has an effect on forecast accuracy and safety stocks and therefore also on the customer service level.
Ouyang & Wu (1997) state the importance of lead time reduction in inventory management. “By shortening the lead time, we can lower the safety stock, reduce the loss caused by stock-out, improve the service level to the customer, and increase the competitive ability in business” (p. 875). The authors describe the lead time as consisting of “order preparation, order transit, supplier lead time, delivery time, and setup time” (Ouyang & Wu, 1997, p. 875).

Lee et al. (1997) describe lead time reduction or just-in-time replenishment as a solution to counteract the bullwhip effect. “With long lead times, it is not uncommon to have weeks of safety stocks. The result is that the fluctuations in the order quantities over time can be much greater than those in the demand data” (Lee, Padmanabhan, & Whang, 1997, p. 95). Therefore reducing the lead time will reduce the bullwhip effect and therefore increase the customer service level.

Ciancimino et al. (2012) study the effect of different lead time settings on the average fill rate for a synchronized supply chain and conclude that long lead times affect the customer service level. The authors suggest to raise the safety stocks to maintain a high service level when having long lead times, which can lead to higher inventory holding costs, but also refer to studies, which have proven the benefits of lead time reduction (Ciancimino, Cannella, Bruccoleri, & Framinan, 2012).

Hopp et al. (1990) list some benefits of lead time reduction from a sales and production perspective. With reduced lead time, companies cannot only deliver faster and reduce inventory, but also reduce the need for a distant forecast (Hopp, Spearman, & Woodruff, 1990).

3.2.4 Product variants

Customization is one of the challenges companies have to cope with nowadays. From a marketing point of view, a wide range of products can raise the overall market share, because a wide group of customer segments can be targeted (Wan, Evers, & Dresner, 2012). However, “marketing research has also suggested that ‘excess’ product variety may lead to selection confusion for customers, thus reducing the marginal benefits from variety” (Wan, Evers, & Dresner, 2012, p. 316). The authors Wan et al. (2012) also state the difficulty of product variety for operations management especially inventory management and operational performance. Stock-outs can result from high product variety and therefore cause a poor customer service level.

Lu, Efstathiou & del Valle Lehne (2006) find that to reach a high customer service level, companies can increase their inventories or reduce their number of SKUs. This will be a trade-off between lost sales and inventory holding costs. The authors also suggest to “maintain a responsive lean and dynamic inventory” (Lu, Efststhiou, & del Valle Lehne, 2006, p. 249), by phasing out SKUs that are not popular.

The aggregation level of the products to forecast has also an influence on the forecast accuracy. “Forecasting single products is considerably more difficult than forecasting groups of products” (Jonsson & Mattsson, 2009, p. 112). From this, it can be concluded that less product variants lead to a better forecast accuracy and therefore to a better customer service level.

Thonemann & Bradley (2002) state that a “[h]igh product variety decreases supply-chain performance measured in terms of replenishment lead time and cost” (p. 549). According to the authors, there is a trade-off between the product variety a company wants to offer their customers and the costs of setup times and higher inventory levels due to longer lead times and different products (Thonemann & Bradley, 2002). Postponement strategies can therefore help reducing lead time to the customer, by placing the customer order decoupling point later in the supply chain, which also reduces the necessity of long term forecasts.

3.2.5 The Customer Order Decoupling Point and late customization

In the previous section, we addressed the challenge of mass customization and the need for companies to act rapidly and flexibly to demand changes. The customer order decoupling point is the point in the supply chain where products are pull driven by order of the costumers. Before that point products are pushed through the chain. Lu, Efstathiou & del Valle Lehne (2006) find that placing the customer order decoupling point late in the supply chain, by having (semi-)finished goods in stock, helps serving the customer within a short time and helps dealing with mass customization, but as a trade-off can increase stock holding costs because of semi-finished good inventory.

Brown et al. (2002) state that manufacturers must hold high levels of inventories due to uncertain demand and long lead times in order to guarantee a certain customer service level, which is costly and risky. In order to increase the service level, or reduce inventories, the authors suggest postponement strategies, also called late customization, where “inventory is held at an intermediate point in a generic, non-differentiated form and is only differentiated when demand is better known” (Brown, Ettl, Lin, Petrakian, & Yao, 2002, p. 284). Brown, Lee & Petrakian (2000) state that “delaying the point of product differentiation can be an effective technique to cut supply-chain costs and improve customer service” (p. 65).

3.2.6 Information sharing

Lee et al. (1997) describe that the bullwhip effect, which can cause poor customer service levels, can be reduced through information sharing. The authors suggest to “make demand data at a downstream site available to the upstream site” (Lee, Padmanabhan, & Whang, 1997, p. 98). Then both sites have the same information to update their forecasts.

Ciancimino et al. (2012) describe information sharing or supply chain collaboration as “the alignment of planning, forecasting and replenishment systems among partners” (p. 49). In their research, the authors conclude that “synchronisation eliminates the bullwhip effect and creates stability in inventories under different parameter settings” (Ciancimino, Cannella, Bruccoleri, & Framinan, 2012, p. 50).

Zhao et al. (2002) study the impact of information sharing and the co-ordination of the replenishment of retailers under demand uncertainty on the supply chain performance. They find that information sharing and order co-ordination have a positive impact on costs and customer service level under all demand patterns (Zhao, Xie, & Zhang, 2002).

3.2.7 Order batching

Lee et al. (1997) describe order batching as one of the four major causes for the bullwhip effect. The authors describe that order batching results in higher fluctuations of order sizes upstream the supply chain (Lee, Padmanabhan, & Whang, 1997). However, there are common reasons for companies to order in batches, such as the costs for processing an order, which can increase exponentially when ordering frequently instead of periodically or transportation costs which are optimal for full truck loads, which is why suppliers want to supply batches at the size of a full truck load (Lee, Padmanabhan, & Whang, 1997). Lee et al. (1997) suggest that “companies need to devise strategies that lead to smaller batches or more frequent resupply” (p. 100).

Also Moyaux et al. (2007) also suggest that a lot-for-lot type ordering policy can eliminate the bullwhip effect. On the other hand, the authors also emphasize that “many reasons, such as inventory management, lot-sizing, and market, supply or operation uncertainties, motivate companies not to use this strategy" (Moyaux, Chaib-draa, & D'Amours, 2007, p. 396).

3.2.8 Echelon-based inventory control

Van der Heijden & Diks (1999) describe that in a local inventory control system, where inventories are controlled locally, if no demand or inventory information is shared between locations, a bullwhip effect may be caused and high safety stocks are required to reach a high service level. An integral control system however, where inventories are controlled centrally and are balanced with the use of the concept of echelon stocks, can reach the same customer service level with much less safety stock in the supply chain (Van der Heijden & Diks, 1999). Also Lee et al. (1997) states that "[e]chelon inventory — the total inventory at its upstream and downstream sites — is key to optimal inventory control" (p. 99) and suggests it as a way to improve operational efficiency to counteract the bullwhip effect.

3.3 Methods for improving the customer service level

After understanding the factors, which influence the customer service level, we can focus on finding ways for improvement. Regarding the scope of the case study business, we decide to exclude the opportunities in improving the forecast accuracy and the safety stock settings and instead focus on ways to act more flexibly and rapidly to demand changes. Options for synchronization, order batching and echelon-based inventory control have already been evaluated and partly implemented by the company. Possibilities for reducing the amount of product variants is left for further research since it is only relevant for one of the two studied product groups. Therefore, in this section we will focus on lead time reduction and late customization.

3.3.1 Lead time reduction

The replenishment lead time for the case study business is the lead time from the placement of an order until the arrival of the order in the RDC. This lead time consists of the order planning lead time, the production lead time and the transportation lead time. In this chapter, we will study ways to reduce those three lead times.

Safety stocks buffer against uncertainty in demand during the lead time plus review period. Therefore, a shorter review period at the same safety stock level should lead to higher customer service levels. Jonsson and Mattsson (2009) state that "[t]he review interval also influences the total lead time and thus the reaction time for covering material requirements as they arise" (p. 213) and that more safety stock is required with longer periodic review intervals (Jonsson & Mattsson, 2009). Therefore, to reach a higher customer service level, a continuous review period would be desirable. However, with a periodic replenishment system "planning of new orders can be carried out for a large number of items together, thereby making administration more efficient" (Jonsson & Mattsson, 2009, p. 213). Therefore choosing the length of the interval in periodic interval is a trade-off between costs for higher safety stocks or lower customer service level on the one hand and efficiency and cost downstream the supply chain due to administrative processes, transportation and production costs on the other hand.

According to Johnson (2003), the reduction of production lead time, also called manufacturing throughput time, reduces forecast error. The author provides a framework for reducing the production lead time. To reach a reduction in product lead time, he suggests to reduce:

1) setup times or the number of setups required,
2) processing times including scrap, rework, the number of operations needed and the time for the operations,
3) move times or the number of moves,

4) and/or waiting times through reduction in step 1,2 and/or 3, the right batch sizing, reducing arrival variability of orders, reducing machine utilization, increasing available resources or reducing the number of queues (Johnson, 2003).

Hopp et al. (1990) also suggest strategies to reduce the production lead time, like quality management, reducing the WIP, reducing the setup time, splitting batches, introducing transfer batches, maintaining shorter queues, smoothing the work flow and eliminating manufacturing variability.

The method of transportation influences the time it takes to transport products from the supplier to the inventory in the warehouse. Inventory management decisions like deployment strategies, control policies and safety stock settings influence the choice of transportation methods (Ganeshan & Harrison, 2002). "[T]he best choice of mode is often found by trading-off the cost of using the particular mode of transport with the indirect cost of inventory associated with that mode" (Ganeshan & Harrison, 2002, p. 3). Air transportation is often fast, reliable and requires low safety stocks, but is also quite expensive. Shipping via sea is much less expensive, but requires higher safety stocks due to longer lead times and unreliability (Ganeshan & Harrison, 2002). Naturally, the choice of transportation mode also depends on the geographic locations of suppliers and warehouses, the size and frequency of the shipments and the desired customer service level (Ganeshan & Harrison, 2002).

To find a good trade-off between long lead times of cheap transportation methods and the costs of faster methods, researchers suggest dual sourcing as a solution. This means that companies should "get the bulk of their materials from a cheaper *regular* supplier at a lower cost (and longer lead time) but turn to premium *expedited* channels when needed" (Veeraraghaven & Scheller-Wolf, 2008, p. 850).

3.3.2 Late customization

Next, we discuss forms of postponement strategies, also called late customization. Brown et. al (2000) state that companies from several industries like semiconductor and automobile industries "have been able to delay the point of product differentiation, either by standardizing some components or processes or by moving the customization steps to downstream sites, such as distribution centers or retail channels" (p. 66). The authors describe product postponement as delaying the customization of the product through using standardized components, that can be customized at a later stage in the chain, when the demand for the product is more known (Brown, Lee, & Petrakian, 2000). "In process postponement, the firm designs the manufacturing and distribution processes so that it can delay product differentiation, often by moving the push-pull boundary or decoupling point toward the final customer" (Brown, Lee, & Petrakian, 2000, pp. 67-68).

Research proves for different industries the positive effects of postponement strategies on customer service levels and inventory levels. Brown et. al (2002) describe the reduction of inventory levels and better service levels due to product and process postponement for the semiconductor firm Xilinx. Whitney (1993) describes a company in the automobile industry that changes the design of their product, which enables them to react more flexibly to demand by first producing standard parts and then customizing the products later in the production. Christopher (as cited in (Huanga & Li, 2008)) describes that advantages in postponement strategies are reduction in inventory, increased flexibility (because a standardized component can be used for several finished goods) and the ease of forecasting on a more generic level. Furthermore, we have already seen that postponement strategies help shorten the lead time to the final customer, which reduces the need for long term forecasts and can rely more on real time demand. Therefore, we find it interesting to explore options for late customization for the case study firm.

3.4 Simulation

Simulation is a tool to analyze different settings and controls for a supply chain by generating a mathematical model for the real system and running experiments with different control settings in the model, which can be done in seconds for several years of simulation time (Axsäter, 2006). Although it has limitations due to making assumptions for the model and using historical data and experiences, simulation can give valuable insights at low cost and time compared to real world experiments (Axsäter, 2006). Therefore, we will use simulation as a tool for analyzing different settings for lead times and inventory control to gain more insight in what effect those changes can have on the customer service level.

3.5 Conclusions

In this chapter, we have performed a literature study to find several factors that influence the customer service level. First of all, due to the logic of the calculation, to reach a higher order fill rate, a company should then prioritize the order lines with a small number of products. However, sales and profitability targets give a higher prioritization to customers who order larger amounts.

High uncertainty in demand combined with the customers' expectation of fast delivery lead times and highly customized products is a big challenge for companies nowadays. Often this uncertainty leads to order variation throughout the supply chain and the bullwhip effect. There are several ways to cope with demand uncertainty, like improving forecast accuracy, increasing safety stocks or reducing the number of product variants. Reducing the lead time and review period while keeping the safety stock at the same level will increase the customer service level.

Regarding the scope of the case study business, we decide to exclude the opportunities in improving the forecast accuracy and the safety stock settings and instead focus on ways to act more flexibly and rapidly to demand changes. Options for synchronization, order batching and echelon-based inventory control have already been evaluated and partly implemented by the company. Possibilities for reducing the number of product variants is left for further research since it is only relevant for one of the two studied product groups. Therefore, we will focus on lead time reduction and late customization and address simulation as a tool to measure the effect of a certain setting on the customer service level.

4 Solution design

In this chapter, we evaluate which types of solutions are most suitable to research for the case study business at Philips. We develop solutions for lead time reduction and late customization through ideas we got from literature or ideas from stakeholders at Philips. For each idea, we evaluate the relevance, the ease of implementation and the expected impact on the CSL-availability through stakeholder discussions. Then we choose a few types of solutions for each product group to investigate in more detail. In this chapter, we thus answer research question 4.

4.1 Lead time reduction

The transportation lead time is in most of the cases the biggest contributor to the total replenishment lead time, up to seven weeks for product group A and six weeks for product group B. Therefore, the biggest opportunity lies in reducing the transportation lead times. As discussed in the previous chapter, factors which influence the transportation lead time are the mode of transportation and the distance of the factory to the warehouse. Let us discuss opportunities in reduction of those factors for product group A first and for product group B second.

4.1.1 Transportation lead time reduction product group A

Air Freight North America

The products for North America are currently transported by road and sea. Only for urgent matters Air Freight is used. Air Freight reduces the transportation time drastically from up to seven weeks to one week. The actual transportation time via air is one day, but due to administrative processes, the actual lead time is one week. Therefore it can have a great positive effect on the CSL-availability to use air instead of sea transportation. Air Freight is currently only used for firefighting due to the high transportation costs. However, it is unclear if an improvement in the CSL-availability will outweigh the added costs for Air Freight. Another possibility would be to ship most of the products by sea and if the inventory level runs low to also transport a certain amount of products by air to reduce the chance of a stock out, which is called dual sourcing. We will investigate this solution to see if the CSL-availability improvement can outweigh the higher costs for transportation. Since we reduce the lead time with this solution, the inventory in the pipeline will decrease and therefore we expect to save inventory holding costs. Implementing this solution would not require many changes. Air freight is used already for firefighting and therefore using it more often would only require the right training of the shipping team and probably adjustments in the replenishment lead time, which will have to be communicated to the responsible planners. We will therefore neglect any initial costs for implementing the solution.

Production North America

All products of product group A are currently produced in one factory in Europe. Therefore the distance to e.g. America is really long. A way to shorten the transportation lead time to those far away countries is to open another factory nearby that country. This is of course only suitable for countries with a high demand. Since Philips is a global company and therefore has its factories spread out globally, it could be interesting to research the impact of adding an extra production line for this product group in one of the current factories. We therefore evaluate the effect of lead time reduction of this solution on the customer service level.

Factory warehouse: Supply Center for Europe

An idea for Europe could be to use the factory warehouse, which is currently used as Supply Center for the country where the factory is placed, also as a Supply Center for whole Europe. However this is not desirable since most customers expect very short delivery times for their orders. It could be an idea to transport big orders, which are known far enough in advance, directly to the customer from the factory and not through the RDC of the country. Due to time limitations, we decide to exclude this idea from our analysis and suggest Philips to research this further.

Combined shipping Europe and sharing production plan with shipping department

After production, the finished goods arrive at the factory warehouse. They are not shipped directly, except for a few countries which have daily shipping. Finished goods are first stored in the warehouse. The shipping planner waits until a certain amount of goods arrives from the production facility and then plans a shipment for the next day. 24 hours are reserved for picking and packing. Because the factory also produces other products, truck loads combine different types of finished goods and drive them to one specific warehouse. For Europe, on average two or three trucks are sent out weekly. To reduce the waiting time of finished goods in the factory warehouse, the shipping to the three warehouses in Europe could be combined. Also the waiting time of products would decrease if the production plan was available and followed by the factory. However, to change the process of the shipping, the scope should consist of the whole factory and not only one product group of the factory. In addition to that, the effect on the CSL-availability is expected to be very low, because we will only reach a lead time reduction of a few hours or maybe a day with these kinds of solutions. Therefore we exclude this part of the transportation process from further investigation.

Furthermore, an idea could be to merge some warehouses into a hub. This has already been done recently in Europe and will therefore not be investigated further. The impact that this solution had on the customer service level is unfortunately not measured by Philips, since the main goal of this project was to reduce inventory levels.

4.1.2 Transportation lead time reduction product group B

Air Freight

The transportation lead time is also the biggest factor of the lead time for this product. Products are shipped by sea from Asia to Europe. As discussed earlier, transportation by air can reduce the transportation lead time drastically to one week. Again, the actual lead time for air freight is only one day, but due to administrative processes it takes one week for the product to actually be checked in as available at the warehouse. We will investigate if the costs for Air Freight can be outweighed by an improvement in the CSL-availability.

Supplier in Europe

Another possibility could be to search for a supplier which is closer to the RDCs, i.e. a supplier in Europe. This solution has some implications. Since the product is produced at a supplier, Philips only partly owns the product design and therefore will have to develop a whole new product or buy the rights when switching to another supplier. Furthermore, the case study business already investigated the option to buy at another supplier in Europe, but no supplier could be found in Europe for that type of product since it is more cost effective to produce in Asia. We can still investigate the effect of lead time reduction on the CSL-availability, which can help for future decision making in the choosing of a supplier.

Insourcing

The case study business also investigated if insourcing the product could be a good investment, but since the product is not a core product and there is not enough expertise inside Philips for that type of product it would be too costly.

4.1.3 Production and planning lead time product group A

The lead time for production of product group A is one week at most. The length of the production lead time depends on the production schedule. Orders which are planned first are finished in the beginning of the week and orders which are planned last are finished at the end of the week. When placing an order, planners always assume the worst case scenario of a production lead time of one week. However, priority orders can be finished faster through scheduling them in the beginning of the week.

Throughput time
The production is organized as a line production and the products are produced in batches. Since optimizing the production lead time through reducing processing times, setup times, scrap, rework, move times per part e.g. through automation will only reduce the total lead time with a really small percentage and are constantly worked on within a team at the factory, we will not investigate those possibilities. Also the machine reliability is already researched by an expert team and will not be investigated. Furthermore, possibilities for synchronization will not be investigated, since a lot has been implemented already. However, the information sharing between the shipping team and the factory could be more visible, but that would require the scope of all products of the factory, since truckloads of different products are combined.

Lot and batch sizes
From our literature study in section 3.2.7, we learned that reducing lot sizes for ordering can reduce the bullwhip effect. Furthermore, sending out orders more frequently reduces the total lead time due to lower waiting times before transportation. However there is a trade-off between lead time reduction and costs for transportation. The factory recently tried to deliver in smaller lot sizes, which resulted in higher transportation costs. Due to the small impact on the lead time and high costs, the factory decided to increase the lot sizes again to save costs and therefore we decide to not investigate this option.
Reducing the batch sizes in production can also shorten the production lead time due to shorter waiting times. However there is also a trade-off between batch sizes and setup times. This trade-off is continuously evaluated and optimized by stakeholders at the factory. Therefore, and because we expect a small impact on the lead times, we will not investigate this decision further.

Components
Furthermore, the availability of components can have an impact on the lead time. If components are not available, then the orders will have to wait until the components arrive before production can start. Therefore the material replenishment process can influence the total lead time from placing the order at the factory until the arrival of the product at the warehouse. However, the amount of orders which could not finish production in one week due to unavailability of components are almost zero according to the Inventory Planner Procurement at the factory. On top of that, there have been recent projects helping 2nd tier suppliers to improve their flexibility and lead time. Therefore, we will not investigate improvement possibilities for material replenishment.

Daily replenishment
Safety stocks buffer against uncertainty in demand during the lead time plus review period. Therefore, a shorter review period at the same safety stock level should lead to higher customer service levels. The case study business currently uses a periodic replenishment method, where replenishment sizes are evaluated once a week. Every week on Monday, the demand planner places an order and has to wait one week to place a next order. Thus if new demand information comes in on Tuesday, this information waits one week before being able to react on. Therefore, a more frequent replenishment, like daily replenishment, could let the supply chain react faster to changes in demand. However, the factory now aggregates the demand for early production steps as much as possible and schedules production for one week. Daily changes to that schedule can lead to less capacity due to more changeovers or even require having a daily production schedule instead, which also leads to less capacity.
Furthermore, implementing this solution for a single product group is however not advisable, because the factory produces more products and would then have two different procedures of making production schedules and material replenishment. Also the markets would then have different ways

of forecasting and ordering for these products. Therefore this project would need a bigger scope, namely all product groups that are produced at that factory, and is not further investigated in this research. Moreover, for countries with long transportation lead times, this solution would probably have a rather small effect on the customer service level, because the total lead time of eight weeks plus one week review period is only reduced by a few days. For countries that have already shorter lead times of two weeks, reducing the review period could be more interesting. If this is combined with daily production and daily shipping, the lead time could be reduced to 4 days having a review period of 1 day.

Scheduling

The production schedule also influences the lead time of a product. In the current situation, there is no optimal schedule made. Therefore it can happen that products are scheduled to be produced, but the right components are not available or the workers do not accept the schedule, which results in waiting times and machines standing still or doing additional non-planned changeovers. Therefore, sometimes orders are not finished on time during the week, which can cause a lower customer service level due to late arrival at the RDC. An optimal schedule can help to improve the customer service level. However the impact is expected to be quite low, since the number of failures due to production issues is only 3% according to the CSL availability OTTR report. But having a better production schedule also has other benefits for other KPIs of the factory and is therefore currently implemented.

4.1.4 Production and planning lead time product group B

The production lead time of the supplier of product group B is currently three weeks, however the actual production time is much shorter. The manufacturer of the products of product group B plans some extra time for production to secure against risks in demand fluctuations. The planning lead time is the same as for product group A and has the same issues. We will therefore not discuss the planning lead time again.

Shorten production lead time

The supplier wants to be able to deliver products to Philips within the agreed lead time. Due to demand variability, the supplier keeps safety stock of components, which also have a certain lead time to arrive. Since some components have really long lead times, the supplier would have to hold really high safety stocks to ensure a short lead time for Philips since there is demand variability, which cannot be forecasted accurately. Consequently, the production lead time is a trade-off between costs and the risk of not being able to supply in time and it would be quite risky for the supplier to ensure a shorter lead time, since the demand forecast is not quite accurate. Therefore, we will only investigate what kind of effect a reduction in the lead time would have on the customer service level and will not focus on the implementation.

4.2 Late customization

Product group A

To find opportunities for late customization, we asked several stakeholders, see Appendix A.3, for ideas on this topic. This product group has more than 150 different customizations and has to be packed in packages with different languages, which leads to almost 300 different SKUs. Although late customization seems to be a good opportunity in this case, it is hard to implement due to health restrictions in production. All products are produced and sealed in packages, which are not allowed to be opened again unless they are in an environment, which satisfies the health requirements. Therefore, the possibilities for postponement are limited. One project about postponing the packaging process, which makes a SKU country specific has already been stopped due to high costs. Therefore, we will not investigate this idea further.

Product group B

Late customization can reduce the replenishment lead time to the RDC, through customizing the product at a later stage in the supply chain closer to the RDC. An opportunity for late customization for this product is to pack the products at a later stage in the supply chain closer to the RDC. Therefore we will investigate different options of late packaging in Europe for this product. We investigate two types with high volumes and 3 different country versions. See Figure 4.1 for a state mapping for the late customization solution. The total replenishment lead time is three weeks and consists of one week planning lead time, one week packing lead time and one week transportation lead time.

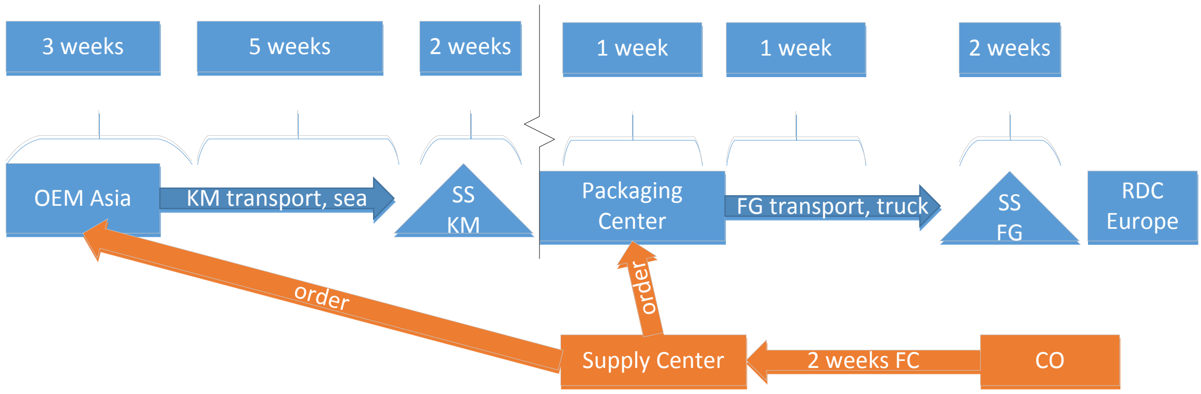

Figure 4.1 State mapping late customization

Option 1 – Key Modules

One idea is to send key modules to a packaging center in Europe and then pack the key modules together to a finished good and send them out to the RDC. However, this idea is not possible in practice, since some key modules are already country specific due to plugs, which are directly connected to the key module. Furthermore, two key modules have to be paired during production. One of those key modules is already country specific due to software, which is installed in a pre-defined language. This makes the combination of the two key modules country specific and leaves us with less flexibility. Therefore, this option will not be further investigated.

Option 2 – Sleeving

Another option for late customization is to send the finished good to a packaging center in a plain box, and put a sleeve around it to make it country specific. There are currently three different country versions. To make this solution cost effective, the product has to have high volumes. Therefore, we investigate the costs for the solution for two products with the highest volume in this category. This solution considers some changes in the supply chain lay-out. The supplier in Asia has to pack the finished good in a plain box instead of the designed box. The packaging center will need to have the right amount of safety stock to be able to deliver the finished goods in two weeks lead time. If a stock-out occurs, the total lead time will increase. Of course, this also depends on the safety stock settings at the warehouse. In this research, we ignore the calculation of the right amount of safety stocks and start out with setting the safety stocks to two weeks. In the cost analysis, we will evaluate what happens to the costs if we change the amount of safety stocks. Furthermore, extra communication is needed between the packing center and the Supply Center (SC).

4.3 Solutions to address

In this chapter, we evaluated different solution ideas on the ease of implementation and on the CSL-availability effect through stakeholder discussions. In the following tables, we summarize the solutions we decided to investigate further and their expected ease of implementation and effect on CSL-availability. The notation "++"means that the solution will be quite easy to implement, i.e. does not need difficult organizational and/or physical changes in the supply chain or that we expect a very high positive impact on CSL-availability. Solutions which reduce the lead time drastically are expected to have a high impact on the CSL-availability. The size of the impact will be evaluated in the next chapters through a simulation model.

