Align the inventory of ingredients with the new production planning at Ben & Jerry's

Frank Buisman

MSc Industrial Engineering and Management
University of Twente

17-11-2021

Align the inventory of ingredients with the new production planning at Ben & Jerry's

Master thesis report

17-11-2021

Auteur

Frank Buisman
S 1862502
frank.buisman@gmail.com

University of Twente

Dr. M. C. Van der Heijden
Faculty of Behavioural Management and Social Science

Dr. E. Topan
Faculty of Behavioural Management and Social Science

Ben & Jerry's

Wilco Oldemaat
Teamleader planning

Management Summary

This research is conducted at Ben & Jerry's in Hellendoorn. The plant produces the whole portfolio of Ben & Jerry ice creams for the Netherlands, Europa and some countries in Asia. Due to a high demand in 2020 and 2021, there is a lot of pressure on the production lines. Making full use of the production capacity is very important because Ben & Jerry's has a target service level of 98,5%, which is currently not reached. Therefore, Ben & Jerry's implements since the beginning of 2020 a fixed and repeating production cycle to minimize the changeover time. However, due to the lack of ingredients that are used in the production process, the planners are not able to stick to the fixed and optimal production cycle. In addition to this, the lack of material affects the amount that could be produced, so the planned production runs are shortened. The consequences are that the production planning becomes suboptimal and this affects the line efficiency and the service level. In this research, we only focus on the chunks and swirls (C&S), because those types of ingredients are the major problem. This problem statement results in the following research question:

What improvements to the inventory policies of chunks and swirls should Ben & Jerry's make and implement to be able to stick to the optimal production cycle?

Important uncertainties that influence the production quantity and the availability of C&S's are the supply quantity, supply timing and demand uncertainty. Only the last two uncertainties are taken into account since the supply quantity is not a structural problem. For the supply timing, we have seen that lead times can increase up to 20% longer than expected. Analysing the demand uncertainty shows that almost every finished product is structurally overforecasted. The last important aspect is the complexity of the portfolio. For a production run, all C&S of the finished product should be available and therefore it is important to optimise those C&S together.

Since the situation of Ben & Jerry's is very complex, due to (i) the fixed and repeating production cycle, (ii) many correlations between C&S, (iii) uncertainty in the lead time and (iv) uncertainty in the demand pattern. Therefore, we have not found an existing model that suits the situation at Ben & Jerry's. A simulation study is used to model the processes of Ben & Jerry's. It is a dynamic, stochastic and discrete event simulation. The settings that we can change are the safety stock (SS) and Safety Lead Time (SLT). Those two parameters should cover the uncertainty in the demand pattern and lead time to be able to reach the target service level of Ben & Jerry's. Furthermore, literature states that SLT is usually only preferable to SS when forecasts are accurate and otherwise SS is more robust in coping with fluctuation in demand. Furthermore, we describe several validation and verification techniques that have been used. Since the result of a simulation run should be considered as an estimate of the 'true' outcome, a