Solution A	**Ease of implementation**	**Qualitative impact assessment**	**Conclusion for investigation**
Air Freight North America (4.1.1, p. 27)	++	+	Investigate costs and effect on CSL-availability for: - Using only air freight - Combining shipping and air freight (dual sourcing)
Production North America (4.1.3, p. 28)	-	+	Investigate effect on CSL-availability

Solution B	**Ease of implementa-tion**	**Qualitative impact assessment**	**Conclusion for investigation**
Late customization (Option 2) (4.2, p. 30)	-+	+	Investigate costs and effect on CSL-availability
Supplier in Europe (4.1.2, p. 28)	--	++	Investigate effect on CSL-availability
Shorten production LT (4.1.4, p. 30)	--	-	Investigate effect on CSL-availability
Air Freight (4.1.2, p. 28)	++	++	Investigate costs and effect on CSL-availability for: - Using only air freight - Combining shipping and air freight (dual sourcing)

4.4 Conclusion

In this chapter, we have described different solution designs for lead time reduction and late customization. Each solution is discussed with several stakeholders and we evaluated if that solution is suitable for Philips and what the expected improvement in CSL-availability is.

The biggest opportunity in lead time reduction lies in the transportation lead time. Product group A has a transportation lead time of seven weeks to North America compared to a production lead time of one week. Product group B has a transportation lead time of six weeks compared to a production lead time of three weeks. For both product groups, there is an opportunity to use air freight instead of sea freight to shorten the transportation lead time. The lead time for air freight is one week for both product groups, due to administrative processes. The costs for implementing this solution are not yet investigated by Philips. Therefore, we will investigate the costs for shipping everything via air and a dual sourcing approach, where we ship part of the order via air and the other part via sea in Chapter 6.

Another idea to reduce the transportation lead time is to move the production closer to the country in which the products are sold, by opening a new production site or finding a supplier. For product group A this could be a possibility since there is already a production site close to North America and only a new production line would be needed. We will investigate if the improvement in CSL-availability weighs out the costs of that solution. For product group B, the implementation is not possible, since there is no supplier close to Europe currently and insourcing of the product requires more expertise inside Philips.

We evaluated different methods of reducing the production and planning lead time. Product group A is produced in a factory where also other product groups are produced. This leads to a lot of dependencies, e.g. different product groups are shipped together. Therefore a reduction in production lead time does not always lead to a reduction in the total lead time but could just increase waiting times for the products. Therefore, we recommend to scope an improvement project for all product groups in that factory. Due to time limitations and our earlier defined scope, we do not investigate solutions for production and planning lead time reduction. The production lead time of product group B is three weeks. The supplier of this product does not want to reduce the lead time since it reduces the supplier's risk of not being able to deliver the promised amount of products. Therefore, we will only investigate what kind of effect a reduction in the lead time would have on the customer service level and will not focus on the implementation in Chapter 6.

Finally, we look at opportunities for late customization. For product group A, there were limited options due to health restrictions see section 4.2. After discussions with stakeholders, see Appendix A.3, at the production site and different project owners, we decided that none of the options was suitable for Philips. For product group B, we developed an idea which we call "sleeving". That means that the packing of the product is postponed to a later stage in the supply chain to a packing center in Europe. This reduces the lead time to the markets to two weeks, assuming that we have enough safety stock at the packing center. We investigate the costs for this idea of late customization in Chapter 6.

5 Model development and validation

In this chapter, we build a simulation model, which will help us to evaluate the effect of lead time reduction on the customer service level, the CSL-availability. To develop this model, we follow the procedure described in Chapter 1.7 in Law (2007). We will start by formulating the problem, goal of the study and scope. Then we collect data and define a model. Furthermore we discuss whether the model assumptions are valid. The next step is to construct a computer program and verify and validate the program. After these steps, we will continue with designing experiments in the following chapter.

5.1 Problem formulation

The case study business would like to increase the customer service level by 6% in the next 4 years. To achieve that, the business needs to improve the availability of the products at the warehouses in different markets, which is measured in the CSL-availability performance indicator. Our previous root cause analysis has shown that lead time reduction in the form of late customization or dual sourcing is a potential solution for improving the lead time. The company is interested in having a tool to measure the effect of lead time reduction on the CSL-availability.

5.1.1 Goal of the simulation study

The goal of this study is to develop a simulation model that can test the effect of different lead time and inventory control settings on the CSL-availability in order to find the required setting to ensure a service level improvement of 6%.

5.1.2 Scope

We will develop the simulation model for a specific SKU from product group A for one specific market and for one SKU from product group B for one specific market. In Appendix B.1, we will discuss what changes need to be made for applying the simulation in other contexts. We only use one SKU from both products groups, because there was no more data available for us. We discuss in Chapter 7 what consequences this can have on the conclusions that we get from the results.

The model will be used to simulate the inventory level at the market warehouse (RDC). The inventory is consumed by the customer through placing order lines. In section 5.1.5, we describe how we model the input demand for the model. The CSL-availability is calculated on a weekly basis through dividing all order lines that were fulfilled in full and on time during that week by the total requested order lines during that week. Every week, the replenishment size is calculated depending on the forecast and the inventory level. After one week of planning, we have one week of production and seven weeks of transportation lead time for product group A and three weeks of production and five weeks of transportation for product group B, which results in a total of nine weeks of replenishment lead time including review period for both groups. In addition to that, there is two weeks of safety stock at the RDC for both SKUs. See Figure 5.1 for a state mapping of the product flow. The model only simulates the inventory level and not the flow in the factory or on the truck.

The costs of a solution are dependent on the type of solution and will not be calculated within the model. The results of the cost calculation are presented in Chapter 6.2.

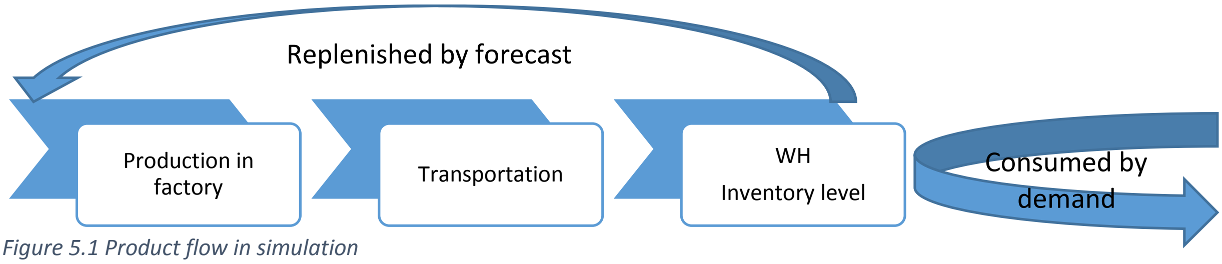

Figure 5.1 Product flow in simulation

5.1.3 Output

The output of the simulation is the CSL-availability level for the certain SKU given the defined input setting for the total replenishment lead time. To calculate the CSL-availability level, we need to know the total number of order lines and the number of order lines which could be delivered directly from inventory.

5.1.4 Input

As input for the simulation, we need information on demand per order line, the number of order lines and the replenishment lead times, i.e. production and transportation lead time of the SKUs to investigate. We collected data from one year, week 32 2015 to week 32 2016, about the number of order lines per week, the number of products per order line, the safety stock settings and the set lead times for the two SKUs. The demand data is used to create input distributions for the simulation model, see section 5.1.5 for further explanation. The safety stock settings will be a static parameter and are described as weeks of average demand. The replenishment lead time will be the experiment factor and is measured in weeks.

5.1.5 Demand distribution

To model the demand input, we have three different choices, namely using the data directly as input, finding an empirical distribution or fitting a theoretical probability distribution on the data, see Chapter 5 in (Law, 2007). Each of the approaches has advantages and disadvantages.
The advantage of using the data directly in the model is that complicated correlations or time-varying parameters are included in the data, which are often complicated to model. However, the data represents historic demand and not what could happen in the future. Also we have only a limited amount of data, which limits the total amount of simulation runs.
Using an empirical input distribution generates demand data that are not in the current data set but that could occur as well. This gives us the opportunity to define the simulation length as long as needed for generating valid results, which do not depend on the initialization of the model. The empirical distribution simply models demand patterns, which are different from but comparable to the real data. This is preferable for analyzing possible variations in the future, but also has the risk of excluding important correlations in the demand.
Finally, fitting a theoretical probability distribution to the data gives a much smoother distribution than using an empirical distribution and also gives us the possibility to generate data which are similar to but not in the current data set. It also gives us the opportunity for easy sensitivity analysis to input parameters. Unfortunately, we also have the risk of ignoring important correlations in the data with this approach.

Fitting a theoretical probability distribution is generally preferred in research. Therefore we start with this approach for modeling our data. We try to fit a theoretical probability distribution on the number of order lines per week and on the amount of order lines per order line for each of the two SKUs. Note that we assume independence between the number of order lines and the amount of products per order line in this scenario. The detailed procedure of finding a distribution is described in Appendix B.2. The distribution of the amount of products per order line for both SKUs is however heavily skewed to the left and has a great variation, e.g. for the SKU of product group B, 90% of all order lines have an amount of smaller than 10 units per order lines and the other 10% is almost evenly distributed between 10 and 2000. Therefore it is impossible to fit a theoretical probability distribution and we decide to use an empirical distribution for that data. For the SKU of product group A, the number of order lines per week is estimated by a normal distribution $\mathcal{N}(\mu, \sigma^2)$ with parameters $\mu = 20$ and $\sigma^2 = 10$. For the demand per order line we describe the empirical distribution in Appendix B.3.2. For the SKU of product group B, the number of order lines per week is estimated by a normal distribution $\mathcal{N}(\mu, \sigma^2)$ with parameters $\mu = 190$ and $\sigma^2 = 31$. For the demand per order line we describe the empirical distribution in Appendix B.4.2.

We test the validity of the input distribution by simulating the demand and comparing the results to the real data. For this model run, we use a lead time of nine weeks and a safety stock of two weeks as in practice. The initial stock is two weeks of average demand and the first nine weeks we expect replenishment orders of average demand. Since we want to be independent of these initial settings, we let the simulation run for 5000 weeks. In the following tables, we compare the average and standard deviation of the total demand per week, the number of order lines per week and the amount of products per order line for each SKU, see Table 5.1 and Table 5.2. Note that for the SKU of product group B we first filtered out weeks from which we knew that there were promotions. For product group A, those weeks were not known and therefore we considered all data.

Table 5.1 Input distribution validation SKU A

SKU A	Data	Model output	Percentage of deviation
Average total demand per week	3736	3839	3%
Std. dev. total demand per week	1997	2333	17%
Average # order lines	20	20	0%
Std. dev. # order lines	10	10	0%
Average # products per order line	190	189	-1%
Std. dev. # products per order line	344	344	0%

Table 5.2 Input distribution validation SKU B

SKU B	Data	Model output	Percentage of deviation
Average total demand per week	2512	2537	1%
Std. dev. total demand per week	1235	1005	-19%
Average # order lines	190	190	0%
Std. dev. # order lines	31	31	0%
Average # products per order line	13	13	0%
Std. dev. # products per order line	65	65	0%

We can see in Table 5.1 and Table 5.2, that for both SKUs most of the values from the model output only deviate slightly from the data that we want to model, namely 0%-3%. However for both SKUs, the standard deviation of the total demand per week deviates around almost 20% from the original data. Making the empirical distribution more accurate by assigning smaller intervals changed the standard deviation by only 1-2%. In our approach for modeling the demand, we assume independence between the number of order lines and the amount of products per order line. Since the standard deviation of the modeled demand deviates that much from the real demand, we expect that there is some correlation between the number of order lines per week and the amount of products per order line. We test some assumptions about that in Appendix B.4.4 and Appendix B.3.4 for both SKUs. For example, one could think that the demand of the previous week influences the demand of the current week. Another example is that the amount of products per order line can be dependent on the total number of order lines.

For the SKU of product group A, we see from the demand data that the correlation between the average amount of products per order line and the number of order lines are negatively correlated with a value of -0,46. So if the number of order lines in one week is high, then the amount of products on an order line is likely be low on average. That would mean that customers either place one order line with a large amount of products in one week or they split their order into multiple order lines with a smaller amount of products on an order line. With this insight, we start modeling the demand for this SKU differently. The distribution for the number of order lines stays the same, but we make the distribution of the amount of products per order line dependent of the amount of order lines. Since it is way too time consuming to create an empirical distribution of the amount of products on an order line for every possible amount of order lines, we make two different distributions, namely one for if the number of order lines is above average and one for if the number of order lines is below or equal to the average. The two empirical distributions can be found in Appendix B.3.1. Table 5.3 shows the demand output of the model after these changes and compares it with the real demand. We see that the difference between the standard deviation of the real and the modeled demand is only 5%. We decide to use this demand model approach for this SKU. The slightly lower standard deviation however lets us expect that the customer service level output of the model will be slightly higher than in reality.

Table 5.3 Second input distribution validation SKU A

SKU A	Data	Model output	Percentage of deviation
Average total demand per week	3736	3749	0%
Std. dev. total demand per week	1997	1892	-5%
Average # order lines	20	20	0%
Std. dev. # order lines	10	10	0%
Average # products per order line	190	185	-3%
Std. dev. # products per order line	344	330	-4%

The SKU of product group B however shows no such correlations as we have seen for the SKU of product group A, see Appendix B.4.4. Therefore, we try to use the data directly as input for the model. This allows us to include the correlations, which are hard to find, but also limits the amount of data we have to simulate. We also decide to include the weeks with promotions again, since this gives us more available data and a more valid representation of the reality, since there are probably correlations between the weeks and the amount of products per week. When we run the model and compare the outcome for the CSL-availability, we do not see a big difference. The differences that we see can be explained by the fact that we do not use weeks of promotion in our theoretical and empirical distribution, but include promotions when we use the demand data directly as input. See Appendix B.4.4 and Appendix B.6 for a comparison between the two demand models. Using a theoretical and empirical distribution for a simulation is preferred to using the data directly, because it gives us the possibility to generate data which are similar to but not in the current data set. Therefore, we decide to use the theoretical and empirical distributions as demand input in the following.

5.1.6 Intermezzo: Big order lines vs. small order lines

When analyzing the demand data, there was one thing that was noticeable. We saw that the CSL-availability is lower for order lines with a larger amount of products per order line. In Table 5.4, we show the CSL-availability for different numbers of products per order line.

Table 5.4 CSL-availability for different order line sizes

Product group A		Product group B	
Products per Order line	CSL availability	Products per Order line	CSL-availability
1 to 100	89,0%	1 to 99	82,9%
100 to 400	82,8%	100 to 399	77,1%
400+	79,2%	400+	66,2%
600+	74,7%	600+	64,5%

When looking at the CSL-availability per customer, we see that customers with a CSL-availability of lower than 80% order together 56% of the total volume that is ordered for SKU A and 67% of the total volume that is ordered for SKU B. This, together with the fact that the number of products per order line is heavily skewed and spread out, lets us suggest to research if handling big order lines in a different way than small order lines. An idea could be to serve big order lines directly from the factory, since demand is often known a lot earlier than the demand for small order lines. Due to time limitations, we decide to exclude this idea from our analysis and suggest Philips to research this further.

5.2 Model definition

In this section, we describe the logic of the simulation program and the assumptions we made. We also discuss the validity and the influence of the chosen assumptions on the model.

5.2.1 Simulation logic

In this section, we define the logic of the simulation model. All assumptions stated here are further discussed in the next section.

For the simulation model, we define three events stated below. For a flow chart of this model, see Appendix B.7. The three events happen in the stated order during one week. Then the program goes to the next week and the three events occur again. This repeats until the program has reached a pre-defined number of weeks. The three events are:

1) Arrival of a forecast replenishment at the CO inventory at the beginning of the week
2) Demand for the product in order lines throughout the week
3) Inventory evaluation and possible ordering at the end of the week

Each week, these three events happen in this order. At the beginning of the week, we evaluate if a forecast replenishment arrives at the CO inventory and the inventory level will be incremented by the ordered amount. We ignore lead time variability and the lead times for order handling on arrival, and assume that there will be always enough capacity in the warehouse. Then the program moves to the next event.

Throughout the week, customers arrive and will order products. The program reads the available input data of the demand and assigns the number of order lines and the amount of products per order line in the given order. We do not use prioritization of customers and just serve the order lines first come first serve. Then, order line for order line, the inventory level is decremented by the amount of products of the specific order line and it is immediately checked if we were able to fulfill that order line in full to be able to measure the CSL-availability after the simulation. If an order line cannot be delivered in full, then it is delivered partly. All orders which could not be fulfilled are backordered and delivered immediately when stock is available. An order line can only count as a hit once.

In the last event, the inventory evaluation event evaluates how much we will have to order and schedules the arrival of that order. The inventory is reviewed on a weekly basis. A replenishment order will be placed when there is at least one product needed. Only full pallets can be shipped and therefore the size of the replenishment order is rounded up according to the pallet size. The replenishment size will be calculated by subtracting the inventory level and the products in pipeline from the safety stock and the forecasted demand over the lead time:

$$R[i + LT] = \left(\sum_{j=i+1}^{i+LT} F_j\right) + SS - \left(\sum_{j=(i+1)}^{i+LT-1} PS_j\right) - I$$

where $R[i]$ is the replenishment size for week i, LT is the replenishment lead time, F_j is the forecast for week j, SS is the safety stock, PS_j is the stock in the pipeline which arrives in week j and I is the current inventory level, which is the on-hand inventory minus backorders. For simplicity reasons, we decide to use as a forecast F_j for all weeks j the weekly average demand of the period week 32 2015 to week 32 2016. As soon as one product is needed we will replenish. The minimum order quantity is one pallet, which contains 2544 products for the SKU of product group A and 180 products for the SKU of product group B. The replenishment size $R[i]$ is always rounded up to the next multiple of 2544 or 180 respectively. The safety stock is defined as a safety lead time of two weeks, which means that the forecasted demand of the two weeks after the replenishment horizon should also lie on stock. In this case it means that the safety stock is two times the average weekly demand of the period week 32 2015 to week 32 2016.

5.2.2 Assumptions

In this section we give an overview of our model assumptions. These assumptions are validated with several stakeholders that we report in Appendix A.3. In the next section, we discuss the validity and impact on the output of these assumptions.

- Every CO has its own inventory (inventory cannot be shared between COs).
- The warehouse has infinite storage capacity and goods receipt capacity.
- Customer orders are served First Come First Serve per week. That means that the stock will not be reserved for some weeks for a certain customer.
- A customer order is always in multiples of twelve products for product group A or three products for product group B.
- If an order line cannot be delivered in full, then part of the order is delivered and other part is backordered.
- All other unfulfilled demand is also backordered and backorders are served FCFS in the next week. However, for the SKU of product group A, if an order line cannot be (partially) fulfilled, there is a 70% chance that the order line is cancelled completely by the customer and does not need to be delivered, not even partly.
- Weeks with special promotions are treated the same as other weeks for the SKU of product group A. For the SKU of product group B, we filter out promotion weeks.
- The replenishment lead time (production and transportation lead time) does not vary. The shipping of the order is done weekly and the full order can be shipped. Also no incidental Air Freight is used to shorten the lead time when the inventory levels get low.
- The factory of the SKU of product group A has unlimited capacity. The factory of SKU of product group B has a maximum capacity of 5040 products per week. If an order exceeds this amount, then only 5040 products will be shipped and nothing is backordered.
- All components are always available.
- As soon as one product is needed we will replenish. Replenishment takes places in multiples of one pallet size.

5.2.3 Discussion of validity

In the previous section, we formulated some assumptions for our model. These assumptions have influence on the performance and output of the model. In reality, not all of these assumptions are always true. Therefore, we will discuss how these assumptions can influence the output of the model and how valid the model is compared to reality.

Every CO has its own inventory (inventory cannot be shared between COs).

The market we chose for the SKU of product group A has two physical stock locations in reality whereas we model only one stock location, since the available demand data is aggregated for both stock locations. When in reality one stock location is out of stock and the other still has stock it can lead to CSL-availability hits at one stock location, whereas in the model it would not lead to a CSL-availability hit since there is still enough total stock. Of course in reality sometimes lateral trans-shipments are possible, but we still expect that our model gives a much higher CSL-availability than in practice for this SKU.

We assume that every CO has its own inventory and that COs cannot share or exchange parts of their inventory. In practice this is not always true. The market of the SKU of product group B we chose to model shares a warehouse with two other markets. They all have stock allocated to their market but also there is some safety stock unallocated which can be shared between the three markets. Also when one market has shortage and there is no more unallocated safety stock left, then the market can ask for additional inventory from other markets if available. However, the SKUs of product group B cannot be shared between markets due to language regulations. Therefore, this assumption will not influence the OTTR of model for the chosen SKU compared to practice. If we choose a SKU of another product group, then it can happen that the model will give a slightly smaller CSL-availability than in reality, because in reality markets can have some extra buffer stock against shortages. However the impact for the chosen market will be small, because it sells way more than the other two markets and therefore the extra safety stock is not big enough to save a lot of order lines from getting hit.

The warehouse has infinite storage capacity and goods receipt capacity.

Next, we assume that the warehouse has infinite storage capacity and goods receipt capacity. These assumptions were discussed with the warehouses in question and can be assumed to be valid. If the warehouse is full, then there are possibilities to rent extra space to store the products. This leads to extra costs, which we will ignore in this analysis for simplicity reasons. This assumption has therefore no impact on the model CSL-availability.

Customer orders are served First Come First Serve per week. That means that the stock will not be reserved for some weeks for a certain customer.

The following assumption is that the customers are served with a first come first serve rule. That means that we serve the customers in a week in order of when the customer order was known in the system. In reality, each customer is categorized in different prioritization categories, which depends on the product and market. In general, promotions and special deals always have the highest priority. When order lines of a promotion are fulfilled, it can happen that our stock levels are temporarily low. Therefore there will be a higher chance of having a shortage and therefore more hits on the CSL-availability.

Another prioritization rule is that customers, whose inventory is managed by Philips, have a higher priority than other customers. For those customers, the market will decide when and how much to replenish. If the customer decides to place an additional order on its own then he has the same prioritization as all other customers. Also there can be some qualitative prioritization rules which are customer- and market-specific.

Since the prioritization rules are quite complicated and market- and customer-specific it is really hard to include those rules into the model. For simplicity reasons, we decide to follow the simple first come first serve rule. In the sensitivity analysis in section 5.4, we test what effect different priority rules have on the customer service level. If mostly order lines with a high amount of product are preferred, then the CSL-availability is lower than that if mostly order lines with a small amount of products are preferred. However, if sometimes small order lines and sometimes big order lines are preferred, the CSL-availability is not clearly higher or lower than using a first come first serve strategy.

We also exclude possible stock reservations, which means that for a certain customer, part of the stock can be reserved for several weeks, so that the market can be sure that the customer gets his order on time and in full. How often and when this is done is not measured and therefore this effect cannot be included in the model. Of course, this means that less stock is available for the rest of the customers and can therefore increase the risk of lower customer service level. This means that our model can give a higher service level than in reality. Unfortunately, due to the logic of our simulation model, this effect is hard to include in the model and to test the sensitivity of the output to this factor. We generate each week the demand for that current week and do not know how long before the delivery date the order has been placed. We do not have any information on how long stock should be reserved in the system and for which order lines stock will be reserved and for which not.

A customer order is always in multiples of twelve for product group A or three for product group B. Furthermore, we assume that orders are only delivered in multiples of twelve for the SKU of product group A and in multiples of three for the SKU of product group B. In our demand model we round all customer orders up to these multiples. This is an agreement for the chosen product groups between Philips and the customer. The products are packed in multiples of three in shipping boxes and it is not desirable to break those boxes. Only in special cases, for example when sending out samples, those boxes are opened and single products are sent. Since these single cases will have almost no influence on the output and performance of the model, we will exclude these from further analysis.

If an order line cannot be delivered in full, then part of the order is delivered and other part is backordered.
And
All other unfulfilled demand is also backordered and backorders are served first in the next week. However, for the SKU of product group A there is a 70% chance that the order line is cancelled completely by the customer and does not need to be delivered, not even partly.
These two assumptions are about handling a stock-out or almost stock-out. When an order arrives and we are not able to ship the full order, then we will send out the part of the order we have and the rest is backordered. This is sometimes done in practice, but it depends on the customer. In practice, when a big order cannot be fulfilled immediately and the customer chooses to have the order delivered later when we have enough stock, then following orders can possibly still be served in full, which prevents a CSL-availability hit from happening. In our model, this is not possible. Therefore, we expect to have some weeks where the customer service level in the model will be lower than it would have been in practice.

Other orders of the SKU of product group B, which cannot be fulfilled due to a stock-out are always backordered. This also happens in practice most of the time. It can happen that a customer cancels the order and places it again a few days later or order at a competitor, but these situations are very rare and can be left out from the model without having a big influence on the outcome. For the SKU of product group A, 70% of the orders, which cannot be delivered in full on time are cancelled. We can see this in our demand data since we know when zero products are delivered and an extra column reports the reason, which can be cancellation or e.g. double order. A cancellation still counts as a CSL-availability hit, but leaves us with more inventory for next order lines or with less backorders for the next week.

All backordered demand is served directly the next week as soon as stock arrives, so it has the highest priority. This is also done in practice except for a few exceptions.

Weeks with special promotions are treated the same as other weeks for the SKU of product group A. For the SKU of product group B, we filter out promotion weeks.
Our demand data includes also weeks with promotions. That means that a certain customer or several customers can have a special offer to the consumer for which the customer(s) will need more products than usual. For the SKU of product group A, there is no information about promotions and therefore we decide to not take this effect into consideration in our analysis.
For the SKU of product group B we decide to exclude promotion weeks in the demand input. In practice, those promotions are often known on the forehand and planned carefully with the customer and can be taken into account in the forecast. It can happen that safety stock is reduced due to promotions and special deals. Although promotions are often planned with customers a long time before the promotion starts, it happens that customers want a promotion quite quickly. Due to the long lead time, the forecast within that time cannot be changed. Therefore either the promotion has to be cancelled or if enough safety stock is expected to be available, markets can choose to accept the promotion anyways. This can lead to reduced stock levels after promotions and therefore a higher chance of CSL availability hits when demand is higher than expected. Therefore we expect our model to give a higher customer service level after a promotion week. However promotions do not occur that often for that product, i.e. around 7 times a year and therefore we can expect the total effect on the model as limited.

The replenishment lead time (production and transportation lead time) does not vary. The shipping of the order is done weekly and the full order can be shipped. Also no incidental Air Freight is used to shorten the lead time when the inventory levels get low.
Our next assumption is that the replenishment lead time (production and transportation lead time) does not vary. A truck and boat is sent out every week and is able to ship the full order. In practice, the capacity of the truck and ship is never fully used, but low truck loads are sent out weekly. However, it can happen that the transportation lead time is a little longer than expected due to delays on the road or sea. In practice, for the SKU of product group A 10% and respectively for the SKU of product group B 20% of all order lines, which are late, are still delivered in the same week, which means that the replenishment order came in a few days late. Therefore, with the assumption that the replenishment lead time does not vary and the whole order is delivered in the beginning of the week before the demand comes in, the model will give a better customer service level than there is in reality. For simplicity reasons, we keep this assumption, since we do not see big influences on the outcome in the sensitivity analysis in section 5.4.
Furthermore, in practice the company can decide to ship an order via air when the inventory levels run really low and demand increase is expected. This incidental firefighting occurs rarely and is therefore excluded from analysis. It helps in practice to reduce the amount of hits and improve the customer service level. Because this occurs rarely, it will not have much impact on the outcome of the model.

The factory of the SKU of product group A has unlimited capacity. The factory of the SKU of product group B has a maximum capacity of 5040 products per week. If an order exceeds this amount, then only 5040 products will be shipped and nothing is backordered.
The factory, which produces the SKU of product group A, is most of the time able to fulfill the placed replenishment orders. We therefore do not limit the capacity for this SKU. In the model, we find that 2% of the replenishment order exceeds 10.000 products. We will also see later, that this phenomenon does not influence the outcome, see section 5.4.

Furthermore, for product group B the factory has a maximum weekly capacity of 5500 products for all markets which order the product. Since 90% of the total volume of the product is sold in the market we chose to analyze, we assume that the factory can always make 4950 products, which is 90% of the total capacity, for this market. Since 4950 is not a multiple of 180, which is the amount of products on one pallet and therefore the minimum order quantity, we round the maximum capacity to 5040. In reality, it is not guaranteed that we get this amount when we place an order, but it can also happen that the market can get more products when other markets are ordering less. We expect that the negative aspect of not being able to order more is canceled out by the fact that it does not happen that we get promised less if other markets demand is higher. Furthermore, as soon as one product is needed, a whole pallet will be shipped, which is also done in practice. In the model, only 4% of the replenishment orders exceed 5500 products before taking into account capacity restrictions. Therefore, we do not expect a great influence on the overall CSL-availability.

All components are always available.
The amount of problems reported with component availability at the factories for the two SKUs is almost zero, which we verified with the Inventory Planner Procurement at the production site. We therefore assume that all components are always available. If it is not the case in practice, then the capacity of the supplier is limited and it can happen that the replenishment order cannot be delivered in full. However this does not lead to CSL-availability hits immediately, since there is some safety stock at the warehouse. We therefore expect that this assumption will have almost no consequence for the CSL-availability.

As soon as one products is needed we will replenish. Replenishment takes places in multiples of one pallet size.
This is precisely as it is done it practice and should therefore not influence the outcome. If an order size is smaller than a pallet size, we always round up to the next highest multiple of the pallet size.

5.2.4 Conclusion and further approach

Let us now discuss, what we expect from the model compared to reality given the structure of the program and the assumptions we made in the previous sections.

Most of the assumptions we made, simplify our model in comparison to reality in such a way, that we expect our model to give a higher CSL-availability output than in reality. We assume that suppliers can deliver everything we order up to a maximum capacity and that the components are always available. Also everything is delivered on time, which also means that the warehouse can accept all products and has no delays between receiving a delivery and assigning the stock into the system. However, we see that around 15% of the CSL-availability hits of those two products are caused due to supply and delivery issues. Since our model ignores those issues, we expect a higher CSL-availability from our model.

Another aspect which is difficult to model but influences the outcome of our model is the prioritization of customers. First of all, some customers are of such importance that stock reservations are made, which means that the stock is lying ready in the warehouse, but is not available for other customers. In our model, we assume that the stock is always available for all customers and therefore the outcome of the model is expected to be higher than in reality. In addition to that, if an order line with a large amount of products is preferred, it can mean that all following order lines cannot be served directly from stock and lower the CSL-availability.

Finally, a factor that can influence the CSL-availability is promotion combined with forecasting. We use the average demand of the period week 32 2015 to week 32 2016 as forecast. However, for the SKU of product group A, the demand data includes weeks with promotions and therefore the forecast of the model for weeks without promotions is most likely a little higher than the average forecast for those weeks in practice, since promotions are often known to demand planners.

In the following sections, we will run our model and evaluate the outcome with regards to our expectations. Furthermore, we will conduct a sensitivity analysis to test the influence of different parameter settings on the outcome of the model.

5.3 Program construction, verification and validation

The described simulation model is constructed in C# using Microsoft Visual Studio. During and after programming, we verified the model through debugging. For the validation of the program, we discussed the assumptions and our expectations for the model in the previous chapter. We expect the OTTR output of the model to be higher than the OTTR in reality.

In the following table, Table 5.5, we compare the CSL-availability YTD measured in practice to the CSL-availability YTD output from our simulation model for both SKUs. For this model run, we use a lead time of nine weeks and a safety stock of two weeks as in practice. The demand distribution described in Chapter 5.1.5 is used for the SKU of product group A. Real data from week 32 in 2015 to week 32 in 2016 is used as demand input for the SKU of product group B. The initial stock is two weeks of average demand and the first nine weeks we expect replenishment orders of average demand. Since we want to be independent of these initial setting, we let the simulation run for 5000 weeks.

Table 5.5 CSL results

CSL-availability	In practice	Model output
SKU A	86,8%	97,3%
SKU B	82,7%	99,2%

As expected, the CSL-availability output from the model is higher than the OTTR in reality. However the difference between the model and reality is still quite large. This can be caused by various reasons. We discussed some hypotheses of possible reasons for the CSL-availability being higher in the model than in practice in section 5.2.3. For the following factors, we conduct a sensitivity analysis in the next section to understand the influence of these factors on the outcome:

1) Forecast
2) Lead time variability
3) Capacity of the factory
4) Customer prioritization (e.g. FCFS)
5) Order cancellation
6) A combination of those factors

5.4 Sensitivity analysis

In this section, we perform a sensitivity analysis in order to test the robustness of the model by studying the relationship between the input parameters and the output. This can help us understand why the model gives a higher output than the CSL-availability in reality. We test how sensitive the model reacts to forecast variability, replenishment lead time variability, supplier capacity, the order of serving order lines, order cancellation and a combination of those.

For the following simulation runs, we again use 5000 weeks of simulation to be independent of the chosen initial conditions. We start with an initial inventory of one or two weeks safety stock depending on the safety stock parameter and the first nine weeks we expect replenishment orders of average demand. The replenishment lead time is nine weeks for both SKUs. If not stated otherwise, the production capacity for the SKU of product group B is 5040 and the production capacity for the SKU of product group A is unlimited. Furthermore, the order lines are served FCFS and for SKU B there are no order cancellations and for SKU A 70% of all order lines which cannot be fulfilled are cancelled, if not stated otherwise. We test the sensitivity of the model output to different input settings with a normal safety stock of two weeks, called *2 weeks SS* in the following figures, and a reduced safety stock of one week, called *1 week SS* in the following figures, to see if the output is more sensitive to input parameters when stocks are low. For more detailed tables and additional settings see Appendix B.5.

In the following figure, Figure 5.2, we see the output for both SKUs with two safety stock settings, when we change the forecast input. We look at how robust the outcome is to changes in the forecast, like constant under-forecasting or constant over-forecasting. For example, -2% in the figure means that we use a constant forecast which is 2% smaller than the average demand.

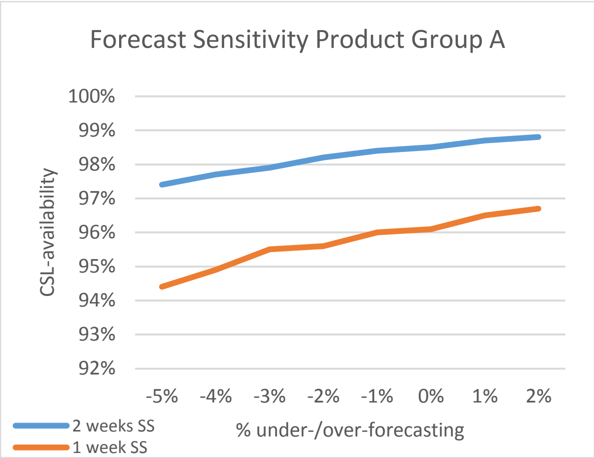

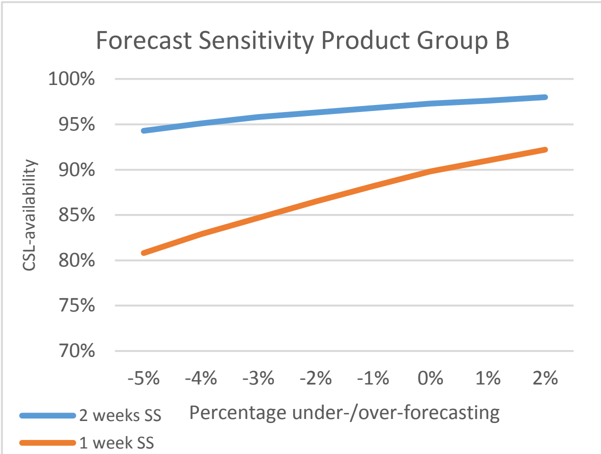

Figure 5.2 Sensitivity to forecast

First of all, when we look at the output of the SKU of product group B in Figure 5.2, we see that reducing the safety stock by one week results in a reduction of the CSL-availability by 7,5%, whereas the reduction in CSL-availability of the SKU of product group A is 2,4%. This shows us that the CSL-availability of SKU B is quite sensitive to reduced stock levels. This can be explained with the fact that orders which cannot be delivered on time are often cancelled for the SKU of product group A. This results in a lower backlog situation than for the SKU of product group B.
Looking at the graph for SKU B, we see that the more under-forecasting, the lower the CSL-availability gets. A constant under-forecast of 5% under the average demand leads to a CSL-availability reduction of 3-9% depending on the amount of safety stock. For the SKU of product group A we also see a reduction in CSL-availability when the forecast is decreased. However, this reduction is less significant compared to the other product group, namely 1-2%, which can be explained again with the fact that a lot of order lines that cannot be delivered are cancelled, which leaves that product group with less backorders. From these graphs, we can conclude that under-forecasting can lead to a lower CSL-availability. This can be part of the reason that our model has a higher CSL-availability, because in practice under-forecasting happens. However, unfortunately we lack the data to get insights in how often this happens.

Next, we test the sensitivity of the output to variability in replenishment lead time. In Figure 5.3, we show the results of varying the percentage of order lines which have to be served from the current stock due to a late arrival of replenishment. We do this for the chances of 20% and 30% for a delivery being late. In our model, we let deliveries be late at random with a chance of 20% or 30%. The predefined percentage of order lines now have to be served from the current stock before the replenishment arrives, which can lead to additional availability hits. We assume for the SKU of product group A, that now no order lines are cancelled, because they are expected to be delivered just a few days late. For these simulation runs, we take the average demand as forecast.

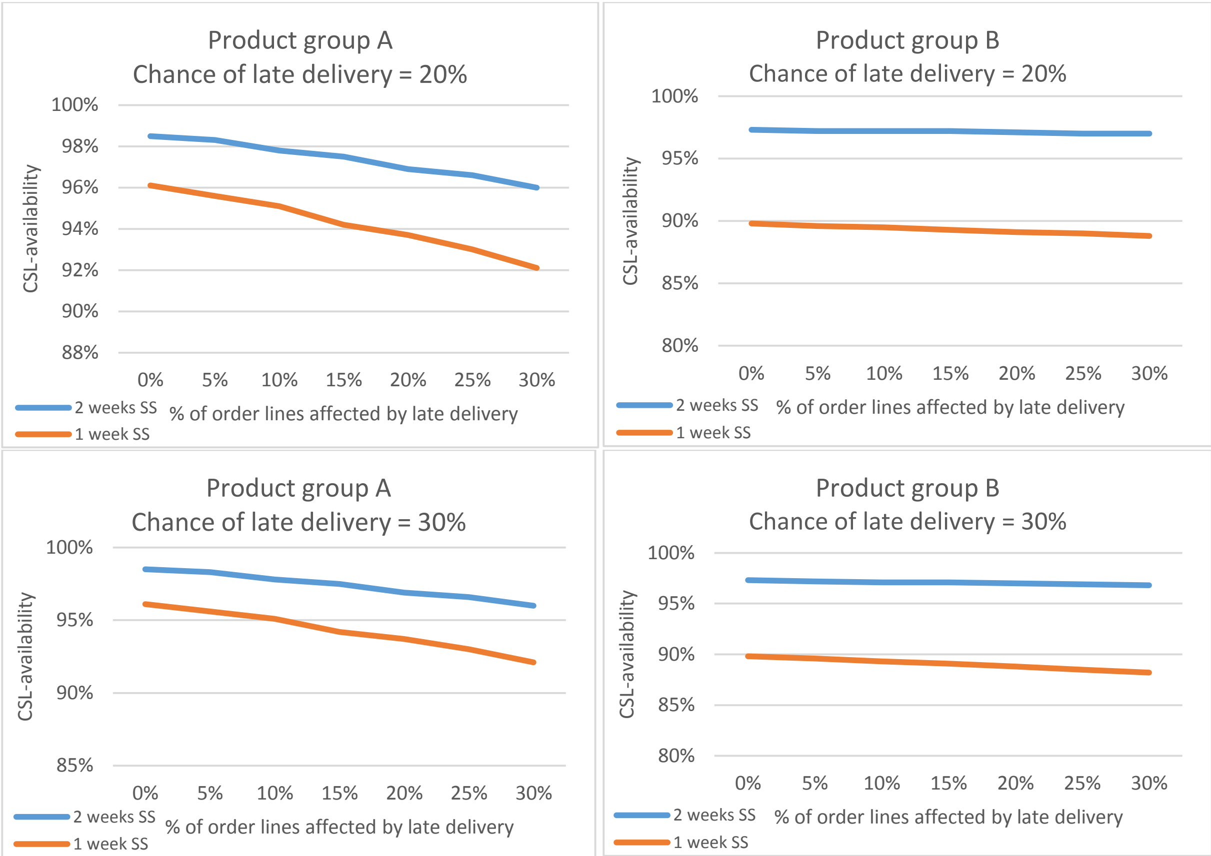

Figure 5.3 Sensitivity to lead time variability

In Figure 5.3, we see that the CSL-availability of the SKU of product group A decreases by 2%-4% or 4%-7% when the chance of products arriving late is 20% or 30% respectively. Nevertheless, the CSL-availability of the SKU of product group B does only decrease by around 1%. This can be explained by the assumption for the SKU of product group A, that the order lines in the beginning of the week are not cancelled and that the number of order lines per week are much smaller than for the SKU of product group B. We therefore can conclude that lead time variability has almost no effect on the CSL-availability output.

Let us now take a look at the production capacity. The capacity of the SKU of product group B is 5040 per week, which we show as 100% in the graph. We estimate the capacity for the SKU of product group A as 10176 per week, which is more than twice the average demand.

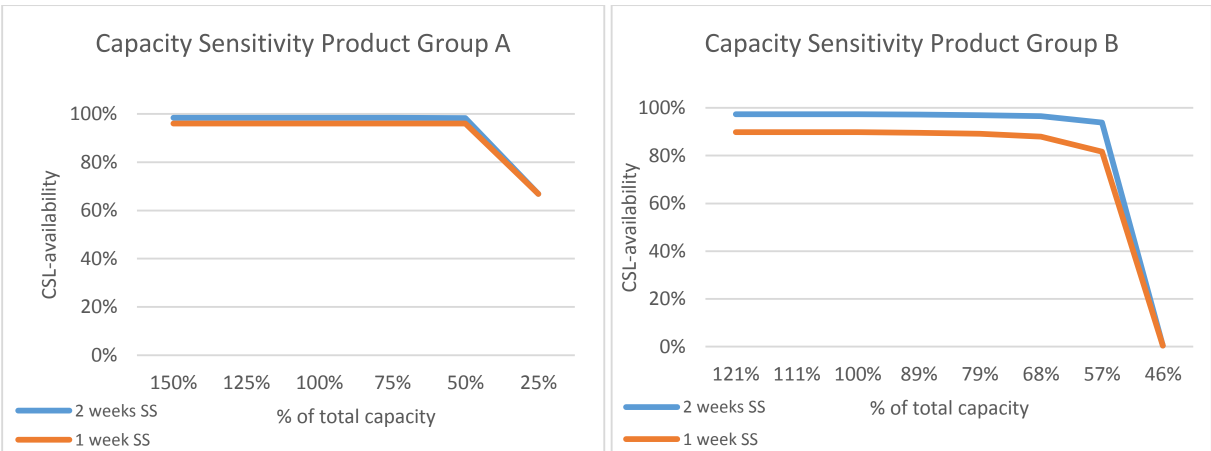

Figure 5.4 Sensitivity to capacity

The results in from show that the CSL-availability only significantly decreases if the capacity falls below 60%. When the capacity falls under 60%, the factory produces less than the average demand per week, which leads to CSL-availability hits. Therefore, we can conclude that the amount of capacity does not influence the CSL-availability as long as it stays above 60%.

Next, we will test how the order of serving customers influences the output of the system. We test our FCFS approach against two sorting methods, descending and ascending. In the descending sorting method, order lines with a large amount of products are served first and the small ones last. For the ascending sorting method it is the other way around. See Figure 5.5 for the outcomes.

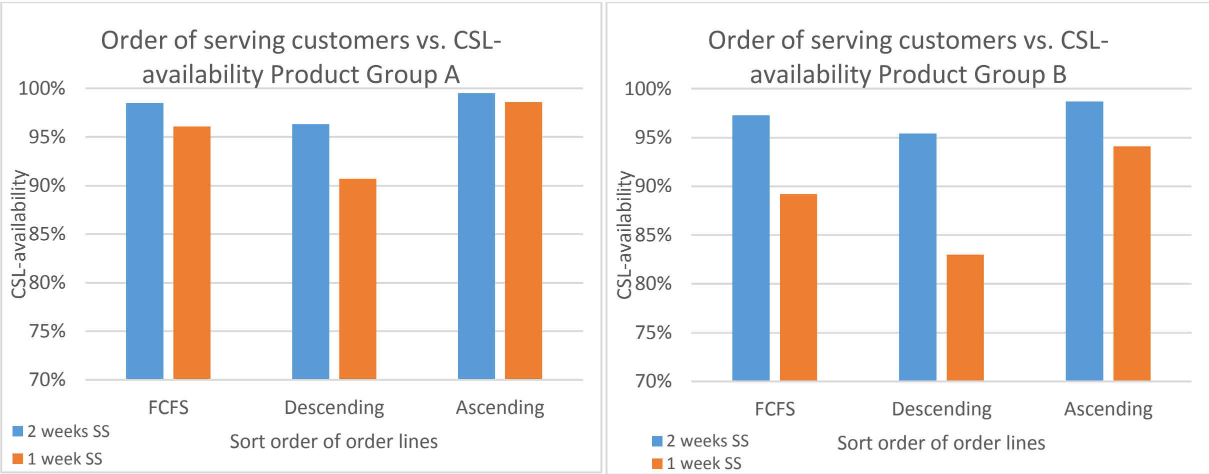

Figure 5.5 Sensitivity to order of serving customers

Due to the logic of the calculation of the CSL-availability, we already expect that the CSL-availability is higher for an ascending than for a descending order. This behavior can also be seen in Figure 5.5. Serving order lines with a large amount of products first lead to a reduction in CSL-availability of 2,4% - 5,4% for the SKU of product group A and 1,9% - 6,2% for the SKU of product group B compared to a FCFS prioritization. Therefore, we can conclude, that the order of serving customers matters significantly.

In the previous tables we have seen that the sensitivity of the output of SKU B to the input parameters was higher than that of SKU A. The reason for this is that 70% of the orders of the SKU of product group A that get an availability hit are cancelled. Therefore, we test how the output of SKU A reacts to changes in the percentage of order cancellations, see Figure 5.6.

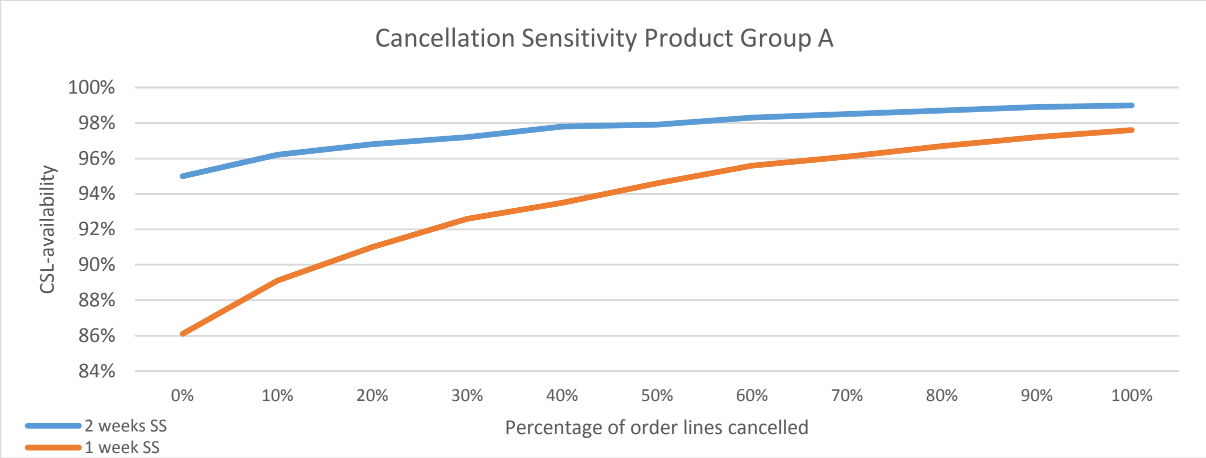

Figure 5.6 Sensitivity to order line cancellations

Figure 5.6 shows that the CSL-availability output decreases when decreasing the chance of cancellation. This is logical since the amount of backorders increases if less orders are cancelled and we remain with less stock. Also the difference between the outputs of two weeks safety stock and one week safety stock increases when decreasing the chance of cancellation of an order line. From this we can conclude that when more customers decide to keep their orders as backorder and have them delivered as soon as products are available, the CSL-availability is more sensitive to low stock situations.

Finally, we test how a combination of the previous factors influences the outcome. For both SKUs, we let the forecast be 2% less than the average demand, the demand is served FCFS, the chance of a product delivered late is 30% and 30% of the order lines are affected by the late delivery. The SKU of product group A has a 60% chance of products being cancelled instead of a 70% chance. For the SKU of product group B, we limit the production capacity to 4500 instead of 5040 products. The results are shown in Table 5.6 and Table 5.7 for safety stocks of one week and two weeks.

Table 5.6 Sensitivity analysis: combination of factors for SKU A

CSL-availability output	**SKU A**	
Combination	2 weeks SS	1 week SS
Initial assumptions	98,5%	96,1%
- 2% under-forecasting - Cancellation: 60% - Chance of delivering late: 30% - #OL affected when late: 30%	92,4%	86,6%

Table 5.7 Sensitivity analysis: combination of factors for SKU B

CSL-availability output	**SKU B**	
Combination	2 weeks SS	1 week SS
Initial assumptions	97,3%	89,2%
- 2% under-forecasting - Max Replenishment Cap: 4500 - Chance of delivering late: 30% - #OL affected when late: 30%	95,5%	84,3%

If we sum up all drops in CSL-availability for the SKU of product group A that we have seen for 2% under-forecasting (0,3% - 0,5% drop), 60% order cancellation (0,2% - 0,5% drop) and a chance of delivering late of 30% with 30% of the order lines affected when late (4,4% - 6,7% drop), we would expect a drop of 4,9% - 7,7% if the factors do not strengthen each other. The decrease in CSL-availability is 1% - 2% more than only adding up the factors. If we sum up all drops in CSL-availability for the SKU of product group B that we have seen for 2% under-forecasting (1% - 3,3% drop), a production capacity of 4500 (0,1% - 0,2% drop) and a chance of delivering late of 30% with 30% of the order lines affected when late (0,5% - 1,6% drop), we would expect a drop of 1,6% - 5,1% if the factors do not strengthen each other. Here, the decrease in CSL-availability is only 0,2% more than only adding up the factors. Therefore, a combination of the factors seem to only add up and not to intensify the decrease in the outcome.

5.5 Conclusions

In this chapter, we have built a simulation model for two different SKUs. With this model we want to analyze how solutions like lead time reduction and dual sourcing help improve the CSL-availability. We here describe our most important finding about modeling the demand, the assumptions we made and the sensitivity analysis.

The scope for the simulation study is one SKU from product group A and one SKU from product group B. We only use one SKU from both products groups, because there was no more data available for us. We discuss in Chapter 7 what consequences this can have on the conclusions that we get from the results. The demand is modeled by fitting a theoretical distribution for the number of order lines per week and an empirical distribution for the amount of products per order line. For the SKU of product group A, the number of order lines per week is estimated by a normal distribution $\mathcal{N}(\mu, \sigma^2)$ with parameters $\mu = 20$ and $\sigma^2 = 10$. For the demand per order line we describe two empirical distributions in Appendix B.3.1. For the SKU of product group B, the number of order lines per week is estimated by a normal distribution $\mathcal{N}(\mu, \sigma^2)$ with parameters $\mu = 190$ and $\sigma^2 = 31$. For the demand per order line we describe the empirical distribution in Appendix B.4.2.

A full list of our assumptions, which are validated with the corresponding stakeholders, see Appendix A.3, is given in section 5.2.2. We now summarize two important assumptions, which we expect to have an influence on the outcome. First of all, the demand data for the SKU of product group A is aggregated over two different stock locations. Of course, with two stock locations, the risk of getting out of stock is higher than for one stock location. In our model, we use only one stock location, which means that out CSL-availability output of the model will be higher than in reality. However due to the unavailable data for the separate stock locations, it is impossible to analyze the effect of this assumption. Secondly, markets sometimes make stock reservations for customer orders. That means that the stock of an order will be reserved from the moment of ordering and is not available for serving other customer orders. Unfortunately, due to the logic of our simulation model and the unavailability of information on the handling of stock reservations, this effect is hard to include in the model and to test the sensitivity of the output to this factor. However, stock reservations can lower the customer service level in reality and therefore our model gives a higher output.

We have tested the sensitivity of the model to demand forecast, replenishment lead time variability, production capacity, the order of serving customers and order cancellation for two and one weeks of safety stock. We list our findings about these factors.

- We have seen that reducing safety stocks by one week reduces the CSL-availability, namely 2,4% for the SKU of product group A and 7,5% for the SKU of product group B.

- We can conclude that under-forecasting has a low effect on the decrease of CSL-availability, namely a 2%-under-forecast leads to a reduction of CSL-availability of 0,3% - 0,5% for the SKU of product group A and 1% - 3% for the SKU of product group B.
- We have seen that lead time variability has the biggest effect on the CSL-availability output of the SKU of product group A, namely 4,4% - 6,7%, which is due to the small number of order lines. For the SKU of product group B, the decrease is only 0,5% - 1,6%.
- We have seen that the amount of capacity does not influence the CSL-availability as long as it stays above 60%.
- The order of serving customers plays an important role on the total amount of CSL-availability hits, especially when in a low stock situation. Serving order lines with a large amount of products first lead to a reduction in CSL-availability of 2,4% - 5,4% for the SKU of product group A and 1,9% - 6,2% for the SKU of product group B compared to a FCFS prioritization.
- About order cancellations for the SKU of product group A, we can conclude that the more customers decide to keep their orders as backorder and have them delivered as soon as products are available, the more the CSL-availability is sensitive to low stock situations. Also, less order cancellations lead to less CSL-availability.
- A combination of the factors, with 2% under-forecasting, 60% order cancellation for the SKU of product group A or 4500 production capacity for the SKU of product group B and a chance of delivering late of 30% with 30% of the order lines affected when late, only seem to only add up and not to intensify the decrease in the outcome.

6 Experimental design and results

In this chapter, we perform a theoretical analysis of what effect lead time reduction and dual sourcing can have on the output of our model, and we evaluate the results and costs. For that, we use the distributions described in section 5.1.5 as demand input. This has the advantage that we can create unlimited random demand data, which is based on and somewhat similar to the real demand, instead of using the limited historical data, which only describes what happened in the past. We also keep all assumptions we made in section 5.2.2 for all following experiments.

To begin with, we design the experiments for lead time reduction and dual sourcing. For this we explain what kind of dual sourcing method we use. Then, we describe the type of our simulation, the warm-up period and the number of replications needed. Next, we run the simulation and discuss the results and costs of the solutions. These steps give us the results we need to answer research question 5.

6.1 Experimental design

We start by designing the experiments for lead time reduction. Then we will describe the logic we use for experimenting with dual sourcing and also design experiments for this solution. Before we start giving results, we discuss the type of our simulation and define the warm-up period, the length of one simulation run and the number of replications needed.

6.1.1 Replenishment lead time reduction

The first solution method we want to assess is simple lead time reduction. We are interested in how the CSL-availability is influenced by a reduction in lead time, which can be achieved by the described solutions in Chapter 4, e.g. late customization, shortening production lead time or transporting all products by air. We are also interested in how the effect on the CSL-availability is influenced, when having a low stock situation and therefore we want to test different values of the safety stock. This is also interesting, because when lead time is reduced in practice, the safety stock will most probably be reduced too, since another objective of the company is to minimize inventories.

So, we have two experiment factors, namely the replenishment lead time and the safety stock. We design our experiment using a 2^k factorial design described in Chapter 12.2 in (Law, 2007). To calculate the main effect of a factor, we need to create output for one low level and one high level of each factor, which results in four experiments. Since we are interested in more than two values for the factor replenishment lead time, we design more experiments. The experiments for both SKUs are defined in the following table, Table 6.1.

Table 6.1 Experiment configurations studied for lead time reduction

Experiment	Replenishment Lead Time (in weeks)	Safety stock (in weeks)
1	9	2
2	7	2
3	5	2
4	4	2
5	3	2
6	2	2
7	9	1
8	7	1
9	5	1
10	4	1
11	3	1
12	2	1

6.1.2 Dual sourcing: Method and experiment design

The other solution we want to research for both SKUs is dual sourcing, which in our case means sourcing our products via sea transport regularly and via air transport when stock levels run low. For this approach, we design experiments following the method described in (Veeraraghavan & Scheller-Wolf, 2008).

Let i be the index for the current week in the simulation model. We call all orders shipped via sea regular orders with a lead time l_r of nine weeks and orders shipped via air expedited orders with a lead time l_e of five weeks and let $l = l_r - l_e$. In our simulation model, we calculate the replenishment orders at the end of each week. Therefore I_i is the on-hand stock at the end of week i. X_i^e is the expedited order that we order in week i and it "is based on the on-hand inventory plus the expedited and regular orders that will arrive within l_e periods[…]. This expedited order [X_i^e] tries to restore the expedited inventory position [IP_i^e] to some target parameter level z_e" (Veeraraghavan & Scheller-Wolf, 2008, p. 852). Respectively, X_i^r is the regular order in week i and "is based on the regular inventory position (sum of on-hand inventory and all outstanding orders, including [X_i^e]), [IP_i^r], and tries to restore it to some target parameter level z_r" (Veeraraghavan & Scheller-Wolf, 2008, p. 852). We summarize the notations in Table 6.2.

Table 6.2 Dual sourcing notations

Variable	Description
$\boldsymbol{i}$	Week index
$\boldsymbol{l_r}$	Regular order (sea) lead time
$\boldsymbol{l_e}$	Expedited order (air) lead time
$\boldsymbol{l}$	$l = l_r - l_e$
$\boldsymbol{I_i}$	On-hand stock at the end of week i
$\boldsymbol{X_i^e}$	Expedited order size in week i
$\boldsymbol{IP_i^e}$	Expedited inventory position in week i
$\boldsymbol{z_e}$	Order-up-to-level for expedited orders
$\boldsymbol{X_i^e}$	Regular order size in week i
$\boldsymbol{IP_i^r}$	Regular inventory position in week i
$\boldsymbol{z_r}$	Order-up-to-level for regular orders

We start in our model by adding the expedited orders from five weeks ago and the regular orders from nine weeks ago to our inventory position. Then the demand is subtracted from the inventory position. After that we start planning what we need to order, following the method described in (Veeraraghavan & Scheller-Wolf, 2008). We start by calculating the expedited inventory position, which "is comprised of on-hand inventory and all the orders due to arrive in the next l_e periods" (Veeraraghavan & Scheller-Wolf, 2008, p. 852):

$$IP_i^e = I_i + \left(X_{i-l_e}^e + \cdots + X_{i-1}^e\right) + \left(X_{i-l_r}^r + \cdots + X_{i-l-1}^r\right).$$

Next we calculate the order size of the expedited order. Note that the order size has to be positive, since we will not fly back any products. The expedited order X_i^e arrives in week $i + l_e$ and so will X_{i-l}^r. To not let the inventory position exceed the target level z_e, we take that order into account and calculate X_i^e as follows:

$$X_i^e = Max(0, z_e - IP_i^e - X_{i-l}^r).$$

After the expedited order is made, we calculate the regular inventory position, which "is comprised of on-hand inventory and all the orders that will arrive in the next l_r periods" (Veeraraghavan & Scheller-Wolf, 2008, p. 852):

$$IP_i^r = I_i + \left(X_{i-l_e}^e + \cdots + X_{i-1}^e\right) + \left(X_{i-l_r}^r + \cdots + X_{i-1}^r\right).$$

Finally, we calculate the regular order, which also cannot be less than zero. To calculate the order size we also include information on the expedited order that we just made. The following formula shows how to calculate the regular order size.

$$X_i^r = z_r - (IP_i^r + X_i^e)$$

After understanding the method of dual sourcing, we can start designing our experiments. We again design our experiment using a 2^k factorial design as described in Chapter 12.2 in (Law, 2007). Our factors are the order-up-to-levels z_e and z_r. The values of z_r and z_e are calculated in the following way:

$$z_r = l_r * Weekly\ forecast + 1\ or\ 2\ weeks\ of\ safety\ stock \text{ and}$$
$$z_e = l_e * Weekly\ forecast + x\ weeks\ of\ expedited\ safety\ stock.$$

We test different values for x varying from 0 to 3 with both one and two weeks of safety stock for z_r. See Table 6.3 and Table 6.4 for an overview of the values of $z_e$and z_r for the SKU of product group A and B respectively.

Table 6.3 Experiment configurations studied for dual sourcing for SKU A

Experiment	z_e	z_r
13	26852	42196
14	24934	42196
15	23016	42196
16	21098	42196
17	26852	38360
18	24934	38360
19	23016	38360
20	21098	38360

Table 6.4 Experiment configurations studied for dual sourcing for SKU B

Experiment	z_e	z_r
21	17794	27962
22	16523	27962
23	15252	27962
24	13981	27962
25	17794	25420
26	16523	25420
27	15252	25420
28	13981	25420

6.1.3 Number of Replications and warm-up period

Before running the experiments, we need to define the type of our simulation and find the length of each simulation run. Furthermore, we have to define the length of the warmup period if necessary and the number of replications needed.

First we need to discuss the type of simulation. A simulation can be terminating or non-terminating. The definitions are given in Chapter 9.3 of the book of Law (2007). We define our simulation as a non-terminating simulation "for which there is no natural event E to specify the length of a run" (Law, 2007,

p. 495). We simulate several numbers of weeks and ignore weekends, shifts and breaks by pasting together the weeks into one long simulation run of n weeks. This works since the system state at the end of one week gives the initial conditions for the next week.

The desired output of the program is the CSL-availability YTD, which is the average CSL-availability over the whole time horizon. Recall that the CSL-availability YTD is calculated with the following formula.

$$CSL\ availability\ YTD\% = \frac{\sum Order\ lines\ delivered\ OT\&IF}{\sum Order\ lines\ requested}$$

Note that due to the logic of the formula, we cannot simply calculate the CSL-availability for each week and take the average of those outcomes to get the CSL-availability YTD since the CSL-availability per week is a percentage of the total order lines in that week, but the amount of order lines changes each week. Therefore, let us define $Y_1, Y_2, Y_3, \ldots$ as output of our simulation program, which is the CSL-availability YTD of week $i = 1,2,3, \ldots$, which means that

$$Y_i = \frac{\sum_{k=1}^{i} Order\ lines\ deliverd\ OT\&IF\ in\ week\ k}{\sum_{k=1}^{i} Order\ lines\ requested\ in\ week\ k}.$$

With our simulation program, we are interested in estimating the steady-state mean $v = E(Y)$, which is also defined by $v = \lim_{i\to\infty} E(Y_i)$ (Law, 2007, p. 508). Law states that because of the initial conditions, the model has a warm-up period and if we want to estimate the mean by calculating the average for a specified simulation length, we need to delete the initial warm-up period first. Let m be the length of the simulation and l the length of the warm-up period, then we can estimate the steady-state parameter v by the following formula (Law, 2007, p. 509).

$$\bar{Y}(m,l) = \frac{\sum_{i=l+1}^{m} Y_i}{m-l}$$

We evaluate the warm-up period by using the method of Welch, which is described in Law (Law, 2007, p. 509). The warm-up period is $l = 200$ and we choose $n = 5000$ as the simulation run length. How we evaluated this is described in Appendix C.

To determine the number of replications needed for each experiment we follow the sequential method, which is described by Law (2007, p. 505). We determine the number of replications needed for each experiment and then use the maximum of the minimal required number of runs as the number of runs for calculating the results, see Appendix C.2. The number of replications we use is 4 for the SKU of product group B and 2 for the SKU of product group A.

6.2 Results

In this section, we will present the results of the experiments and analyze the output. First, we run the experiments for the different settings of the factors lead time and safety stock and then for the dual sourcing solution. For all experiments, we use the replication/deletion approach described in Chapter 9.5.2 in (Law, 2007). We calculate 2 replications for the SKU of product group A and 4 replications for the SKU of product group B for each experiment and delete the warm-up period. Then we take the average CSL-availability YTD of the different replications to give our result. We also give the 95% confidence intervals. In Appendix C.3 and Appendix C.4, you can find detailed tables with the results of our analysis. In this section, we report the most important results and conclusions from our analysis.

To analyze the cost of the solution, we first recall the solutions that we want to discuss from Chapter 4. For SKU A, we are interested in a cost analysis for using air freight either as only transport mode or combined with shipping. Also, we are interested in the option of having a production line at the production site in North America. For SKU B, we are interested in a cost analysis for a late customization scenario and for using air freight either as only transport mode or combined with shipping. Moreover, we are interested in the effect of having a shorter production lead time at the current supplier or having a supplier in Europe. Note that due to confidentiality, all prices and numbers we use for this cost analysis are different from the prices in reality, but still representative. To be able to compare the solutions, we count the number of pallets that are ordered, the number of products sold and the number of products delivered late.

6.2.1 Replenishment lead time reduction

First, we look at the results for experiments 1-8 for both SKUs. See a visualization of the improvement in CSL-availability when reducing the lead time in Figure 6.1. In the graph, we can clearly see that the CSL-availability improves when reducing the lead time. This improvement in CSL-availability is stronger for the SKU of product group B, because for the other SKU we have to deal with a lot of order cancellation. Looking at the graph, one might think, that lead time reduction has not a big effect on the CSL-availability of the SKU of product group A. But if we take a look at the percentage of order lines that are cancelled, we see a decrease when we reduce lead time, which results in more sales.

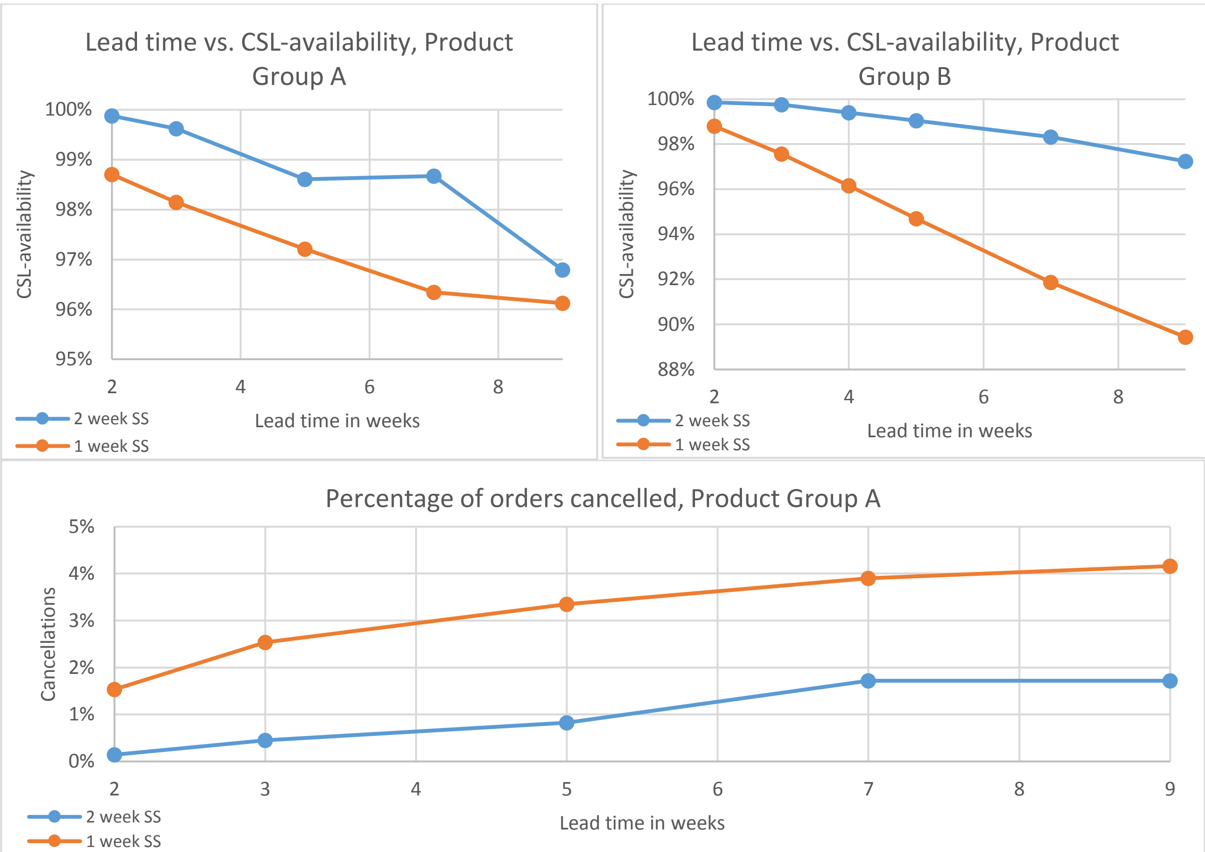

Figure 6.1 Result: Lead time reduction vs. CSL-availability

To analyze the main effect of the change of a factor we follow the procedure within the 2^k factorial design described in Chapter 12.2 in (Law, 2007). We choose experiments 1, 2, 7 and 8 for this evaluation. Let the CSL-availability YTD be denoted by R_i, where i is the number of the experiment. To calculate the main effect of the lead time we use the following formula.

$$e_{LT} = \frac{(R_2 - R_1) + (R_8 - R_7)}{2}$$

When we calculate the effects for all experiments with a two weeks lead time reduction, we see that there is a CSL-availability improvement of 0,8% for the SKU of product group A and 1,7% for the SKU of product group B. For product group B, the improvements in CSL-availability are almost linear. The SKU of product group A does not have a completely linear improvement, which results from the fact that a lot of late orders are cancelled. However, we see that the amount of orders that is cancelled decreases by 0,5% on average for a lead time reduction of two weeks. When we calculate the effects for all experiments with a one week safety stock reduction, we see that there is a CSL-availability improvement of 1,5% for the SKU of product group A and 4,1% for the SKU of product group B. The amount of orders that is cancelled for product group A reduces by 2,2% on average. Therefore it might seem more attractive for Philips to set higher safety stocks for reaching the improvement in CSL-availability. However, we excluded this analysis from our research scope.

Air freight product group A

We analyze the improvement in customer service level and the corresponding costs for transporting all products via air for the SKU of product group A. In Table 6.5, we list prices for the cost price per pallet for the finished goods, the selling price to the customer per product, the transport costs for sea and air transport and the percentage of CSL-availability hits that we consider as lost sales.

Table 6.5 Prices for the SKU of product group A

	Price	
Cost price per pallet	-	1.833,00 €
Sell price per product		2,00 €
Transport pallet (sea)	-	54,00 €
Transport pallet (air)	-	135,00 €
Lost sales		-30%

Before we can start with the cost analysis, we need to know how many pallets we order, how many products we sell and how many products are a CSL-availability hit. In addition to that, we collect data of the amount of products which are cancelled during the simulation run. The cancelled products are directly subtracted from the sales profit. Because of the cancellations, the amount of pallets ordered can also vary. In Table 6.6, we can for example see that more products are sold, but also more pallets are ordered.

To get this information, we run our simulation model again and take the average of 2 replications. Since one simulation run is 5000 weeks long, we get large numbers and therefore divide everything by 100 just to have smaller numbers to work with. Now we know the number of pallets and the demand for 50 weeks, which is almost a year. For all products, including cancelled orders, that we cannot deliver on time, which is a different amount for sea freighting and air freighting, we assign a penalty cost of 30% of the selling price per late delivered product, which we call lost sales. We now describe the procedure of evaluating the costs, which is shown in Table 6.6. We calculate the amount of profit we have through sales by multiplying the selling price with the total amount of sold products. The cost price for buying a pallet of products and the price for transportation is multiplied with the number of pallets ordered in total. If we sum up these costs with the sales and lost sales, we get the absolute Integral Gross Margin (IGM). We see in Table 6.6 that the change in Absolute IGM is +0,3%, which means that transporting all products via air leads to 0,3% more income. Additionally, we have inventory savings because there is less inventory in the pipeline, when the lead time is lower.

To evaluate the inventory savings, we calculate the Cost of Goods Sold (COGS) per week, which is the cost price of all pallets plus the transport price divided by the amount of weeks, which is 50. We then evaluate the amount of weeks that one pallet stays in the supply chain. Remember that the replenishment lead time of nine weeks consisted of one week of planning lead time, one week of production lead time and seven weeks of transport. Additionally there is two weeks of safety stock in the RDC. Therefore a products stays on average ten weeks in the supply chain. When transporting the product via air, the transportation lead time is only one week. When multiplying the COGS per week with the amount of weeks that the product stays in the supply chain, we get the net inventory value. The costs of spending money on inventory can be calculated with Weighted Average Cost of Capital (WACC). The WACC is 6,7% in our fictive example. We subtract 6,7% of the net inventory value from the absolute IGM to evaluate the total profit for our product.

Table 6.6 Cost analysis: Air Freight SKU A

SKU A	Sea Freight	Air Freight	Percentage change
Sales	€ 355K	€ 361K	
Cost price all pallets	€ -128K	€ -130K	
Cost price transport	€ -3,7K	€ -10K	
Lost sales	€ -2,7K	€ -218	+92%
Absolute IGM	€ 220K	€ 221K	**+0,3%**
COGS	€ -132K	€ -140K	
COGS per week	€ -2,6K	€ -2,8K	
Amount of weeks in Supply Chain	10	4	
Net inventory value	€ -26K	€ -11K	
WACC (6,7%)	€ -18K	€ -7,5K	-58%
Total per year	€ 203K	€ 213K	**+5,3%**

Now that we have described the procedure of how we calculate the costs of a solution, we can focus on looking at the results. The detailed cost tables for all following solutions can be found in Appendix C.3 and Appendix C.4. Let us now take a look at the result for transporting the SKU of product group A via air instead of sea in Table 6.7.

Table 6.7 Result: Air Freight product group A

	Sea Freight	Air Freight	Profit change
2 weeks SS			
Total profit	€ 203K	€ 213K	**+5,3%**
CSL-availability	96,8%	99,9%	
Confidence Interval	[96,7%, 96,9%]	[99,9%, 99,9%]	
Percentage of order lines cancelled	1,7%	0,1%	
1 week SS			
Total profit	€ 196K	€ 210K	**+7,1%**
CSL-availability	96,1%	98,7%	
Confidence Interval	[96,1%, 96,1%]	[98,7%, 98,7%]	
Percentage of order lines cancelled	4,2%	1,5%	

We see that changing the mode of transportation results in a profit increase of 5,3% for two weeks of safety stock and an increase in profit of 7,1% for one week of safety stock. The improvement in CSL-availability is 3,1% and 2,6% respectively. Therefore, implementing this solution seems desirable for

Philips. In section 6.2.2, we analyze if transporting part of the orders via sea and the other part via air results in a higher profit increase.

Production site in North America for product group A

Another idea to reduce the transportation lead time is to move the production closer to the country in which the products are sold. For product group A this could be a possibility since there is already a production site close to North America and only a new production line would be needed. The lead time for this solution consists of one week of production lead time and one week of transportation lead time. Before looking at the costs of the solution, we look at the cost benefits that we get from the lead time reduction. We estimate the costs for road transportation being half the costs for sea transportation per pallet. The savings from less lost sales due to CSL-availability improvement and a reduction in inventory costs due to a smaller lead time leave us with € 15K – € 20K of additional profit depending on the amount of safety stock as a budget for the project. Using these cost benefits for implementing the production line, it would take years before the project would start being profitable again, which is not desirable for Philips. Therefore, we can conclude that further investigation about the costs for this project is not necessary, because of the low benefits.

Air freight product group B

Now, we calculate the costs for the option that all products are transported via air for SKU B and compare that to the current costs. This SKU has different characteristics and therefore different prices. In Table 6.8, we list prices for the cost price per pallet for the finished goods, the selling price to the customer per product, the transport costs for sea and air transport and the percentage of CSL-availability hits that we consider as lost sales.

Table 6.8 Prices for SKU B

	Price
Cost price per pallet	- 1.865,00 €
Selling price per product	30,00 €
Transport pallet (sea)	- 54,00 €
Transport pallet (air)	- 360,00 €
Lost sales	-30%

The cost price for buying a pallet of products and the price for transportation is multiplied with the number of pallets ordered in total. The costs for the solution are calculated in almost the same manner as described above. The only difference is that we do not have to deal with cancellations. Note that we do not sell more when air freighting, because everything is backordered and eventually sold. We show the results for transporting everything via air in Table 6.9.

Table 6.9 Result: Air Freight product group B

	Sea Freight	Air Freight	Profit change
2 weeks SS			
Total profit	€ 2.163K	€ 2.032K	**-6,1%**
CSL-availability	97,2%	99,0%	
Confidence Interval	[97,2%, 97,3%]	[99,0%, 99,1%]	
1 week SS			
Total profit	€ 2.072K	€ 1.998K	**-3,5%**
CSL-availability	89,4%	94,7%	
Confidence Interval	[89,4%, 89,5%]	[94,6%, 94,7%]	

We see that the total profit decreases by 6,1% for two weeks of safety stock and by 3,1% for one week of safety stock. The improvement in CSL-availability is 1,8% and 5,3% respectively. However, we can conclude from the cost analysis that transporting all products via air is too costly even when in a low stock situation. The result is different from that of product group A, because for this product the costs of transporting via air are much higher compared to transporting via sea. Furthermore, the sales of SKU A increase due to less order cancellations. In section 6.2.2, we analyze if transporting part of the orders via sea and the other part via air results in a profit increase.

Late customization product group B

Next, we look at the solution of late customization for SKU B. In this scenario, we buy the finished good in a plain box, which we call Key Module (KM) from now on. The KMs are sent to a supplier for putting a sleeve around the products to make them country specific. We assume that the sleeving supplier has two weeks of safety stock and that this is enough to never run out of stock. We use this assumption to determine if this solution setup is profitable. If it seems that it is not then there is no further research needed for the optimal amount of safety stock. Furthermore, we ignore any cost increases due to additional communication between suppliers.

Let the additional handling costs at the sleeving supplier be €100 per pallet and the additional transport price from the sleeving supplier to the RDC €10 per pallet. The price of the finished good at the first supplier decreases by €50 per pallet, because all products are put in plain boxes instead of country specific packaging. The other prices remain as stated in Table 6.8. The additional handling and transportation costs are added to the cost price for the pallets and the transport, see Appendix C.4. For evaluating the costs of inventory, we need to take a closer look at the supply chain. First, there are now KMs in the beginning of the supply chain, which have a slightly reduced COGS. The amount of weeks in the supply chain of the KMs is ten weeks consisting of eight weeks production and transport and two weeks safety stock at the sleeving supplier. The COGS for the finished product is slightly higher than in the current situation due to additional handling and transportation costs. However, the lead time in the supply chain is only four weeks consisting of one week reserved for sleeving at the supplier, one week of transportation to the RDC and two weeks of safety stock at the RDC. Table 6.10 shows the result of the cost analysis.

Table 6.10 Result: Late customization product group B

	Sea Freight	**Late customization**	**Profit change**
2 weeks SS			
Total profit	€ 2.163K	€ 2.053K	**-5,1%**
CSL-availability	97,2%	99,0%	
Confidence Interval	[97,2%, 97,3%]	[99,0%, 99,1%]	
1 week SS			
Total profit	€ 2.072K	€ 2.046K	**-1,2%**
CSL-availability	89,4%	94,7%	
Confidence Interval	[89,4%, 89,5%]	[94,6%, 94,7%]	

We see that this solution results in a profit loss of 5,1% when having two weeks of safety stock and a profit loss of 1,2% when having one week of safety stock. The improvement in CSL-availability is 3,1% and 2,6% respectively. The cost savings of the KM COGS and the lower penalty costs due to lost sales do not weigh out the increased handling costs and the increase in inventory. We assumed a safety stock of two weeks at the sleeving supplier. Even reducing the safety stock to one week without assuming more penalty costs does also not result in a profitable solution.

Supplier in Europe product group B
Since no supplier for this product exists in Europe, we can only create a hypothetical situation. We assume one week of transportation lead time and two weeks of production lead time. If producing the product at the new supplier would cost the same as at the old one, then the profit would increase by 5%-9%. However, production in Europe is most probably more expensive than in Europe and therefore the cost price for the product would be higher. The profitability does not decrease for an increase in cost price per pallet of 8% for two weeks of safety stock and 14% for one week of safety stock. If the increase in cost price per pallet stays under those percentages, then it is advisable to switch to the new supplier, assuming the transportation costs stay the same.

Shorten production lead time supplier product group B
The production lead time at the supplier is currently three weeks, because the supplier wants to lower the risk of not being able to deliver due to long lead time components. Assuming all costs stay the same and the supplier will not run low on stock for components, the profit would increase by 5% if the production lead time is shortened by two weeks and the CSL-availability increases by 1%-2%. We recommend Philips to negotiate the possibility of this lead time reduction. Philips can offer the supplier an increase in the cost price for up to 6% and still stay profitable.

6.2.2 Dual sourcing
In this section, we take a look at the cost analysis for implementing a dual sourcing strategy. For this we count the amount of pallets ordered via the regular and expedited channel. The regular channel will be our current supplier with a replenishment lead time of eight weeks plus one week planning lead time. The expedited channel is the same supplier, but then the products are transported via air with a replenishment lead time of two weeks for SKU A (one week production lead time plus one week transportation lead time) and four weeks for SKU B (three weeks production lead time plus one week transportation lead time) plus one week planning lead time for both SKUs.

First, we look at costs for SKU A for implementing dual sourcing. Also here, we need to subtract the cancelled order lines from the total sales profit. The cost prices are the same as listed in Table 6.5. We vary z_e, the expedited order-up-to-level, for two values of z_r, the regular order-up-to-level. In the following graph, Figure 6.2, we plot the percentage of profit increase against the percentage of pallets that we send via air.

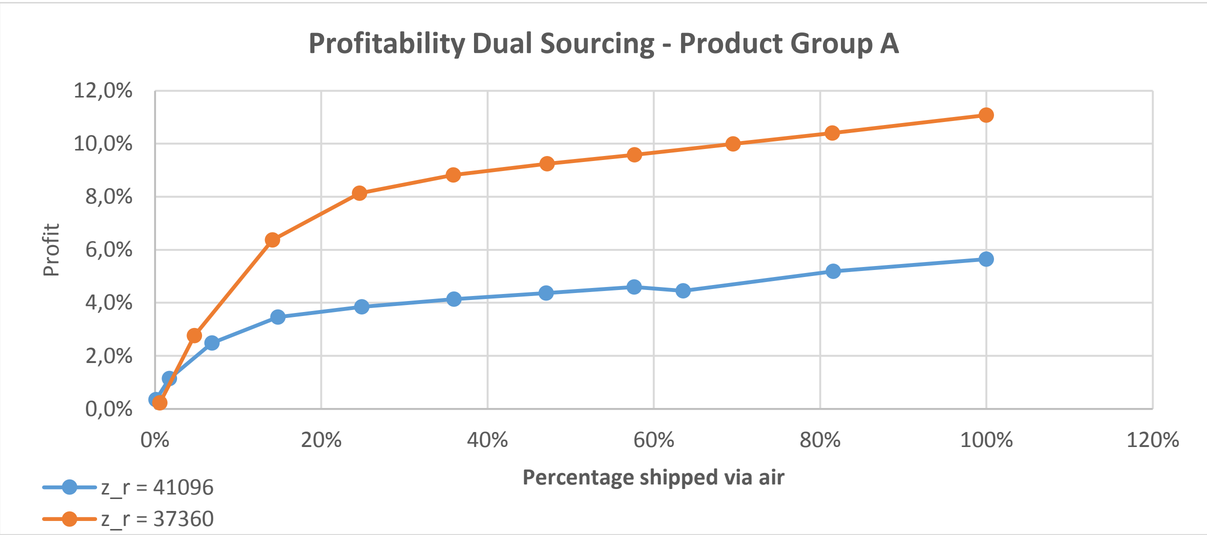

Figure 6.2 Profitability dual sourcing product group A

Remember that we have seen earlier, that sending all products via air is indeed profitable. We are interested if the profit increases even more if we only send part of the pallets via air and the other part via sea. In Figure 6.2, we can see that the profitability increases the more we ship via air. Therefore, we can conclude that for this product group it is best to ship everything via air instead of sea.

Next, we evaluate the costs for SKU B. The cost prices are the same as listed in Table 6.8. The difference in costs for transporting via sea instead of air is much bigger for this product group. We again vary z_e, the expedited order-up-to-level, for two values of z_r, the regular order-up-to-level. In the following graph, Figure 6.2 we plot the percentage of profit increase against the percentage of pallets that we send via air.

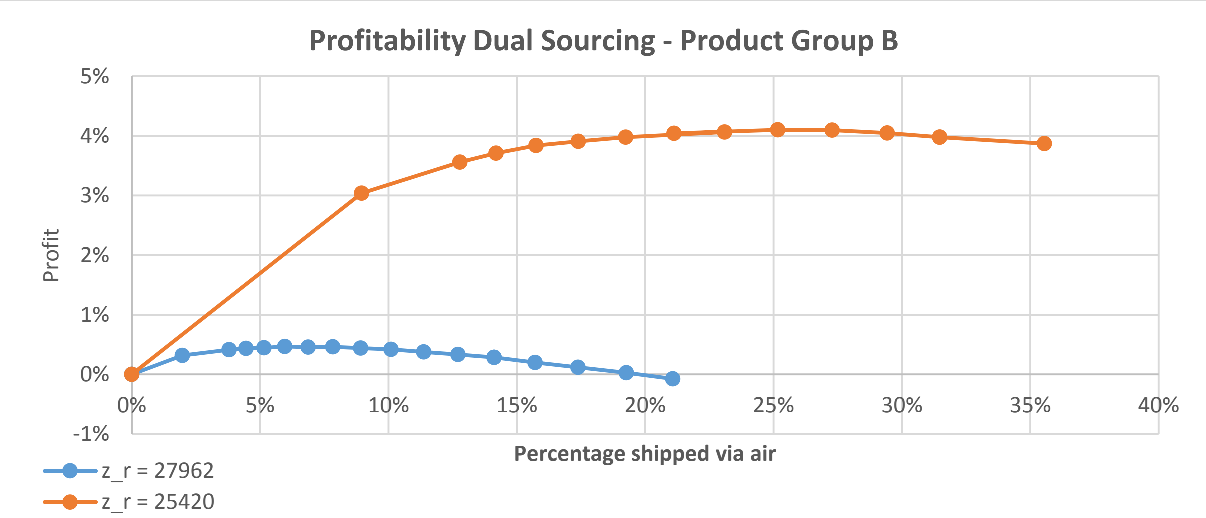

Figure 6.3 Profitability dual sourcing product group B

Remember that we have seen earlier, that sending all products via air is not profitable at all. We are interested if the profit increases if we only send part of the pallets via air and the other part via sea. In Figure 6.3, we see indeed that there are proportions of combining sea and air freight that are profitable. The most profitable solution that we found is sending 6% of all shipments via air, which corresponds with having an expedited order-up-to-level of $z_e = 15760$, when having a regular order-up-to-level of $z_r = 27962$ and sending 25% of all shipments via air, which corresponds with having an expedited order-up-to-level of $z_e = 16523$, when having a regular order-up-to-level of $z_r = 25420$. We advise Philips to implement these settings, which will result in a profit increase of 0,5% and 4,1% respectively and an improvement in CSL-availability of 1,6% and 9,2% respectively.

6.3 Conclusions

In this chapter, we evaluated the impact on the customer service level and the costs for the different solutions that we have formulated in Chapter 4. The cost analysis is done with different price information, but comparable price proportions. We summarize our findings for each solution.

For our experiments, we choose a warm-up period of 200 and a the simulation run length of 5000. The number of replications we use is 2 for the SKU of product group A and 4 for the SKU of product group B. When reducing the lead time by two weeks, the CSL-availability improves on average by 0,8% for the SKU of product group A and 1,7% for the SKU of product group B.
When increasing the safety stock by one week, the CSL-availability improves on average by 1,5% for the SKU of product group A and 4,1% for the SKU of product group B. The CSL-availability of SKU A does not seem to improve as much as the CSL-availability of SKU B. However, SKU A has a lot more cancellations than SKU B, which decrease by 0,5% on average when decreasing the lead time by two weeks and 2,2% on average when increasing the safety stock by one week. Therefore it might seem

more attractive for Philips to set higher safety stocks for reaching the improvement in CSL-availability. However, we excluded this analysis from our research scope.

A lead time reduction for SKU A can be reached through changing the transport mode from sea to air. The improvement in CSL-availability YTD is 3,1% for two weeks of safety stock and 2,6% for one week of safety stock. The profit increase due to less inventory costs, less lost sales and more sales due to less cancellations is 5,3% for two weeks of safety stock and 7,1% for one week of safety stock. We also researched a dual sourcing strategy, where products are partly shipped via sea transport and partly via air transport. However, it the most profitable solution is to ship everything via air. Therefore, we can recommend Philips to implement this solution.
Another idea to achieve a shorter transportation lead time is to build a production line at the production site near the RDC in North America. The benefits of this solution, namely savings in inventory costs and less lost sales, are € 15K- € 20K. Without calculating the precise costs of implementing a new production line, we can estimate that the costs are way more than the benefits. Therefore, we can conclude that this solution would be too costly to implement for Philips.

The lead time for SKU B can be reduced by transporting goods via air instead of sea either completely or partly with a dual sourcing strategy. Changing the transport mode completely from sea to air is however not profitable and will result in a profit loss of 6%. With a dual sourcing strategy, we found settings, which do lead to a profit increase. The most profitable settings for the expedited order-up-to-level z_e are $z_e = 15760$ when the regular order-up-to-level is $z_r = 27962$ and $z_e = 16523$ when the regular order-up-to-level is $z_r = 25420$. This corresponds with sending 6% and 25% of all shipments via air respectively.. We can advise Philips to implement these settings, which will result in a profit increase of 0,5% and 4,1% respectively and an improvement in CSL-availability of 1,6% and 9,2% respectively.
Implementing the idea of late customization results in a profit loss of 5,1% when having two weeks of safety stock and a profit loss of 1,2% when having one week of safety stock. The cost savings of the KM COGS and the lower penalty costs due to lost sales do not weigh out the increased handling costs and the increase in inventory. We assumed a safety stock of two weeks at the sleeving supplier. Reducing this to one week does also not result into a profitable solution.
Switching to a supplier in Europe decreases the lead time. Since there is no supplier available, we assume a lead time consisting of two weeks of production and one week of transportation. A supplier in Europe will probably have a higher cost price per pallet. We conclude that if the increase in cost price per pallet stays under 8% for two weeks of safety stock and 14% for one week of safety stock, then it is advisable to switch to the new supplier, assuming the transportation costs stay the same. This will then result in an improvement in CSL-availability of 2,5% and 8,1% respectively.
Finally, shortening production lead time at the supplier by two weeks and assuming all costs stays the same and the supplier will not run low on stock for components, the profit would increase by 5% if the production lead time is shortened by two weeks and the CSL-availability increases by 1%-2%. We recommend Philips to negotiate the possibility of this lead time reduction. Philips can offer the supplier an increase in the cost price for up to 6% and still stay profitable.

7 Overall conclusions

In this chapter, we summarize our most important findings of this research point by point.

- Due to CSL OTTR reports of the Customer Collaboration Team of Philips, 78% of all CSL OTTR failures are due to unavailability of products at the warehouses.
- Most of the CSL-availability hits, 84% in the first half of 2016 and 74% in 2015, are due to forecast errors.
- Product group A has a quite high CSL-availability YTD of 93%. But small improvements in this product group have a big impact, up to 1,5%, on the overall CSL-availability of the case study business.
- Product group B has structurally low CSL-availability YTD of 87%. This CSL-availability YTD leaves much room for improvement.
- Due to the logic of the calculation, to reach a higher order fill rate, a company should then prioritize the order lines with a small number of products. However, sales and profitability targets give a higher prioritization to customers who order larger amounts.
- There are several ways to cope with demand uncertainty, like improving forecast accuracy, increasing safety stocks or reducing the number of product variants. Reducing the lead time and review period while keeping the safety stock at the same level will increase the customer service level.
- The biggest opportunity in lead time reduction for product group A and B lies in the transportation lead time.
- We have seen that reducing safety stocks by one week reduces the CSL-availability, namely 2,4% for the SKU of product group A and 7,5% for the SKU of product group B.
- The order of serving customers plays an important role on the total amount of CSL-availability hits, especially when in a low stock situation. Serving order lines with a large amount of products first lead to a reduction in CSL-availability of 2,4% - 5,4% for the SKU of product group A and 1,9% - 6,2% for the SKU of product group B compared to a FCFS prioritization.
- When reducing the lead time by two weeks, the CSL-availability improves on average by 0,8% for the SKU of product group A and 1,7% for the SKU of product group B.
- When increasing the safety stock by one week, the CSL-availability improves on average by 1,5% for the SKU of product group A and 4,1% for the SKU of product group B.
- The preferred solution for implementation for product group A is transporting all replenishment orders via air instead of sea. The improvement in CSL-availability YTD is 3,1% for two weeks of safety stock and 2,6% for one week of safety stock. The profit increase due to less inventory costs, less lost sales and more sales due to less cancellations is 5,3% for two weeks of safety stock and 7,1% for one week of safety stock.
- Building an additional production line for product group A near the RDC in North America is not cost beneficial.
- For product group B, sending 6% for two weeks of safety stock and 25% for one week of safety stock of all shipments via air will result in a profit increase of 0,5% and 4,1% respectively and an improvement in CSL-availability of 1,6% and 9,2% respectively.
- If the increase in cost price per pallet, when switching to a supplier in Europe for product group B, stays under 8% for two weeks of safety stock and 14% for one week of safety stock, then it is advisable to switch to the new supplier, assuming the transportation costs stay the same. This will then result in an improvement in CSL-availability of 2,5% and 8,1% respectively.

- Finally, shortening production lead time at the supplier by two weeks and assuming all costs stays the same and the supplier will not run low on stock for components, the profit would increase by 5% if the production lead time is shortened by two weeks and the CSL-availability increases by 1%-2%. We recommend Philips to negotiate the possibility of this lead time reduction. Philips can offer the supplier an increase in the cost price for up to 6% and still stay profitable.

8 Recommendations for further research

Finally, in the last chapter we discuss all our findings during this research. Moreover, we give recommendations to Philips for further research.

We started this research with the goal of finding ways to improve the customer service level of a certain part of Philips, the case study business. We looked into how the CSL OTTR YTD is calculated and quickly we raise the question if it is the right measurement. The CSL OTTR YTD measures how many order lines can be fulfilled on time and in full. To get a nice KPI, all order lines with a small amount of products could be prioritized. However, Philips needs customers who order a large amount of products to stay profitable. If those customers are always prioritized last, then the satisfaction of those customers will decrease and opportunities for big sales are lost. Since the amount of products per order line varies heavily, e.g. for product group B from 3 to 1745 products on one order line, we find it questionable if it is fair to weigh such order lines the same. To us, it sounds more logical to use a volume fill rate, because volume has a more direct relationship to profitability than order lines when the volume per order line is so different.

Some stakeholders in the management functions explained us that this measurement is used to give smaller customers a chance to be treated the same as bigger customers. However, stakeholders in the factory find it hard to work with that KPI because prioritizing in that KPI is not intuitive since sending more volume to a warehouse does not always lead to a better CSL OTTR. When using the volume fill rate as KPI, there is a more direct relationship and easier to prioritize.

Furthermore, a comparison between different Markets can be unfair. Take for example one Market, which has a low CSL, because it misses a lot of small order lines. Then another Market can have a much higher CSL, while being short a lot more products, because it only misses one order line. The volume fill rate measures the percentage of products available on time. Markets get more comparable with this measurement, but it still remains a little unfair, since the total volume that a Market sells varies between Markets.

All in all, we recommend Philips to reconsider using the order line fill rate. As an alternative, we suggest to use the volume fill rate, which makes it easier for production to prioritize orders and gives better insight in the performance of the markets in relation to sales and profitability.

The scope of our project was the CSL-availability, because a lot of CSL OTTR hits are caused by unavailability at the RDC. The root cause for unavailability is mostly forecast related. We then started looking into the CSL-availability levels of different product groups. When trying to improve the CSL-availability we recommend Philips to not only look at products groups with a low CSL-availability, but also focus on improving product groups with a high impact on the total CSL-availability of all product groups. Furthermore, we think it is important to focus on core products. Also, we can learn from product groups with a high CSL-availability and maybe copy solutions to other product groups. Therefore, we recommend Philips to keep an improvement sheet with possible solutions, which have worked earlier for product groups or maybe entire production sites. Such a sheet should contain information on improvement projects like: who was the owner of the project, what kind of improvements were made for which product groups, what were the difficulties of the project and what is the measured effect on the CSL-availability after implementation. Furthermore, when Philips wants to improve the CSL-availability of one single product group it is important to completely understand the root causes of CSL-availability hits.

During our literature research, we found an article which indicates that targets for the customer service level have an optimal point regarding to the extra costs needed for reaching a higher target, see (Jeffery, Butler, & Malone, 2008). This article shows for which volume fill rate the balance between inventory costs and costs of lost sales is optimal. The authors use the following formula to calculate the costs for a given service level.

*Cost(Service level) = Inventory units/period*Inventory Holding Cost/unit*
*- Expected Lost Sales units/period*Profit Margin/unit*

We recommend Philips to research if their target setting is indeed still cost-efficient. If reaching the CSL target results in too much costs, it can harm the profitability of the company. But on the other hand, we find that a high target setting encourages employees to come up with new ideas for reaching a higher CSL without increasing costs. This analysis is about finding the optimal balance between the costs for lost sales and costs for safety stocks. Since we excluded finding the right safety stock settings from the scope of our research, we do not derive the optimal service level regarding to inventory costs. Furthermore, the volume fill rate instead of the order fill rate is used in the article. Since the relation between the order line fill rate and the volume fill rate is not clearly defined, it is hard to evaluate the optimal order fill rate using this method.

While designing the solutions we learned that products are constantly phased out and phased in. Therefore, when developing a solution, one should always consider a whole group of products and not a single SKU. In our research, we performed a simulation study for only one single SKU per product group due to data limitations. However, the results can strongly depend on the SKU selection, due to e.g. different demand patterns of SKUs or different volumes of orders and shipments. Due to the wide variation of SKUs of product group A, the chosen SKU represents only 0,5% of all order lines and is therefore not representative for the whole product group. We recommend Philips to also study other SKUs before implementing a solution for this product group. The SKU of product group B represents 43% of all order lines. Therefore, our conclusions about this product group are much more valuable.
Also, when a product group is dependent on one or more other product groups, because e.g. they are all produced in the same production facility, it should be considered to include all product groups of that facility in the scope. Solutions can be more cost-effective when being implemented for a large amount of SKUs.

One solution that we thought of for product group A, we excluded from our analysis due to time restrictions. However, we recommend Philips to research that idea further. The idea is the following: Product group A currently has almost 290 different SKUs, including SKUs which are currently phased in, which are sold in more than 35 different countries. Due to different languages and country regulations, not all countries are able to share the same SKU. The number of SKUs has risen by almost 30% in the last two years. From a marketing point of view, the products are necessary to stay competitive in the market, but due to this high number of SKUs, the supply chain performance of the case study business decreases for this type of products. With every SKU the factory has to produce more in a single week, the capacity decreases due to changeovers. From literature, we know that the inventory levels at the RDC are also impacted by SKU variety, see Chapter 3.2.4. With more SKUs, more safety stock is needed, which results in extra costs. We learned from literature that product variety also influence the forecast accuracy. In the case study business, for each SKU a forecast is made and this process is more challenging if there are more SKUs, because the demand per SKU is more variable than the demand for the product group. We can assume from our learnings from literature and the experience from stakeholders that the demand variability of an SKU decreases when we reduce the amount of SKUs and therefore the forecast accuracy improves. This has a positive effect on the customer service level. Philips needs to perform further research to find a smart way of achieving SKU reduction.

When building our simulation model, we started to analyze the demand of the SKUs. What comes to our attention is that the variation in order size per order line is really big. Order sizes vary from one to 2000 products per order line. We saw in the data, that the CSL-availability for big order lines is lower than that for big order lines. We assume that the CSL-availability for big order lines is lower, because the Market tends to prioritize smaller customers. However this can lead to opportunity losses. Looking at the CSL-availability per customer, we saw that the customers with a CSL-availability of lower than

80% order 56% of the total volume that is ordered for SKU A and 67% of the total volume that is ordered for SKU B. This, together with the fact that the number of products per order line is heavily skewed and spread out, lets us suggest to research if handling big order lines in a different way than small order lines. An idea could be to serve big order lines directly from the factory, since demand is often known a lot earlier than the demand for small order lines. Due to time limitations, we decide to exclude this idea from our analysis and suggest Philips to research this further.

Finally, we perform a simulation study with a theoretical model with similar behaviors as in practice. We analyze the effect of lead time reduction and dual sourcing and on the costs of implementing solutions. For product group A, the best solution is to ship all products via air instead of sea. The improvement in CSL-availability YTD is 3,1% for two weeks of safety stock and 2,6% for one week of safety stock. The profit increase due to less inventory costs, less penalty costs for lost sales and more sales due to less cancellations is 5,3% for two weeks of safety stock and 7,1% for one week of safety stock. The use of a dual sourcing strategy has the most promising result for product group B. The best settings for the expedited order-up-to-level z_e are $z_e = 15760$ when the regular order-up-to-level is $z_r = 27962$ and $z_e = 16523$ when the regular order-up-to-level is $z_r = 25420$. This corresponds with sending 6% and 25% of all shipments via air respectively. It leads to a CSL-availability YTD increase of 1,6% and 9,2% respectively and a profit increase of 0,5% and 4,1% respectively. However, these conclusions are bound to the particular SKUs of this study. We recommend Philips to research the best settings for other SKUs to find the best settings for the whole product group.

9 References

Axsäter, S. (2006). *Inventory Control* (Vol. Second Edition). New York: Springer.

Beutel, A.-L., & Minner, S. (2012). Safety stock planning under casual demand forecasting. *Int. J. Production Economics, 140*, 637-645.

Brown, A. O., Ettl, M., Lin, G. Y., Petrakian, R., & Yao, D. D. (2002). Inventory allocation at a semiconductor company. In J.-S. Song, & D. Yao, *Supply chain structures - coordination, information and optimization* (pp. 283-309). Massachusetts: Kluwer Academic Publishers.

Brown, A. O., Lee, H. L., & Petrakian, R. (2000, July–August). Xilinx Improves Its Semiconductor Supply Chain Using Product and Process Postponement. *INTERFACES, 30*(4), 65–80.

Ciancimino, E., Cannella, S., Bruccoleri, M., & Framinan, J. M. (2012). On the Bullwhip Avoidance Phase: The Synchronised Supply Chain. *European Journal of Operational Research, 221*, 49-63.

Ganeshan, R., & Harrison, T. P. (2002, 9 28). *Ram Ganeshan.* Retrieved from Raymond A. Mason School of Business: http://mason.wm.edu/faculty/ganeshan_r/documents/intro_supply_chain.pdf

Gupta, A., & Maranas, C. D. (2003). Managing demand uncertainty in supply chain planning. *Computers and Chemical Engineering, 27*, 1219-1227.

Hevner, A. R. (2007). A Three Cycle View of Design Science Research. *Scandinavian Journal of Information Systems, 19*(2), 87-92.

Hopp, W. J., Spearman, M. L., & Woodruff, D. L. (1990, June). Practical Strategies for Lead Time Reduction. *Manufacturing Review, 3*(2), 78-84.

Huanga, Y.-Y., & Li, S.-J. (2008). Suitable application situations of different postponement approaches: Standardization vs. modularization. *Journal of Manufacturing Systems, 27*, 111-122.

Jeffery, M. M., Butler, R. J., & Malone, L. C. (2008). Determining a cost-effective customer service level. *Supply Chain Management: An International Journal, 13*(3), 225 - 232.

Johnson, D. J. (2003). A Framework for Reducing Manufacturing Throughput Time. *Journal of Manufacturing Systems, 22*(4), 283-298.

Jonsson, P., & Mattsson, S.-A. (2009). *Manufacturing, Planning and Control.* Berkshire: McGraw-Hill Higher Education.

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

Law, A. M. (2007). *Simulation Modeling and Analysis* (4 ed.). New York: McGraw-Hill.

Lee, H. l., Padmanabhan, V., & Whang, S. (1997). The Bullwhip Effect in Supply Chains. *Sloan Management Review, 38*, 93-102.

Li, C.-L., Erlebacher, S. J., & Kropp, D. H. (1997). Investment in setup costs, lead time, and demand predictability improvement in the EOQ model. *Prodcution and operations management, 6*(4).

Lu, W., Efststhiou, J., & del Valle Lehne, E. (2006). Customer service level in a lead inventory under mass customization. *International Series in Operations Research and Management Science, 87*, 233-250.

Moyaux, T., Chaib-draa, B., & D'Amours, S. (2007). Information Sharing as a Coordination Mechanism for Reducing the Bullwhip Effect in a Supply Chain. *IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS—PART C: APPLICATIONS AND REVIEWS, 37*(3), 396-409.

Ouyang, L.-Y., & Wu, K.-S. (1997). Mixture inventory model involving variable lead time with a service level constraint. *Computers Ops Res., 24*(9), 875-882.

Sabria, E., & Beamon, B. (2000). A multi-objective approach to simultaneous strategic and operational planning in supply chain design. *Omega, 28*(5), 581±598. doi:10.1016/S0305-0483(99)00080-8

Schmidt, M., Hartmann, W., & Nyhuis, P. (2012). Simulation based comparison of safety-stock calculation methods. *CIRP Annals - Manufacturing Technology, 61*, 403-406.

Thonemann, U. W., & Bradley, J. R. (2002). The effect of product variety on supply-chain performance. *European Journal of Operational Research, 143*, 548-569.

Towill, D. R. (1996). Time compression and supply chain management - a guided tour. *Supply Chain Management : An International Journal, 1*(1), 15-27.
Van Aken, J. E. (2004, March). Management Research Based on the Paradigm of the Design Sciences: The Quest for Field-Tested and Grounded Technological Rules. *Journal of Management Studies, 41*(2), 219-246.
Van Aken, J. E. (2005). Management Research as a Design Science: Articulating the Research Products of Mode 2 Knowledge Production in Management. *British Journal of Management, 16*, 19-36. doi:10.1111/j.1467-8551.2005.00437.x
Van Aken, J. E. (2007, March). Design Science and Organization Development. *The journal of applied behavioral science, 43*(1), 67-88. doi:10.1177/0021886306297761
Van der Heijden, M., & Diks, E. (1999). Verdeel en heers: Voorraadallocatie in distributienetwerken. In *Praktijkboek magazijnen en distributiecentra.*
Veeraraghavan, S., & Scheller-Wolf, A. (2008, July-August). Now or Later: A Simple Policy for Effective Dual Sourcing in Capacitated Systems. *Operations Research, 56*(4), 850-864. doi:10.1287/opre.1080.0552
Veeraraghaven, S., & Scheller-Wolf, A. (2008, July-August). Now or Later: A Simple Policy for Effective Dual Sourcing in Capacitated Systems. *Operations Research, 56*(4), 850-864. doi:10.1287/opre.1080.0552
Wan, X., Evers, P. T., & Dresner, M. E. (2012). Too much of a good thing: The impact of product variety on operations and sales performance. *Journal of Operations Management, 30*, 316-324.
Whitney, D. E. (1993). Nippondenso Co. Ltd: A Case Study of Strategic Product Design. *Research in Engineering Design, 5*, 1-20.
Zhao, X., Xie, J., & Zhang, W. (2002). The impact of information sharing and ordering co-ordination on supply chain performance. *Supply Chain Management: An International Journal, 7*(1), 24-40.

Appendix A Background information supply chain

Appendix A.1 Supply chain mapping

In Figure A.9.1, we see a detailed description of the possible product flows in the supply chain. Part (a) of the figure shows how the products and orders flow through the chain for around 80% of the SKUs for the case study business. The products are produced at one of the in-house factories or at the suppliers. The finished goods are delivered to the warehouses of the COs, which are called Regional Distribution Centers (RDCs). From the RDCs, the products are delivered to distributors and the distributors send the products further to the retailers or directly to the consumer.
The order flow looks a bit different. A retailer orders products at the distributor. The distributor order the products he needs at his contact person from the Commercial Organization (CO) of his country. Almost each CO is linked to at least one RDC. There are COs which share one warehouse location, but there are also COs which have two warehouse locations. The CO has administrative tasks, like collecting customer orders and making demand forecasts, whereas the RDC is the location where the products are physically stored. The demand planner at the CO makes forecasts for all SKUs which are sold within that CO and stores them in a system. The supply planner can see all forecasts of all COs and makes a production plan and fills a confirmed production plan for each CO into the system after consultation with the supplier or the in-house factory.

The supply chain structures of part (b) and (c) of Figure A.9.1 are quite similar to part (a). In (b), the finished goods will be sent directly from the RDC to a retailer and the retailer order directly at the CO, instead of ordering at and receiving the products from a distributor. In (c), finished goods will be send directly to the distributor and not through the RDC. This is done exceptionally for the few COs which have no RDC.

The transportation of the products is done by ship, train, truck or air freight depending on the urgency and the location. Each CO is linked to one RDC. In Europe, there are three RDCs, the 3DC's, which store products for different European countries. In that way, if one country has a shortage, the products can be re-allocated to another country taking into consideration packaging and language requirements. Other countries around the world have with some exceptions one RDC per country. Sometimes an RDC can also have a packaging function, which means that finished goods are combined and are sold as a set of products. These sets can be unique per country. Since the RDC's, factories and customers are spread all over the world, the replenishment lead times and the customer lead times vary for different products and different markets.

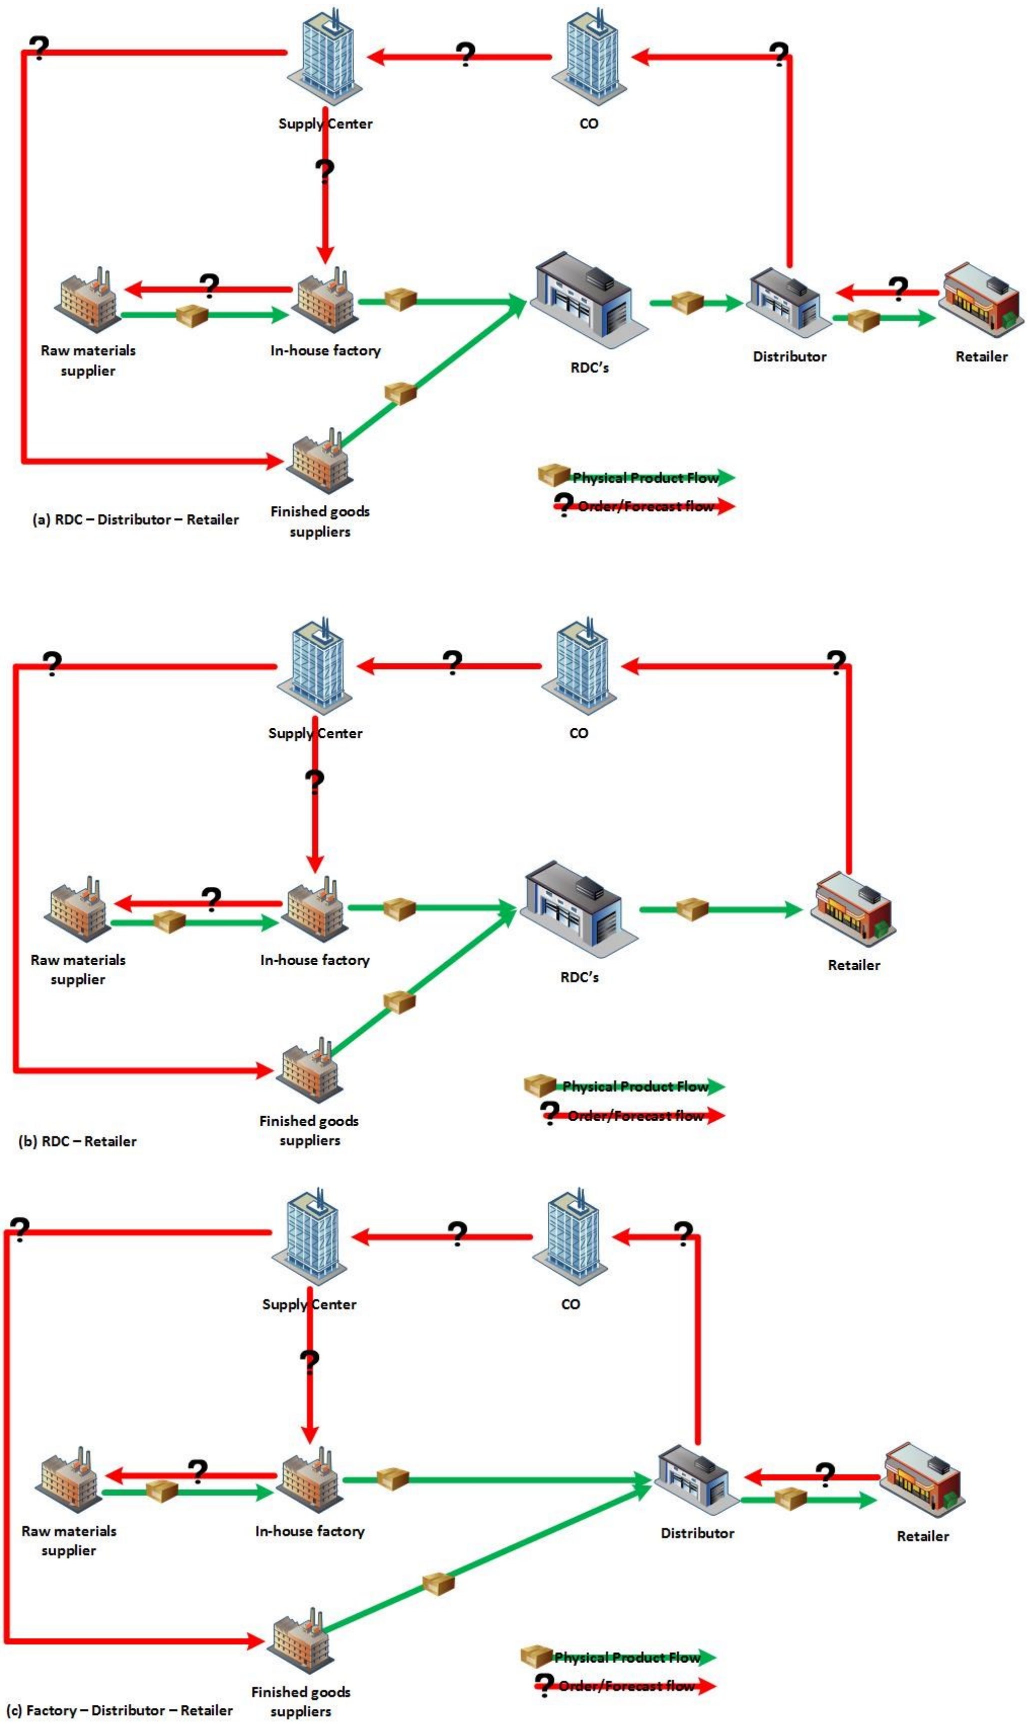

Figure A.9.1 Three different supply chain mappings

Appendix A.2 Detailed root cause analysis

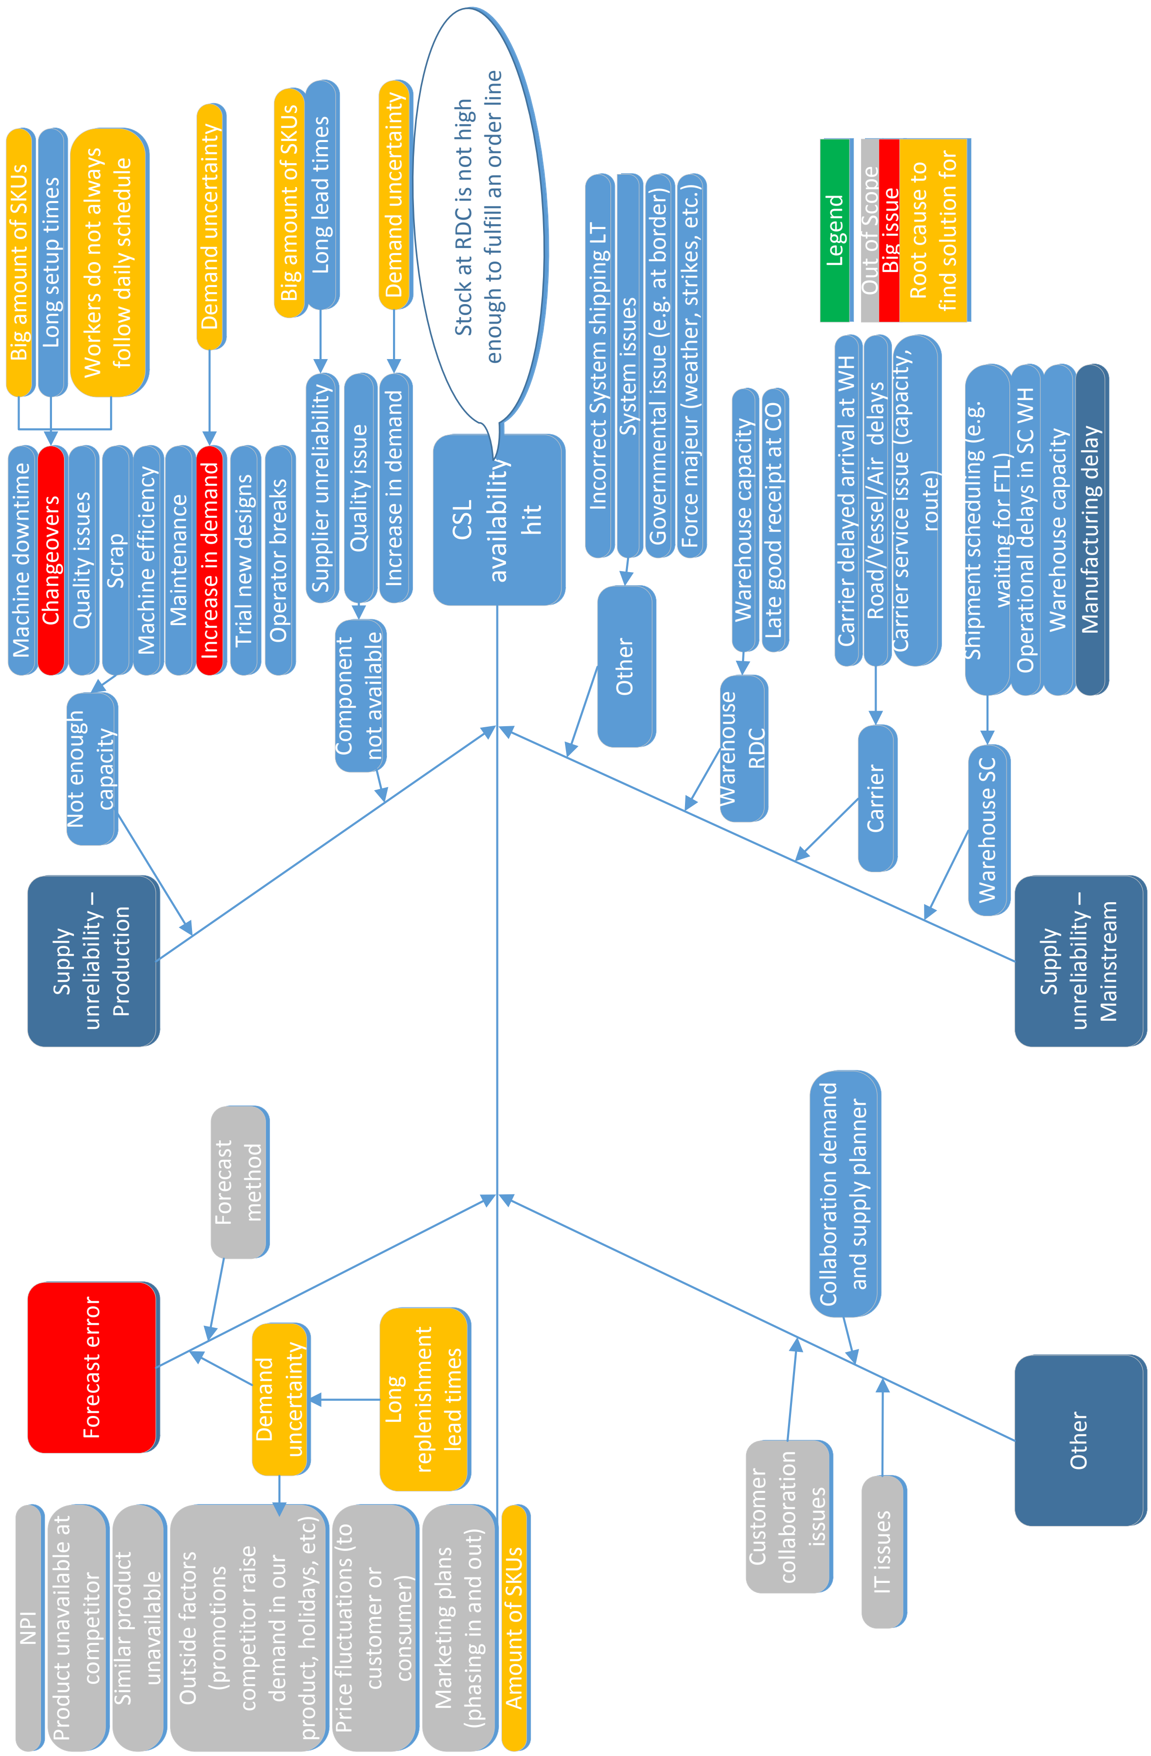

Figure A.9.2 Root cause analysis diagram

In Figure A.9.2 a detailed root cause analysis for a CSL-availability hit is given. We will discuss here in more detail what the root causes for a CSL-availability hit are. A CSL-availability hit means, that for some reason, there is not enough stock available at the warehouse to fulfill an order line of a customer (completely). The reasons for not having enough stock are assigned to four different main categories, namely supply unreliability due to production, supply unreliability due to mainstream, forecast errors or other.

Appendix A.2.1 Supply unreliability due to production

A possible reason for having a shortage of products in the warehouse is that the supply was not sufficient due to some problems on the production site. Generally we categorize two types of problems here, capacity issues and component stock issues.

Capacity

When not having enough capacity at the production site, not all ordered products can be delivered to the warehouses, which leads to reduced safety stocks and can cause stock-outs. Capacity can be lower than usual for many reasons, e.g. machine downtime, operator breaks and maintenance.

Also quality issues can lead to can cause lower capacity since products need to be reworked or even reproduced.

The number of changeovers and the time needed for a changeover also influences the available capacity. If there is a product, which has a lot of different product types, then every time the factory wants to produce another type, the machine needs to be setup differently, which costs time. The factory therefore makes a production schedule to have the right parts for assembly done at the same time. But if workers do not follow the schedule, then not all parts for assembly are ready at the same time and the machine has to wait for the right parts, which reduces the throughput of the factory. Furthermore, the amount of scrap after a changeover also influences the setup time and therefore reduces capacity if a lot of changeovers are needed. But also the design of the machine influences the time needed for changeovers. Some machines can be more efficient than others, but the capacity of the production site is limited by the bottleneck machine.

Another factor which reduces capacity is trial runs for new designs, which reserves some of the total capacity.

Finally, increases in demand due to uncertainty can raise the need for more total capacity. It is not that easy to quickly raise capacity since most probably additional machines are needed which take long to design and implement. Therefore the limited capacity cannot always compete with an increase in demand. The company makes long-term forecasts to prevent this problem from happening and plan to increase production capacity in advance. Therefore, if this problem occurs due to not forecasted increase in demand or of the increase is forecasted later than it actually happens, the root cause for CSL-availability hits is partly production and partly forecast.

Component availability

When components for the products are not available or are not enough to meet all demand, then the production facility cannot produce all products, which are needed at the warehouses, which can lead to stock-outs.

Component stock-outs can occur due to supplier unreliability especially for long lead time components and components with uncertain demand. A big amount of SKUs can lead to a need for a lot of different components. Due to the bullwhip-effect, which we describe in Chapter 3, demand for those components is even more uncertain than for the single SKUs. A sudden increase in demand for a SKU lead to a higher demand in components, which can lead to stock-outs of the components.

Supply unreliability can also mean that something goes wrong when transporting the components, especially with long lead times, like traffic jams.

Finally, quality issues also lead to shortages in components. It can be that components need to be send back to the supplier to reproduce due to e.g. production mistakes or when components are damaged during transport.

Appendix A.2.2 Supply unreliability due to mainstream

Mainstream is a term used by Philips, which stands for transportation. We discuss here reasons for shortages in stock at the RDC due to issues in transportation. These are situations which happen after the products have left the factory.

After finishing all production steps, the finished product enters the warehouse at the supplier, where it is made ready for shipping. However delays can already start here. When finished goods are not directly sent to the warehouse by the factory, which is called manufacturing delay, those products can miss the truck they should go on. Another root cause is that products can be waiting at the entrance of the warehouse to be checked in into the system due to operational delays or not enough warehouse capacity. Only when the product is checked in, shipment plans can be made for that product. When enough products lie in the warehouse to almost have a full truck load (FTL), then a truck is ordered to the warehouse to transport the products to the RDC. Sometimes, products need to wait quite long until a FTL for the RDC where they need to go is reached and therefore they can be too late at the RDC.

When a truck for the product is ordered on time, there can also occur problems with the carrier. If the carrier does not have enough capacity or the route is planned inefficiently, then the product can arrive late at the RDC, which can cause stock-outs. These issues are called carrier service issues. Furthermore the carrier can also arrive delayed at the factory warehouse, but also at the RDC due to traffic.

Sometimes it occurs that products are physically at the RDC, but not available for selling, which also leads to stock-outs in the system. This happens for example if the RDC has reached its capacity and cannot store more products. Some warehouses can rent additional storing space, but that is possibly not always immediately available or products need to be relocated to another storage area. But also simple late good receipts due to operational delays can cause late check-in in the RDC and cause temporary stock-out.

Finally, there are some root causes, which can be categorized as other transporting issues. For example, if there are system issues e.g. when trying to check-in products at the goods receipt or when an incorrect shipping lead time is filled in the planning system. There can also be governmental issues e.g. at the border or force majeure like weather or strikes.

Appendix A.2.3 Forecast errors

The most common root cause for stock-outs is forecast errors. Forecast errors mainly come from either a not advanced enough forecasting method or demand uncertainty. Uncertainty in demand can have many reasons. We list the most common reasons for demand uncertainty for Philips.
First of all, long replenishment lead time increase the chance of forecast errors, because the uncertainty in demand is higher how further in the future the demand has to be forecasted.
Furthermore, the amount of SKUs of a certain product type influences the uncertainty in demand. A lot of different SKUs are much harder to forecast than just one.
In addition to that, new introduced products are hard to forecast since there is no information except for demand planner experience on demand for that product at all.
Sometimes there can be a sudden higher interest in a certain product if a similar product is unavailable or if a product is unavailable at the competitor. Additionally, when a product is unavailable and is backordered by the customer, it can change future demand behavior. It probably takes longer than planned before the customer places a next order, because he received his order too late, or even worse: the customer does not order at all. Moreover, outside factors like promotions of a product of

a competitor, holidays, fashion trends, media or events in the world can influence interest in certain kinds of products.
Price fluctuations of products at Philips and at competitors also influence the order behavior of customers and therefore cause uncertainty in demand.
Finally, marketing plans for phasing in and phasing out products influence the order behavior of customers. If a product is not available anymore, because it is phased out, customers are likely to order similar products. Likewise, if a new product just arrived on the market, customers could order the new product instead of the one they used to order. These shifts in order behavior cause uncertainty in demand.

Appendix A.2.4 Other

Finally, if none of the above is the root cause for a CSL-availability hit, a stock-out at the RDC can have other reasons like IT issues or unknown reasons. Another root cause for stock-outs is miscommunication between the demand planner and the customer, which is called customer collaboration issues and e.g. happens when customers do not share the right information or no information at all on demand, or between the demand planner and the supply planner about e.g. the urgency of receiving products.

Appendix A.3 Stakeholders

The following table presents all stakeholders that were involved in the project. Stakeholders gave us descriptions of the supply chain, insights in root causes and ideas for solutions. We validated our assumptions for the simulation model with different stakeholders. The following table shows the job title and the workplace of each stakeholder and gives a short description of their responsibilities.

Job title	Short Description	Workplace
Director Supply Chain Management	Managing supply chain processes end-to-end	Office Amsterdam
S&OP Manager	Improving business processes including S&OP	Production site product group A
Supply Planner A	Responsible for the supply planning at the in-house factory of product group A	Production site product group A
Supply Planner B	Responsible for the supply planning and collaboration with the supplier of product group B	Office Asia
Supply Chain Engineer	Project-based improvement processes in the supply chain	Office Amsterdam
Supply Chain Manager	Direct responsibility for entire supply chain function, including planning, operational buying, shipping and transport, warehousing and execution.	Production site product group A
Supply Chain Engineer at Production Site	Project-based improvement processes in the supply chain at the factory of product group A	Production site product group A
Supply Chain Manager Customer Collaboration	Improving Global Customer Service Level Metric in Collaboration with Demand Planners	Office Amsterdam
Director Supply Planning	Design, streamline and operate supply chain networks as well as manage operations	Office Eindhoven
Central Forecasting Manager	Managing the Statistical Forecast process: Preparing inputs to statistical forecast, creating a baseline. Review, correct and improve statistical forecast	Office Eindhoven
Regional Distribution Centre Manager	Managing all in- and outbound processes of the distribution center	Distribution Center
Senior Project Manager in Supply Chain	Designing, Implementing & Improving Global Customer Service Level Metric	Office France
Global Inventory Manager	Managing and reviewing inventories including safety stocks	Office Amsterdam
Shipping Coordinator	Coordinating the shipping plan for outbound shipments from the factory	Production site product group A
Inventory planner procurement	Managing inventory for production site product group A	Production site product group A
Demand planner	Planning the demand for a specific region using customer information	Office Germany

Appendix A.4 Detailed supply chain description product group A

The products of product group A are produced in an in-house factory in Europe and are sold globally. The order process for the products is visualized in Figure 2.2. In week 1, the markets fill in the forecasts on Monday into a planning system. The demand planners make a high level production plan manually on Wednesday taking into consideration the production capacity and the prioritization wishes of the markets. The factory planner will confirm or adjust the plan on Thursday and the confirmed plan will be sent out to the markets on Friday. So the total planning lead time is one week.

The production lead time is generally expected to be one week, however prioritization is possible to shorten the lead time. Products which are produced in the beginning of the week can be shipped earlier and therefore arrive earlier at the RDC. The rest of the products is produced during the rest of the week and is finished at the latest on Saturday at 6 am. Sometimes the factory cannot produce all orders during one week. Then the orders are backordered and immediately produced in the beginning of the next week. Some production activities can also be done in the weekend. The factory strives to finish at least 98% of the confirmed orders in one week and to confirm at least 95% of the orders each week. If there is not enough capacity to confirm all orders, then the supply planner decides which market gets less stock, depending on the stock levels at the markets. The supply planner then gives less priority to orders, which only intend to fill up the safety stock.

After finishing production, the products arrive at the warehouse adjacent to the factory. Also other products produced in the factory arrive in that warehouse and are stored in racks. Preferably full truck loads are shipped and therefore products wait until there are enough products for a full truck load. A truck is sent out separately for each RDC, which means that a truck has only one destination and is not visiting multiple RDCs. The shipping planner reserves 24 hours for picking and packing and therefore tries to book a truck as soon as almost a full truck load is lying on the racks. The rest of the products which arrive in the meantime can directly be stored at the packing area. The transportation time varies between one to seven weeks depending on the distance of the factory to the RDC. See
Table A.9.1 for a demonstration of the best case and worst case scenario for transportation to an RDC with one week transportation lead time. Notice that the best case lead time is two weeks and the worst case lead time is three weeks. The case study business assumes the worst case lead time as the total replenishment lead time.

Table A.9.1 Best and worst case lead time scenario

	Week 1					Week 2						Week 3				
	M	T	W	T	F	M	T	W	T	F	S	M	T	W	T	F
Planning	Order placement	Prioritization check	Capacity planning	Factory planning	Confirmation order											
Produc-tion						Best Case					Worst Case					
Arrival WH						Best Case						Worst Case				
Start Shipping							Best Case						Worst Case			
Arrival RDC										Best Case						Worst Case

Appendix B Additional information on the simulation model

Appendix B.1 Using the program for other SKUs and Markets

The model described in the previous sections is developed at the example of two specific SKU and one specific market respectively. Let is discuss in this section, what changes would need to be made to the model if we choose another SKU to analyze.

If we want to analyze the effect of lead time reduction on CSL-availability on a different SKU and/or in another market, we need to adjust the parameters for the lead time, safety stock. Also the demand structure can be different, which needs to be analyzed on the forehand and changed in the program. Furthermore, the pallet size and the production capacity varies for every SKU and also depends on the Market.

The assumptions we made for the order handling and the warehouse of the market can also vary for different markets and need to be validated again. The assumptions about the replenishment lead time and the factory depend on the SKU and can also depend on the market and need to be validated again when changing the SKU or the market.

Since the assumptions we make are very dependent on the type of SKU, market and supply chain, it is hard to generalize the model in such a way that is valid to be used for other SKUs and markets since the supply chain network of Philips is quite complex. Due to the time restriction for this project, we will therefore only analyze the lead time reduction effect and the effect of dual sourcing of the two SKUs.

Appendix B.2 Input distributions

In order to specify random input data for the simulation, we first try to find a theoretical distribution which fits the given historical data. We present the procedure for the specific SKU from product group B delivered to one CO and for the SKU from product group A for one CO. For these estimations we use demand data from 2015 week 32 until 2016 week 32, which includes information on the date of delivery, the customer, the amount of products and the amount of products delivered on time for each order line. We asked stakeholders about information on weeks with promotion. This we could only get for the SKU of product group B. We filter out these weeks from the analysis. We want to model the total demand in our simulation, but we also need to know the number of order lines and the number of order lines, which are not fulfilled on time or in full in order to calculate the CSL-availability. Therefore, our approach is to find a distribution for the number of order lines per week and a distribution for the amount of products per order line. We assume that the number of order lines per week and the amount of products per order line are independent.

Appendix B.3 Input distribution for the SKU of product group A

In this section, we try to fit a distribution for the demand of SKU A. We follow the same approach as for SKU B. Table B.3.1. shows the demand data for SKU A. Also for this SKU, we see that the standard deviation of the total demand is half the size of the average demand. This is probably due to the high variation in the amount of products per order line, which is almost two times bigger than the average. Moreover, the amount of order lines is much smaller than for SKU B, but has a higher demand in total. The standard deviation of the amount of order lines per week is half the size of the average, which also leads to higher variations in the total demand.

Table B.3.1 Demand statistics SKU A

SKU A	Data
Average total demand per week	3736
Std. dev. total demand per week	1997
Average # order lines	20
Std. dev. # order lines	10
Average # products per order line	190
Std. dev. # products per order line	344

Appendix B.3.1 Number of order lines per week

For estimating the demand, we need to estimate the number of order lines per week and the amount of products ordered on one order line. Let us first look at the distribution of the number of order lines. We have no information about promotions. So we try to find a distribution for the whole data set of 53 data points. We start by measuring the location, variability and shape of the data. The average number of order lines per week is 20 with a standard deviation of 10, see Table B.9.2. The data is a little bit skewed to the left, but with a value of skewness of 0,8 almost symmetric. Likewise, the kurtosis is almost zero and lets us expect a normal shape.

Table B.9.2.1.1 Summary statistics #OL SKU A

Number of order lines	
Mean	20
Standard Deviation	10
Kurtosis	0,2
Skewness	0,8
Minimum	5
Maximum	47
Count	53

From the descriptive statistics of the data, we first assume a normal distribution despite the small skewness to the left. We try fitting a normal distribution with mean 20 and standard deviation 10. To test the goodness-of-fit visually, we look at the Q-Q plot of the fitted model distribution versus the real demand data. It can be seen, that the graph is almost a straight line except for low values and high values.

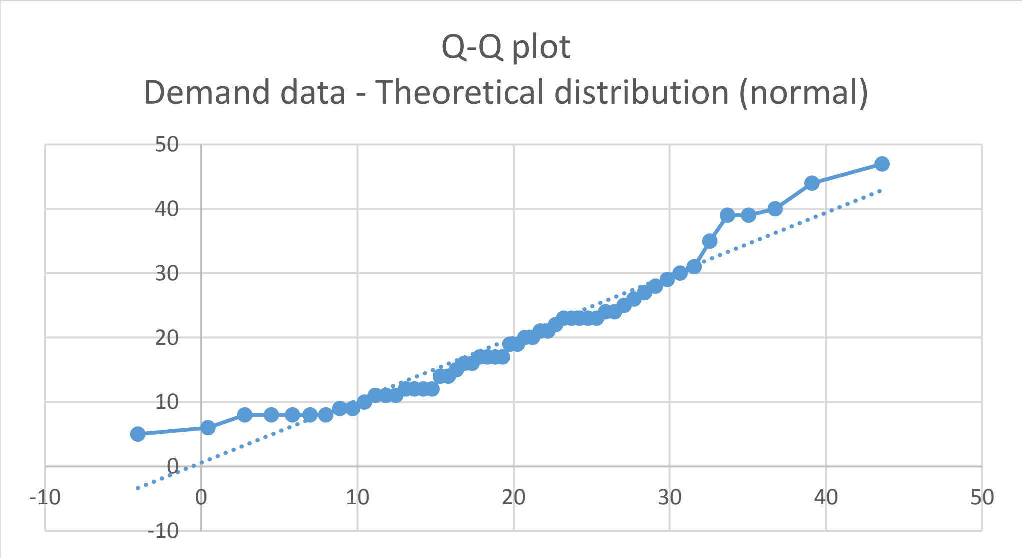

We perform a chi-square test to test the goodness-of-fit of our hypothesized distribution. We follow the same steps as described in Appendix B.4.2. We again choose the amount of intervals to be 20. From the chi-square formula we get the value 28,9 which is a smaller value than the value from the table with a degree of freedom of 19 and a confidence interval of 95%, which is 30,1. Thereby, we conclude that we do not reject the null hypothesis, which states that the normal distribution is a good fit. Therefore, we use the normal distribution with mean 20 and standard deviation 10 for estimating the amount of order lines per week. In the model, we add the condition that the amount of order lines can never be below zero.

Appendix B.3.2 Amount of products per order line

Next we try to estimate the amount of products per order line. Let us look again at the summary statistics. The average amount of products per order line is 13 with a standard deviation of 65. The data is skewed to the left and spread out widely. This, we can read from the values of the skewness and kurtosis in the following table.

Table B.3.2.1. Summary statistics: Amount of products per order line for product group A

Amount of products per order line	
Mean	190
Standard Deviation	344
Kurtosis	45
Skewness	5
Minimum	1
Maximum	4920
Count	1048

For this data, we conclude that it is hard to find a theoretical distribution due to a high skewness and high kurtosis. We therefore will try to find an empirical distribution for the data. We will group the data in intervals, starting with small intervals and ending with bigger intervals and giving each interval a certain probability. See the empirical distribution in the following table.

Value	# observations for value	Probability	Cumulative probability
12	208	0,1985	0,1985
24	87	0,0830	0,2815
36	61	0,0582	0,3397
48	69	0,0658	0,4055
60	48	0,0458	0,4513
72	31	0,0296	0,4809
84	27	0,0258	0,5067
96	44	0,0420	0,5487
108	36	0,0344	0,5830
120	57	0,0544	0,6374
132	41	0,0391	0,6765
144	28	0,0267	0,7032
156	12	0,0115	0,7147
168	23	0,0219	0,7366
180	20	0,0191	0,7557
192	10	0,0095	0,7653
204	8	0,0076	0,7729
216	10	0,0095	0,7824
228	7	0,0067	0,7891
240	15	0,0143	0,8034
252	6	0,0057	0,8092
264	5	0,0048	0,8139
276	1	0,0010	0,8149
288	8	0,0076	0,8225
300	19	0,0181	0,8406
312	2	0,0019	0,8426
324	5	0,0048	0,8473
336	4	0,0038	0,8511
348	5	0,0048	0,8559

360	10	0,0095	0,8655
372	3	0,0029	0,8683
384	6	0,0057	0,8740
396	12	0,0115	0,8855
408	4	0,0038	0,8893
420	4	0,0038	0,8931
432	7	0,0067	0,8998
444	2	0,0019	0,9017
456	2	0,0019	0,9036
468	2	0,0019	0,9055
480	7	0,0067	0,9122
492	1	0,0010	0,9132
504	7	0,0067	0,9198
516	1	0,0010	0,9208
540	2	0,0019	0,9227
552	1	0,0010	0,9237
564	4	0,0038	0,9275
588	2	0,0019	0,9294
600	10	0,0095	0,9389
612	2	0,0019	0,9408
648	2	0,0019	0,9427
672	1	0,0010	0,9437
696	1	0,0010	0,9447
708	1	0,0010	0,9456
720	5	0,0048	0,9504
744	1	0,0010	0,9513
792	2	0,0019	0,9532
804	2	0,0019	0,9552
816	1	0,0010	0,9561
840	2	0,0019	0,9580
852	1	0,0010	0,9590
900	2	0,0019	0,9609
984	2	0,0019	0,9628
1080	2	0,0019	0,9647
1104	1	0,0010	0,9656
1128	1	0,0010	0,9666
1200	13	0,0124	0,9790
1248	1	0,0010	0,9800
1272	1	0,0010	0,9809
1356	1	0,0010	0,9819
1440	1	0,0010	0,9828
1452	1	0,0010	0,9838
1488	1	0,0010	0,9847
1632	1	0,0010	0,9857
1656	1	0,0010	0,9866

1800	10	0,0095	0,9962
1980	1	0,0010	0,9971
2520	1	0,0010	0,9981
2880	1	0,0010	0,9990
4920	1	0,0010	1,0000

Appendix B.3.3 Result of modeling

We now model the distributions in the simulation model described in Chapter 5. We make a run of 5000 weeks and compare the average demand and standard deviation of demand to the real data. We see the results in Table B.9.3. The number of order lines per week and the amount of products per order line are distributed as expected and very close to the summary statistics of the real data. However the standard deviation of the total demand deviates by 17% from the real standard deviation of the demand. From this, we can conclude that our assumption, that the number of order lines per week and the total amount of products per order line are independent, is not true. We expect some correlation in the data. In the next section, we test how the data is correlated.

Table B.9.3 Input distribution validation SKU A

SKU A	Data	Model output	Percentage of deviation
Average total demand per week	3736	3839	3%
Std. dev. total demand per week	1997	2333	17%
Average # order lines	20	20	0%
Std. dev. # order lines	10	10	0%
Average # products per order line	190	189	-1%
Std. dev. # products per order line	344	344	0%

Appendix B.3.4 Correlation Analysis

The results of modeling the demand suggest that there is correlation in the demand, which is not included in the model. Probably our first assumption that the number of order lines and the number of products per order line are independent is not true. It could be that if the number of order lines is high in a week, that a lot of small customers place order lines, which would result in a low average requested quantity per order line. We therefore test the correlation between the average requested quantity and the number of order lines, see Table B.9.4. We see that there is a slightly negative correlation between those two factors. This means that if the number of order lines are high, the number of products ordered per order line is probably low. This could be one reason for the standard variability of the model being higher than in reality.

Furthermore, we see that the number of order lines is slightly correlated with the total demand, but the average requested quantity per order line is somewhat more positively correlated. The amount of order lines in the current week is slightly positively correlated with the amount of order lines in the previous week. Also the number of order lines from 2, 3 or 4 weeks before do have a positive correlation with the current.

Furthermore, the correlation between the total demand of the current and previous week is almost not very high. The same holds for the average requested quantity.

Table B.9.4 Correlation tests SKU A

Correlation Tests	Correlation
Avg. Requested Quantity & #OL	-0,39
#OL & Total demand	0,34
Avg. Requested Quantity &Total demand	0,66
#OL & #OL previous week	0,26
Total demand & previous week	0,14

Avg. Requested Quantity & previous week	0,09
#OL & #OL 2 weeks before	0,47
#OL & #OL 3 weeks before	0,25
#OL & #OL 4 weeks before	0,41

In addition to that, we split the data in months with 4 weeks and months with 5 weeks. For each month, we test if there is a significant difference between the weeks for different factors through ANOVA analysis. We look at differences between the weeks for the total demand and the number of order lines. The result of the ANOVA test lets us conclude that there is no significant difference between the weeks. That means that there is no week in which the demand is significantly higher or lower. Therefore we cannot conclude that customers wait with ordering until the last week of the month. We get the same result for the average requested quantity on an order line. This means that big and small customers place orders throughout the whole months.

Due to the slightly negative correlation between the average requested quantity and the number of order lines, we try to model the empirical distribution differently. This is described in the next section.

Appendix B.3.1 New empirical distribution for the amount of products per order line for product group A

We split the demand in two sets, namely in weeks with the amount of order lines bigger than the average of 20 and in weeks with the amount of order lines smaller or equal to the average. The following table shows the characteristics of our first set.

Amount of products per order line	
Mean	158
Standard Deviation	321
Kurtosis	86
Skewness	8
Minimum	2
Maximum	4920
Count	667

For this data, we conclude that it still is hard to find a theoretical distribution due to a high skewness and high kurtosis, although the mean, standard deviation and skewness are lower those of the whole dataset. We therefore will try to find an empirical distribution for the data. We will group the data in intervals, starting with small intervals and ending with bigger intervals and giving each interval a certain probability. See the empirical distribution in the following table.

Value	**# observations for value**	**Probability**	**Cumulative probability**
12	150	0,2249	0,2249
24	59	0,0885	0,3133
36	34	0,0510	0,3643
48	49	0,0735	0,4378
60	36	0,0540	0,4918
72	19	0,0285	0,5202
84	16	0,0240	0,5442
96	31	0,0465	0,5907
108	20	0,0300	0,6207

120	38	0,0570	0,6777
132	20	0,0300	0,7076
144	19	0,0285	0,7361
156	7	0,0105	0,7466
168	18	0,0270	0,7736
180	15	0,0225	0,7961
192	7	0,0105	0,8066
204	8	0,0120	0,8186
216	7	0,0105	0,8291
228	3	0,0045	0,8336
240	11	0,0165	0,8501
252	5	0,0075	0,8576
264	4	0,0060	0,8636
276	1	0,0015	0,8651
288	4	0,0060	0,8711
300	9	0,0135	0,8846
312	1	0,0015	0,8861
324	1	0,0015	0,8876
336	4	0,0060	0,8936
348	3	0,0045	0,8981
360	2	0,0030	0,9010
372	1	0,0015	0,9025
384	3	0,0045	0,9070
396	5	0,0075	0,9145
408	2	0,0030	0,9175
420	3	0,0045	0,9220
432	4	0,0060	0,9280
468	1	0,0015	0,9295
480	5	0,0075	0,9370
492	1	0,0015	0,9385
504	4	0,0060	0,9445
516	1	0,0015	0,9460
540	2	0,0030	0,9490
564	3	0,0045	0,9535
588	2	0,0030	0,9565
600	2	0,0030	0,9595
612	2	0,0030	0,9625
648	1	0,0015	0,9640
672	1	0,0015	0,9655
696	1	0,0015	0,9670
708	1	0,0015	0,9685
720	1	0,0015	0,9700
744	1	0,0015	0,9715
816	1	0,0015	0,9730
840	2	0,0030	0,9760

900	1	0,0015	0,9775
1080	1	0,0015	0,9790
1200	5	0,0075	0,9865
1272	1	0,0015	0,9880
1488	1	0,0015	0,9895
1800	4	0,0060	0,9955
1980	1	0,0015	0,9970
2880	1	0,0015	0,9985
4920	1	0,0015	1,0000

The following table shows the characteristics of our second set, which includes weeks with a number of order lines of 20 or smaller.

Amount of products per order line	
Mean	246
Standard Deviation	376
Kurtosis	9
Skewness	3
Minimum	1
Maximum	2520
Count	381

For this data, we conclude that it still is hard to find a theoretical distribution due to a high skewness and high kurtosis, although the mean, kurtosis and skewness are lower those of the whole dataset. We therefore will try to find an empirical distribution for the data. We will group the data in intervals, starting with small intervals and ending with bigger intervals and giving each interval a certain probability. See the empirical distribution in the following table.

Value	**# observations for value**	**Probability**	**Cumulative probability**
12	58	0,1522	0,1522
24	28	0,0735	0,2257
36	27	0,0709	0,2966
48	20	0,0525	0,3491
60	12	0,0315	0,3806
72	12	0,0315	0,4121
84	11	0,0289	0,4409
96	13	0,0341	0,4751
108	16	0,0420	0,5171
120	19	0,0499	0,5669
132	21	0,0551	0,6220
144	9	0,0236	0,6457
156	5	0,0131	0,6588
168	5	0,0131	0,6719
180	5	0,0131	0,6850
192	3	0,0079	0,6929
216	3	0,0079	0,7008

228	4	0,0105	0,7113
240	4	0,0105	0,7218
252	1	0,0026	0,7244
264	1	0,0026	0,7270
288	4	0,0105	0,7375
300	10	0,0262	0,7638
312	1	0,0026	0,7664
324	4	0,0105	0,7769
348	2	0,0052	0,7822
360	8	0,0210	0,8031
372	2	0,0052	0,8084
384	3	0,0079	0,8163
396	7	0,0184	0,8346
408	2	0,0052	0,8399
420	1	0,0026	0,8425
432	3	0,0079	0,8504
444	2	0,0052	0,8556
456	2	0,0052	0,8609
468	1	0,0026	0,8635
480	2	0,0052	0,8688
504	3	0,0079	0,8766
552	1	0,0026	0,8793
564	1	0,0026	0,8819
600	8	0,0210	0,9029
648	1	0,0026	0,9055
720	4	0,0105	0,9160
792	2	0,0052	0,9213
804	2	0,0052	0,9265
852	1	0,0026	0,9291
900	1	0,0026	0,9318
984	2	0,0052	0,9370
1080	1	0,0026	0,9396
1104	1	0,0026	0,9423
1128	1	0,0026	0,9449
1200	8	0,0210	0,9659
1248	1	0,0026	0,9685
1356	1	0,0026	0,9711
1440	1	0,0026	0,9738
1452	1	0,0026	0,9764
1632	1	0,0026	0,9790
1656	1	0,0026	0,9816
1800	6	0,0157	0,9974
2520	1	0,0026	1,0000

Appendix B.4 Input distribution for the SKU of product group B

Let us start by trying to fit a theoretical distribution for the demand of the SKU of product group B. We first take a look at the average demand, the standard deviation of the demand, the average number of order lines per week, the standard deviation of the number of order lines per week, the average amount of products per order line and the standard deviation of the amount of products per order line in Table B.9.5. We see that the standard deviation of the total demand is almost half the size of the average demand. This is probably due to the high variation in the amount of products per order line, which is five times bigger than the average. A high variability in the amount of products on an order line can influence the CSL-availability.

Table B.9.5 Demand data SKU B

SKU B	Data
Average total demand per week	2512
Std. dev. total demand per week	1235
Average # order lines	190
Std. dev. # order lines	31
Average # products per order line	13
Std. dev. # products per order line	65

Appendix B.4.1 Number of order lines per week

For estimating the demand, we need to estimate the number of order lines per week and the amount of products ordered on one order line. Let us first look at the distribution of the number of order lines. After filtering out the weeks with promotion, we have 44 weeks of data points left. We start by measuring the location, variability and shape of the data. The average number of order lines per week is 190 with a standard deviation of 31, see Table B.9.6. The data is almost symmetric due to the skewness being almost zero. Likewise, the kurtosis is almost zero and lets us expect a normal shape.

Table B.9.6 Summary statistics #OL SKU B

Number of order lines	
Mean	190
Standard Deviation	31
Kurtosis	-0,5
Skewness	0,2
Minimum	131
Maximum	249
Count	44

From the descriptive statistics of the data, we first assume a normal distribution. We try fitting a normal distribution with mean 190 and standard deviation 31. To test the goodness-of-fit visually, we look at the Q-Q plot of the fitted model distribution versus the real demand data. It can be seen, that the graph is almost a straight line, which lets us expect that the normal distribution fits well.

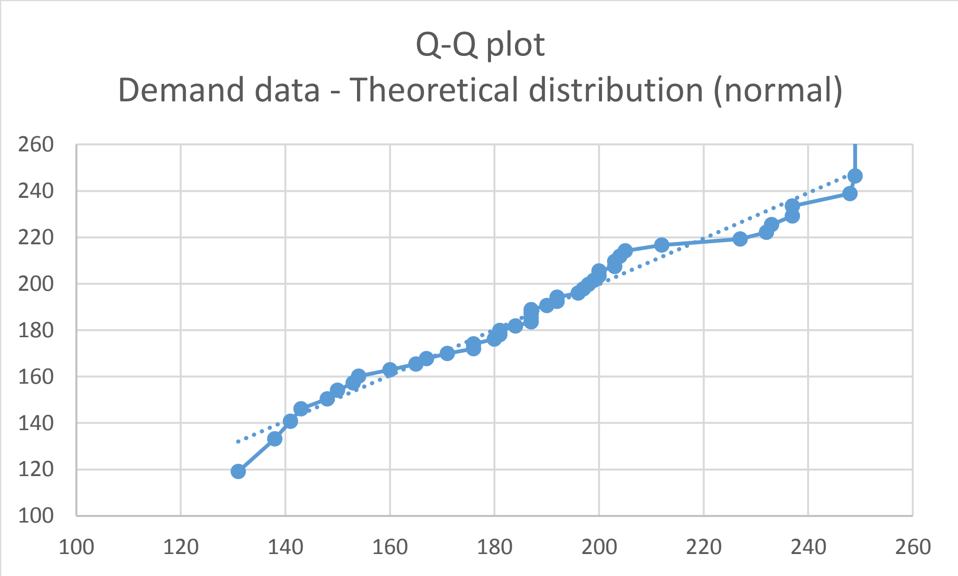

Moreover, we perform a chi-square test to test if the normal distribution fits the data well, described in Chapter 6.6.2 in (Law, 2007). The null hypothesis is that the normal distribution fits our data well. We use the formula for the test statistic $\chi^2 = \sum_{j=1}^{k} \frac{(N_j - np_j)^2}{np_j}$, where N_j is the amount of observations which are in bin j, n the total amount of observations, p_j the expected proportion of the data which fall in the j^{th} interval and k the amount of intervals, which we choose to be 20. From this formula we get the value 14,8 which is a smaller value than the value from the table with a degree of freedom of 19 and a confidence interval of 95%, which is 30,1. Thereby, we conclude that we do not reject the null hypothesis, which states that the normal distribution is a good fit. Therefore, we use the normal distribution with mean 190 and standard deviation 31 for estimating the amount of order lines per week. In the model, we add the condition that the amount of order lines can never be below zero.

Appendix B.4.2 Amount of products per order line

Next we try to estimate the amount of products per order line. Let us look again at the summary statistics. The average amount of products per order line is 13 with a standard deviation of 65. The data is skewed to the left and spread out widely. This, we can read from the values of the skewness and kurtosis in Table B.9.7.

Table B.9.7 Summary statistics amount of products per OL SKU B

Amount of products per order line	
Mean	13
Standard Deviation	65
Kurtosis	221
Skewness	13
Minimum	1
Maximum	1745
Count	11004

With this data, it is impossible to create a smooth histogram due to the high kurtosis and skewness. It is interesting to see that in 77% of the cases an amount of 3 products is ordered. The rest of the order sizes varies between 1 and 1745, which is quite a large range. We take a look at order patterns of different customers, to see if it is possible to divide the customers into groups and analyze the order pattern of the groups. One hypothesis we test is if bigger customers order less frequently, e.g. one order line per month at the end of the month. However, this seems not to be the case. We try to divide customers up into different groups. However, we always remain with at least one customer group with a really large variation in order sizes.

For this data, we conclude that it is hard to find a theoretical distribution due to a high skewness and high kurtosis. We therefore will try to find an empirical distribution for the data. We will group the data in intervals, starting with small intervals and ending with bigger intervals and giving each interval a certain probability. See the empirical distribution in the following table.

Interval	Weighted average # products (multiple of 3)	# observations in interval	Probability	Cumulative probability
[1,3]	3	8628	0,7841	0,7841
[4,6]	6	1330	0,1209	0,9049
[7,9]	9	121	0,0110	0,9159
[10,12]	12	173	0,0157	0,9317
[13,15]	15	30	0,0027	0,9344
[16,18]	18	24	0,0022	0,9366
[19,21]	21	19	0,0017	0,9383
[22,24]	24	25	0,0023	0,9406
[25,27]	27	9	0,0008	0,9414
[28,30]	30	52	0,0047	0,9461
[31,33]	33	1	0,0001	0,9462
[34,36]	36	46	0,0042	0,9504
[37,39]	39	0	0,0000	0,9504
[39,42]	42	19	0,0017	0,9521
[43,45]	45	36	0,0033	0,9554
[46, 60]	51	63	0,0057	0,9611
[61, 80]	75	48	0,0044	0,9655
[81, 99]	90	70	0,0064	0,9718
[100, 120]	108	63	0,0057	0,9776
[121, 160]	144	61	0,0055	0,9831
[161, 200]	180	35	0,0032	0,9863
[201, 299]	255	61	0,0055	0,9918
[300, 399]	330	25	0,0023	0,9941
[400, 499]	438	23	0,0021	0,9962
[500, 599]	555	11	0,0010	0,9972
[600, 699]	642	11	0,0010	0,9982
[700, 999]	828	10	0,0009	0,9991
[1000, 1499]	1233	8	0,0007	0,9998
[1500, 2000]	1671	2	0,0002	1,0000

Appendix B.4.3 Result of modeling

We now model the distributions in the simulation model described in Chapter 5. We make a run of 5000 weeks and compare the average demand and standard deviation of demand to the real data. We see the results in Table B.9.8. The number of order lines per week and the amount of products per order line are distributed as expected and very close to the summary statistics of the real data. However the standard deviation of the total demand deviates by almost 20% from the real standard deviation of the demand. From this, we can conclude that our assumption, that the number of order lines per week and the total amount of products per order line are independent, is not true. We expect some correlation in the data. In the next section, we test how the data is correlated.

Table B.9.8 Input distribution validation SKU B

SKU B	Data	Model output	Percentage of deviation
Average total demand per week	2512	2537	1%
Std. dev. total demand per week	1235	1005	-19%
Average # order lines	190	190	0%
Std. dev. # order lines	31	31	0%
Average # products per order line	13	13	0%
Std. dev. # products per order line	65	65	0%

Appendix B.4.4 Correlation Analysis

The results of modeling the demand suggest that there is correlation in the demand, which is not included in the model. Probably our first assumption that the number of order lines and the number of products per order line are independent is not true. It could be that if the number of order lines is high in a week, that a lot of small customers place order lines, which would result in a low average requested quantity per order line. We therefore test the correlation between the average requested quantity and the number of order lines, see Table B.9.9. We see that there is no correlation between those two factors. Furthermore, we see that the number of order lines is slightly correlated with the total demand, but the average requested quantity per order line is almost perfectly positively correlated. The amount of order lines in the current week is only slightly negatively correlated with the amount of order lines in the previous week. Also the number of order lines from 2, 3 or 4 weeks before do not have a strong correlation. Furthermore, the correlation between the total demand of the current and previous week is almost zero. The same holds for the average requested quantity.

Table B.9.9 Correlation tests SKU B

Correlation Tests	Correlation
Avg. Requested Quantity & #OL	0,00
#OL & Total demand without promotions	0,32
Avg. Requested Quantity &Total demand without promotions	0,94
#OL & #OL previous week	-0,28
Total demand without promotion & previous week	-0,05
Avg. Requested Quantity & previous week	-0,02
#OL & #OL 2 weeks before	-0,05
#OL & #OL 3 weeks before	-0,08
#OL & #OL 4 weeks before	0,16

In addition to that, we split the data in months with 4 weeks and months with 5 weeks. For each month, we test if there is a significant difference between the weeks for different factors through ANOVA analysis. We first change all data for weeks with promotion to the respective average. We look at differences between the weeks for the total demand and the number of order lines. The result of the ANOVA test lets us conclude that there is no significant difference between the weeks. That means that there is no week in which the demand is significantly higher or lower. Therefore we cannot conclude that customers wait with ordering until the last week of the month. We get the same result for the average requested quantity on an order line. This means that big and small customers place orders throughout the whole months. To be sure of this conclusion, we also test the differences between the weeks for the number of order lines below and above 100 products per order line. The test result still concludes no significant difference between the weeks. This lets us reject our hypothesis that big customers order once per month at the end or beginning of the month.

Because, we cannot find logical correlations in the demand it is hard to improve the input distributions of the demand. Using a theoretical and empirical distribution for a simulation is preferred to using the data directly, because it gives us the possibility to generate data which are similar to but not in the current data set. Therefore, we decide to use the theoretical and empirical distributions as demand input in the following. However, we want to know if using the demand data directly as input to the model, results in a output that is closer to reality. In the following table, we see that the demand data, which we call "Real demand", gives even a higher output for a safety stock of two weeks. In practice, the CSL-availability is 82,7%. When reducing the safety stock by one week, the CSL-availability decreases by 17%. So we see that the real demand is quite sensitive to low stock situations. The reason that the modelled demand does not decrease that much in CSL-availability when decreasing the safety stock is that we exclude promotion data when generating the input parameters for the distribution functions.

CSL-availabilty	2 weeks SS	1 week SS
Real demand	99,2%	82,2%
Modeled demand	97,3%	89,8%
	1,9%	-7,6%

Appendix B.5 Sensitivity Analysis SKU A

In the following tables, we show results for different settings for the forecast, lead time variability, capacity, the percentage of orders that are cancelled if stock is unavailable and different orders of serving customers for SKU A. The simulation run length is 5000 weeks. The initial inventory level is one week of safety stock and the first nine weeks have a replenishment size of the average demand. For this SKU, the maximum capacity is 10176 if not stated otherwise. We do the analysis for safety stock levels of one and two weeks of average demand.

Forecasting	Value	2 weeks SS	1 week SS
-5%	3549	97,4%	94,4%
-4%	3587	97,7%	94,9%
-3%	3624	97,9%	95,5%
-2%	3661	98,2%	95,6%
-1%	3699	98,4%	96,0%
0%	3736	98,5%	96,1%
1%	3773	98,7%	96,5%
2%	3811	98,8%	96,7%

For the sensitivity analysis of the lead time variability in the next table, we assume that orders that cannot be fulfilled in the beginning of the week due to the lead time variability are not cancelled, because they can be delivered a few days later.

Chance Late	20%		30%	
OL Affected	2 weeks SS	1 week SS	2 weeks SS	1 week SS
0%	98,5%	96,1%	98,5%	96,1%
5%	98,3%	95,6%	98,1%	95,7%
10%	97,8%	95,1%	97,6%	94,8%
15%	97,5%	94,2%	96,8%	93,2%
20%	96,9%	93,7%	96,1%	91,9%

25%	96,6%	93,0%	95,1%	90,9%
30%	96,0%	92,1%	94,1%	89,4%

The following table is a sensitivity analysis of lead time variability with the assumption of 70% of cancellations also for orders in the beginning of the week, which could be delivered just a few days later.

Chance Late	**20%**		**30%**	
OL Affected	2 weeks SS	1 week SS	2 weeks SS	1 week SS
0%	98,5%	96,1%	98,5%	96,1%
5%	98,4%	95,8%	98,1%	95,5%
10%	97,9%	95,7%	97,7%	95,0%
15%	97,8%	95,1%	97,1%	94,3%
20%	97,2%	94,3%	96,4%	92,9%
25%	97,1%	93,7%	96,0%	92,1%
30%	96,7%	93,2%	95,2%	91,3%

Percentage	**Capacity**	**2 weeks SS**	**1 week SS**
150%	15264	98,5%	96,1%
125%	12720	98,5%	96,1%
100%	10176	98,5%	96,1%
75%	7632	98,5%	96,1%
50%	5088	98,3%	96,1%
25%	2544	66,8%	66,8%

Cancellation	**2 weeks SS**	**1 week SS**
0%	95,0%	86,1%
10%	96,2%	89,1%
20%	96,8%	91,0%
30%	97,2%	92,6%
40%	97,8%	93,5%
50%	97,9%	94,6%
60%	98,3%	95,6%
70%	98,5%	96,1%
80%	98,7%	96,7%
90%	98,9%	97,2%
100%	99,0%	97,6%

Order	**2 weeks SS**	**1 week SS**
FCFS	98,5%	96,1%
Descending	96,3%	90,7%
Ascending	99,5%	98,6%

Appendix B.6 Sensitivity Analysis SKU B

In this part of the appendix, we show additional results of the sensitivity analysis that we perform in Chapter 5.4. Actual demand data is used as demand input for the simulation model as well as the modeled demand distributions. In the following tables, we show results for different settings for the forecast, lead time variability, capacity and different orders of serving customers for SKU B. The simulation run length is 5000 weeks. The initial inventory level is one week of safety stock and the first nine weeks have a replenishment size of the average demand. For this SKU, the maximum capacity is 5040 if not stated otherwise. We do the analysis for safety stock levels of one and two weeks of average demand.

		Real demand		**Modeled demand**	
Percentage	Forecast	2 weeks SS	1 week SS	2 weeks SS	1 week SS
-5%	2640	91,8%	71,4%	94,3%	80,8%
-4%	2668	93,1%	73,5%	95,1%	82,9%
-3%	2696	94,6%	75,7%	95,8%	84,7%
-2%	2723	96,2%	77,0%	96,3%	86,5%
-1%	2751	97,9%	79,6%	96,8%	88,2%
0%	2779	99,2%	82,2%	97,3%	89,8%
1%	2807	99,9%	84,3%	97,6%	91,0%
2%	2835	100,0%	87,6%	98,0%	92,2%

Chance Late = 20%	**Real demand**			**Modeled demand**
OL Affected	2 weeks SS	1 week SS	2 weeks SS	1 week SS
0%	99,2%	82,2%	97,3%	89,8%
5%	99,1%	81,9%	97,2%	89,6%
10%	99,0%	81,6%	97,2%	89,5%
15%	99,0%	81,4%	97,2%	89,3%
20%	98,9%	81,0%	97,1%	89,1%
25%	98,8%	80,6%	97,0%	89,0%
30%	98,6%	80,4%	97,0%	88,8%

Chance Late = 30%	**Real demand**			**Modeled demand**
OL Affected	2 weeks SS	1 week SS	2 weeks SS	1 week SS
0%	99,2%	82,2%	97,3%	89,8%
5%	99,1%	81,8%	97,2%	89,6%
10%	99,0%	81,4%	97,1%	89,3%
15%	98,9%	80,8%	97,1%	89,1%
20%	98,8%	80,5%	97,0%	88,8%
25%	98,5%	79,9%	96,9%	88,5%
30%	98,3%	79,4%	96,8%	88,2%

		Real demand		**Modeled demand**	
Percentage	Capacity	2 weeks SS	1 week SS	2 weeks SS	1 week SS

121%	6120	99,2%	82,2%	97,3%	89,9%
111%	5580	99,2%	82,2%	97,3%	89,9%
100%	5040	99,2%	82,2%	97,3%	89,8%
89%	4500	99,2%	82,2%	97,2%	89,6%
79%	3960	99,2%	82,1%	97,0%	89,2%
68%	3420	99,2%	78,4%	96,5%	87,9%
57%	2880	84,4%	33,4%	93,8%	81,6%
46%	2340	0,3%	0,2%	0,4%	0,4%

	Real demand		**Modeled demand**	
Order	2 weeks SS	1 week SS	2 weeks SS	1 week SS
FCFS	99,2%	82,2%	97,3%	89,2%
Descending	97,4%	77,3%	95,4%	83,0%
Ascending	100,0%	93,4%	98,7%	94,1%

Appendix B.7 Flow chart simulation inventory model

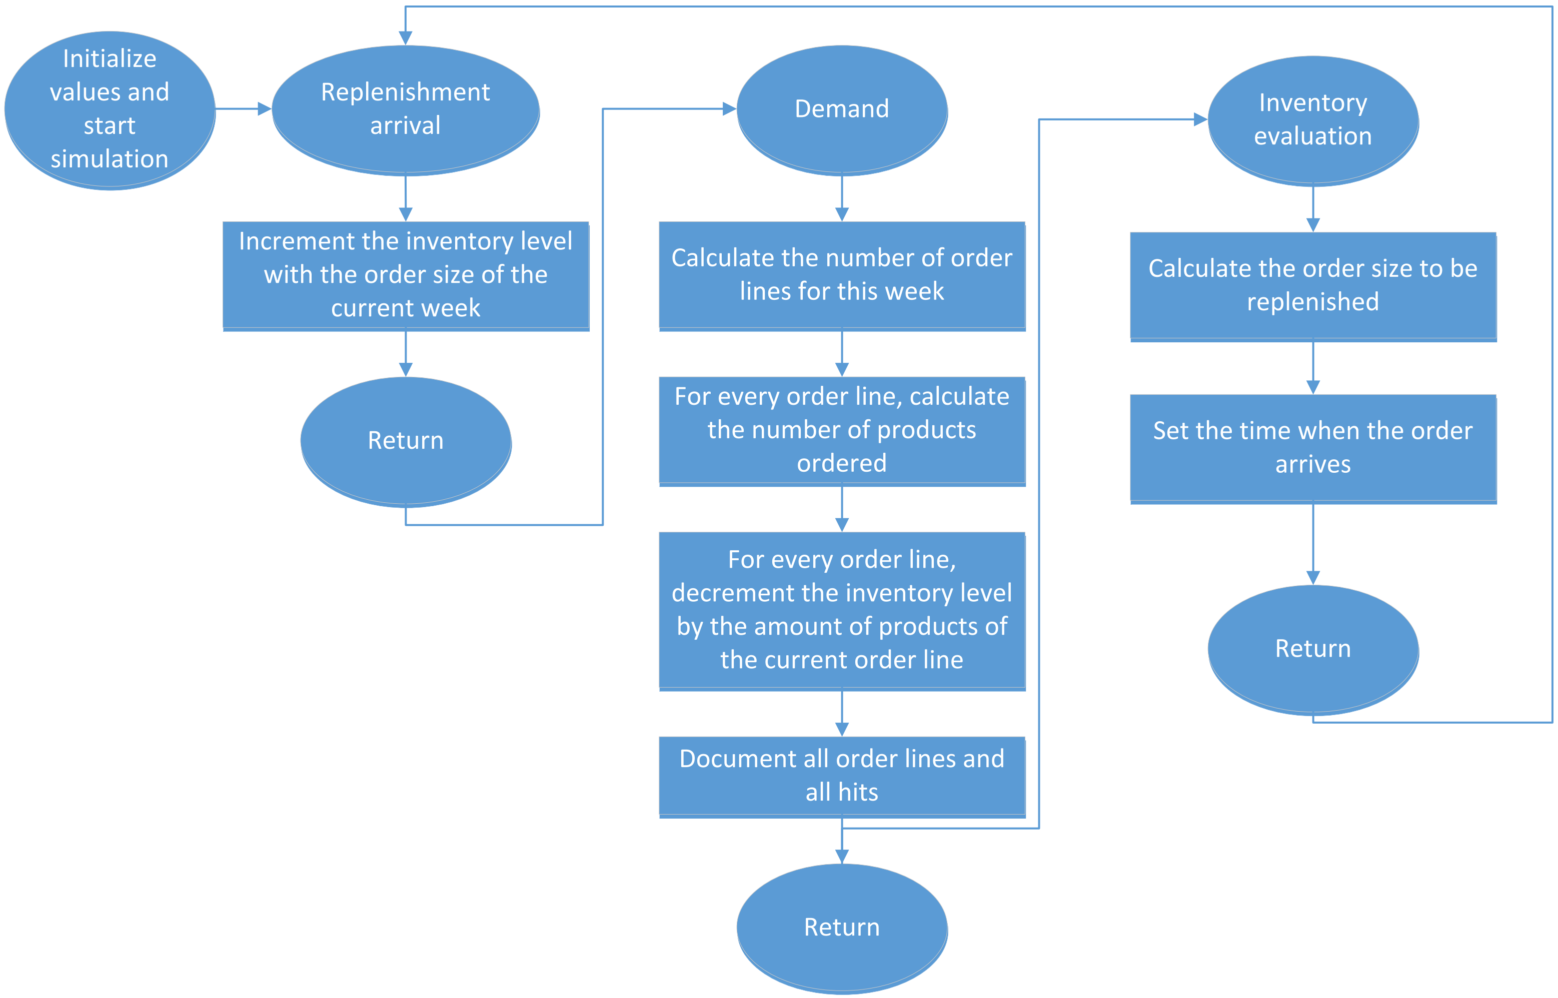

Appendix C Experiment design and results

Appendix C.1 Warm-up period

To determine the warm-up period, we use data created by the simulation model. The input for the demand is generated by the theoretical and empirical distributions that we developed in Appendix B.2.

To identify the warm-up period, we use the method of Welch, which is described in Law (Law, 2007, p. 509). We choose to make $n = 5$ replications which means that we simulate with 5 different initial seed values for the random function and we choose $m = 5000$ which means that we simulate 5000 weeks for each replication. We simulate these replications with the default settings for the replenishment lead time and safety stock, which is 9 and 2 respectively. Y_{ji} is the ith observation from the jth replication, which means that for a given seed value j, Y_{ji} is the average CSL-availability YTD of week i. The following formula shows how Y_{ji} is calculated.

$$Y_{ji} = \frac{\sum_{k=1}^{i} Order\ lines\ deliverd\ OT\&IF\ in\ week\ k}{\sum_{k=1}^{i} Order\ lines\ requested\ in\ week\ k}$$

In the Welch method, $\bar{Y}_i$ is defined as $\bar{Y}_i = (\sum_{j=1}^{n} Y_{ji})/n$ for $i = 1,2, \dots, m$. The following graph shows the moving average for $\bar{Y}_i$ with different window sizes w, which are calculated by the following formula.

$$\bar{Y}_i(w) = \begin{cases} \dfrac{\sum_{s=-w}^{w} \bar{Y}_{i+s}}{2w+1} & if\ i = w+1, \dots, m-w \\ \dfrac{\sum_{s=-(i-1)}^{i-1} \bar{Y}_{i+s}}{2i-1} & if\ i = 1, \dots, w \end{cases}$$

The following graph shows the moving averages of the CSL-availability YTD for SKU B with different window sizes.

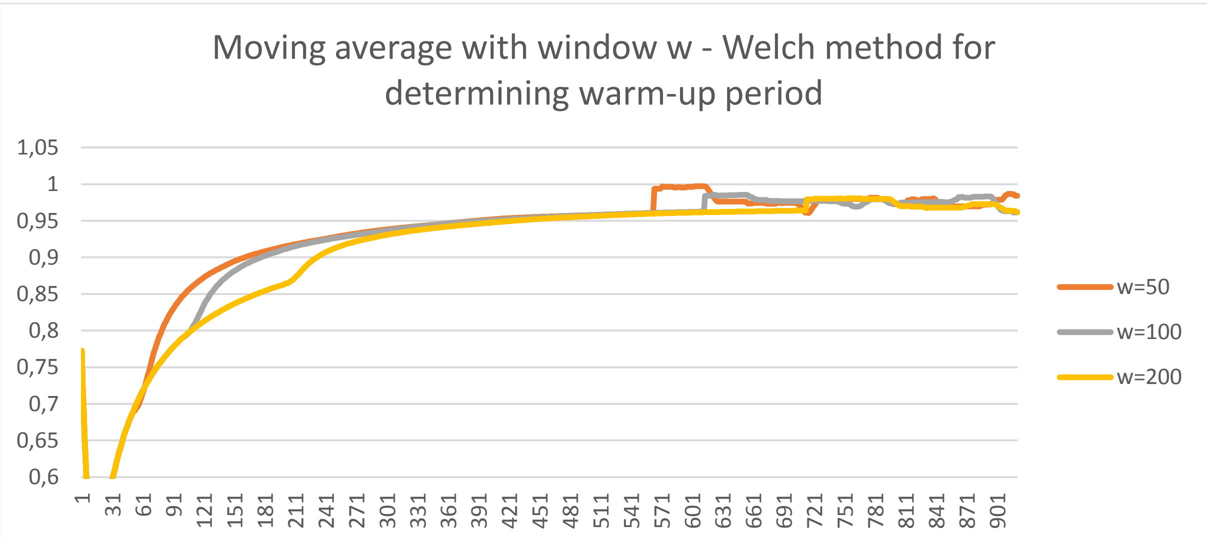

It looks like the steady-state is reached at around 215. We choose the length of the warm-up period to be 300 to have a nice round number.

The following graph shows the moving averages of the CSL-availability YTD for SKU A with different window sizes.

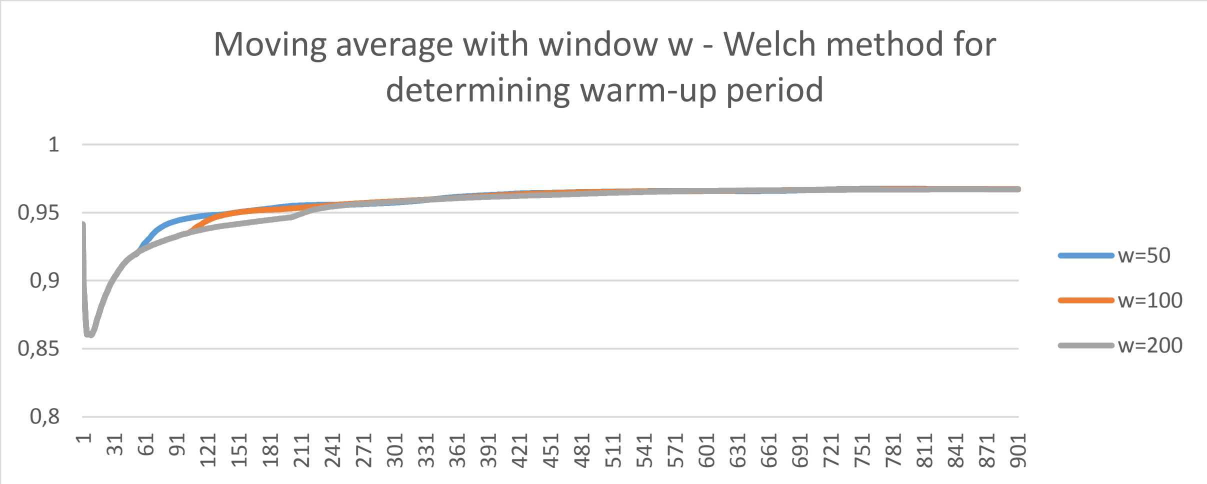

The steady-state in this graph is reach at around 300.

Appendix C.2 Number of Replications per experiment

For determining the number of replications needed for each experiment we follow the "sequential procedure (new replications are added one at a time) for obtaining an estimate of µ with a specified relative error that takes only as many replications as needed" (Law, 2007, p. 505). The mean value that we want to find is the mean CSL-availability YTD.

The first step of the procedure is to make an initial number of replications. We start by making 5 replications of simulation runs with a run length of 5000. We delete the initial warm-up period. We compute the mean CSL-availability YTD of each replication. Furthermore, we calculate the usual confidence interval half-length with the following formula:

$$\delta(n,\alpha) = t_{n-1,1-\frac{\alpha}{2}}\sqrt{\frac{S^2(n)}{n}},$$

with a confidence level $\alpha = 0{,}05$ and $S^2(n)$ the sample variance.

The last step is to check if $\delta(n,\alpha)/\left|\overline{X(n)}\right| \leq \gamma'$, where $\gamma' = \frac{\gamma}{1+\gamma}$ and $\gamma = 0{,}05$ is the relative error of the mean. If the check is positive, then we know the amount of replications needed. The confidence interval is then denoted by

$$\left[\overline{X(n)} - \delta(n,\alpha), \overline{X(n)} + \delta(n,\alpha)\right].$$

In the following tables, the number of replications needed for each experiment are listed. We choose the maximum of all replications needed for a better comparison between the experiments. The number of replications for SKU B is 4 and 2 for SKU A.

Experiment	Replenishment Lead Time (in weeks)	Safety stock (in weeks)	Number of Replications needed (SKU B)	Number of Replications needed (SKU A)
1	9	2	3	2
2	7	2	2	2
3	5	2	2	2
4	4	2	2	2
5	3	2	2	2
6	2	2	4	2
7	9	1	4	2
8	7	1	3	2
9	5	1	3	2
10	4	1	2	2
11	3	1	2	2
12	2	1	2	2

Experiment	Number of Replications needed (SKU B)	Number of Replications needed (SKU A)
13	2	2
14	2	2
15	2	2
16	2	2
17	2	2
18	2	2
19	2	2
20	3	2

Appendix C.3 Results product group A

Air freight

2 weeks SS	Sea Freight	Air Freight	Percentage change	Absolute Value
Sales	355.175,88 €	360.874,44 €		
Cost price all pallets	- 128.310,00 €	- 130.143,00 €		
Cost price transport	- 3.780,00 €	- 9.585,00 €		
Lost sales	- 2.708,93 €	- 217,94 €	92%	2.490,98 €
Absolute IGM	**220.376,95 €**	**220.928,50 €**	**0,3%**	**551,54 €**
COGS	- 132.090,00 €	- 139.728,00 €		
COGS per week	- 2.641,80 €	- 2.794,56 €		
Lead time	10	4		
Net inventory value	- 26.418,00 €	- 11.178,24 €		
WACC (6,7%)	**- 17.700,06 €**	**- 7.489,42 €**	**-58%**	**10.210,64 €**
Total	**202.676,89 €**	**213.439,08 €**	**5,3%**	**10.762,18 €**

1 week SS	Sea Freight	Air Freight	Percentage change	Absolute Value
Sales	346.370,88 €	355.847,04 €		
Cost price all pallets	- 124.644,00 €	- 128.310,00 €		
Cost price transport	- 3.672,00 €	- 9.450,00 €		
Lost sales	- 6.411,31 €	- 2.393,39 €	-63%	4.017,92 €
Absolute IGM	**211.643,57 €**	**215.693,65 €**	**2%**	**4.050,08 €**
COGS	- 128.316,00 €	- 137.760,00 €		
COGS per week	- 2.566,32 €	- 2.755,20 €		
Lead time	9	3		
Net inventory value	- 23.096,88 €	- 8.265,60 €		
WACC (6,7%)	**- 15.474,91 €**	**- 5.537,95 €**	**-64%**	**9.936,96 €**
Total	**196.168,66 €**	**210.155,70 €**	**7,1%**	**13.987,04 €**

Production line at production site North America

2 weeks SS	Sea Freight	Production site budget	Percentage change	Absolute Value
Sales	355.175,88 €	360.874,44 €		
Cost price all pallets	- 128.310,00 €	- 130.143,00 €		
Cost price transport	- 3.780,00 €	- 4.792,50 €		
Lost sales	- 2.708,93 €	- 217,94 €	-100%	2.490,98 €
Absolute IGM	**220.376,95 €**	**225.721,00 €**	**968,9%**	**5.344,04 €**
COGS	- 132.090,00 €	- 134.935,50 €		
COGS per week	- 2.641,80 €	- 2.698,71 €		
Lead time	10	4		
Net inventory value	- 26.418,00 €	- 10.794,84 €		
WACC (6,7%)	**- 17.700,06 €**	**- 7.232,54 €**	**103%**	**10.467,52 €**
Total	**202.676,89 €**	**218.488,45 €**	**146,9%**	**15.811,56 €**

1 week SS	Sea Freight	Production site budget	Percentage change	Absolute Value
Sales	346.370,88 €	355.847,04 €		
Cost price all pallets	- 124.644,00 €	- 128.310,00 €		
Cost price transport	- 3.672,00 €	- 1.890,00 €		
Lost sales	- 6.411,31 €	- 2.393,39 €	-160%	- 6.411,31 €
Absolute IGM	**211.643,57 €**	**223.253,65 €**	**5412%**	**219.203,57 €**
COGS	- 128.316,00 €	- 130.200,00 €		
COGS per week	- 2.566,32 €	- 2.604,00 €		
Lead time	9	3		
Net inventory value	- 23.096,88 €	- 7.812,00 €		
WACC (6,7%)	**- 15.474,91 €**	**- 5.234,04 €**	**103%**	**10.240,87 €**
Total	**196.168,66 €**	**218.019,61 €**	**156,2%**	**21.850,95 €**

Appendix C.4 Results product group B

Air freight

2 weeks SS	Sea Freight	Air Freight	Percentage change	Absolute change
Sales	3.678.675,98 €	3.678.675,98 €		
Cost price all pallets	- 1.270.498,61 €	-1.272.489,50 €		
Cost price transport	- 36.786,56 €	- 245.628,00 €		
Lost sales	- 33.005,81 €	- 6.386,18 €	81%	26.619,64 €
Absolute IGM	**2.338.385,00 €**	**2.154.172,30 €**	**-8%**	**- 184.212,70 €**
COGS	- 1.307.285,17 €	-1.518.117,50 €		
COGS per week	- 26.145,70 €	- 30.362,35 €		
Lead time	10	6		
Net inventory value	- 261.457,03 €	- 182.174,10 €		
WACC (6,7%)	**- 175.176,21 €**	**- 122.056,65 €**	**30%**	**53.119,57 €**
Total	**2.163.208,78 €**	**2.032.115,65 €**	**-6%**	**- 131.093,13 €**

1 week SS	Sea Freight	Air Freight	Percentage change	Absolute change
Sales	3.678.675,98 €	3.678.675,98 €		
Cost price all pallets	- 1.270.503,28 €	-1.272.489,50 €		
Cost price transport	- 36.786,69 €	- 245.628,00 €		
Lost sales	- 141.948,25 €	- 60.588,20 €	-57%	81.360,05 €
Absolute IGM	**2.229.437,76 €**	**2.099.970,27 €**	**-6%**	**- 129.467,49 €**
COGS	- 1.307.289,97 €	-1.518.117,50 €		
COGS per week	- 26.145,80 €	- 30.362,35 €		
Lead time	9	5		
Net inventory value	- 235.312,19 €	- 151.811,75 €		
WACC (6,7%)	**- 157.659,17 €**	**- 101.713,87 €**	**-35%**	**55.945,30 €**
Total	**2.071.778,59 €**	**1.998.256,40 €**	**-4%**	**- 73.522,19 €**

Late customization

2 weeks SS	Sea Freight	Late customization	Percentage change	Absolute change
Sales	3.678.675,98 €	3.678.675,98 €		
Cost price all pallets	- 1.270.498,61 €	-1.340.557,39 €		
Cost price transport	- 36.786,56 €	- 43.661,92 €		
Lost sales	- 33.005,81 €	- 998,19 €	-97%	32.007,62 €
Absolute IGM	**2.338.385,00 €**	**2.293.458,48 €**	**-2%**	**- 44.926,52 €**
COGS KM		-1.238.224,76 €		
COGS per week		- 24.764,50 €		
Lead time		10		
Subtotal		- 247.644,95 €		
COGS	- 1.307.285,17 €	-1.384.219,31 €		
COGS per week	- 26.145,70 €	- 27.684,39 €		
Lead time	10	4		
Subtotal	- 261.457,03 €	- 110.737,54 €		
Net inventory value	- 261.457,03 €	- 358.382,50 €		
WACC (6,7%)	**- 175.176,21 €**	**- 240.116,27 €**	**37%**	**- 64.940,06 €**
Total	**2.163.208,78 €**	**2.053.342,20 €**	**-5,1%**	**- 109.866,58 €**

1 week SS	Sea Freight	Late customization	Percentage change	Absolute change
Sales	3.678.675,98 €	3.678.675,98 €		
Cost price all pallets	- 1.270.503,28 €	-1.340.557,39 €		
Cost price transport	- 36.786,69 €	- 43.661,92 €		
Lost sales	- 141.948,25 €	- 22.133,39 €	-84%	119.814,86 €
Absolute IGM	**2.229.437,76 €**	**2.272.323,28 €**	**2%**	**42.885,52 €**
COGS KM		-1.272.335,64 €		
COGS per week		- 25.446,71 €		
Lead time		10		
Subtotal		- 254.467,13 €		
COGS	- 1.307.289,97 €	-1.384.219,31 €		
COGS per week	- 26.145,80 €	- 27.684,39 €		
Lead time	9	3		
Subtotal	- 235.312,19 €	- 83.053,16 €		
Net inventory value	- 235.312,19 €	- 337.520,29 €		

WACC (6,7%)	**- 157.659,17 €**	**- 226.138,59 €**	**43%**	**- 68.479,42 €**
Total	**2.071.778,59 €**	**2.046.184,69 €**	**-1,2%**	**- 25.593,90 €**

Supplier in Europe

2 weeks SS	**Sea Freight**	**Supplier in Europe**	**Percentage change**	**Absolute change**
Sales	3.678.675,98 €	3.678.675,98 €		
Cost price all pallets	- 1.270.498,61 €	- 1.372.234,18 €		
Cost price transport	- 36.786,56 €	- 36.789,12 €		
Lost sales	- 33.005,81 €	- 2.685,76 €	92%	30.320,06 €
Absolute IGM	**2.338.385,00 €**	**2.266.966,92 €**	**-3%**	**- 71.418,07 €**
COGS	- 1.307.285,17 €	- 1.409.023,30 €		
COGS per week	- 26.145,70 €	- 28.180,47 €		
Lead time	10	5		
Net inventory value	- 261.457,03 €	- 140.902,33 €		
WACC (6,7%)	**- 175.176,21 €**	**- 94.404,56 €**	**46%**	**80.771,65 €**
Total	**2.163.208,78 €**	**2.172.562,36 €**	**0,4%**	**9.353,58 €**

1 week SS	**Sea Freight**	**Supplier in Europe**	**Percentage change**	**Absolute change**
Sales	3.678.675,98 €	3.678.675,98 €		
Cost price all pallets	- 1.270.503,28 €	- 1.448.469,41 €		
Cost price transport	- 36.786,69 €	- 36.789,12 €		
Lost sales	- 141.948,25 €	- 40.857,62 €	-71%	101.090,63 €
Absolute IGM	**2.229.437,76 €**	**2.152.559,83 €**	**-3%**	**- 76.877,93 €**
COGS	- 1.307.289,97 €	- 1.485.258,53 €		
COGS per week	- 26.145,80 €	- 29.705,17 €		
Lead time	9	4		
Net inventory value	- 235.312,19 €	- 118.820,68 €		
WACC (6,7%)	**- 157.659,17 €**	**- 79.609,86 €**	**-50%**	**78.049,31 €**
Total		**2.072.949,97 €**	**0,1%**	**1.171,38 €**

LT reduction at current supplier

2 weeks SS	Sea Freight	LT reduction at current supplier	Percentage change	Absolute change
Sales	3.678.675,98 €	3.678.675,98 €		
Cost price all pallets	- 1.270.498,61 €	- 1.270.582,54 €		
Cost price transport	- 36.786,56 €	- 36.788,99 €		
Lost sales	- 33.005,81 €	- 21.524,27 €	35%	11.481,55 €
Absolute IGM	**2.338.385,00 €**	**2.349.780,19 €**	**0%**	**11.395,19 €**
COGS	- 1.307.285,17 €	- 1.307.371,52 €		
COGS per week	- 26.145,70 €	- 26.147,43 €		
Lead time	10	5		
Net inventory value	- 261.457,03 €	- 130.737,15 €		
WACC (6,7%)	**- 175.176,21 €**	**- 87.593,89 €**	**50%**	**87.582,32 €**
Total	**2.163.208,78 €**	**2.262.186,30 €**	**4,6%**	**98.977,51 €**

1 week SS	Sea Freight	LT reduction at current supplier	Percentage change	Absolute change
Sales	3.678.675,98 €	3.678.675,98 €		
Cost price all pallets	- 1.270.503,28 €	- 1.270.582,54 €		
Cost price transport	- 36.786,69 €	- 36.788,99 €		
Lost sales	- 141.948,25 €	- 112.684,77 €	-21%	29.263,48 €
Absolute IGM	**2.229.437,76 €**	**2.258.619,68 €**	**1%**	**29.181,92 €**
COGS	- 1.307.289,97 €	- 1.307.371,52 €		
COGS per week	- 26.145,80 €	- 26.147,43 €		
Lead time	9	4		
Net inventory value	- 235.312,19 €	- 104.589,72 €		
WACC (6,7%)	**- 157.659,17 €**	**- 70.075,11 €**	**-56%**	**87.584,06 €**
Total	**2.071.778,59 €**	**2.188.544,57 €**	**5,6%**	**116.765,98 €**

Best settings: Dual sourcing

Z_r = 27962	Sea Freight	Dual Sourcing	Percentage change	Absolute change
Sales	3.678.675,98 €	3.678.675,98 €		
Cost price all pallets	- 1.270.498,61 €	- 1.270.615,18 €		
Cost price sea transport	- 36.786,56 €	- 34.595,51 €		
Cost price air transport	- €	- 14.629,50 €		
Lost sales	- 33.005,81 €	- 13.872,06 €	58%	19.133,75 €
Absolute IGM	**2.338.385,00 €**	**2.344.963,74 €**	**0,3%**	**6.578,74 €**
COGS reg.	- 1.307.285,17 €	- 1.319.840,18 €		
COGS reg. per week	- 26.145,70 €	- 26.396,80 €		
Lead time (regular)	11	11		
Lead time (expedited)		6		
Percentage reg. Pallets		94%		
Percentage exp. Pallets		6%		
Net inventory value	- 287.602,74 €	- 282.492,33 €		
WACC (6,7%)	**- 192.693,83 €**	**- 189.269,86 €**	**-2%**	**3.423,97 €**
Total	**2.145.691,16 €**	**2.155.693,87 €**	**0,47%**	**10.002,71 €**

Z_r =25420	Sea Freight	Dual Sourcing	Percentage change	Absolute change
Sales	3.678.675,98 €	3.678.675,98 €		
Cost price all pallets	- 1.270.503,28 €	- 1.270.615,18 €		
Cost price sea transport	- 36.786,69 €	- 27.532,85 €		
Cost price air transport	- €	- 61.724,70 €		
Lost sales	- 141.948,25 €	- 17.273,12 €	378%	24.675,13 €
Absolute IGM	**2.229.437,76 €**	**2.301.530,14 €**	**3,1%**	**72.092,38 €**
COGS	- 1.307.289,97 €	- 1.359.872,72 €		
COGS per week	- 26.145,80 €	- 27.197,45 €		
Lead time (regular)	10	10		
Lead time (expedited)			5	
Percentage reg. Pallets			75%	
Percentage exp. Pallets			25%	
Net inventory value	- 261.457,99 €	- 237.752,94 €		
WACC (6,7%)	**- 175.176,86 €**	**- 159.294,47 €**	**-8%**	**15.882,39 €**
Total	**2.054.260,91 €**	**2.142.235,67 €**	**4,10%**	**87.974,77 €**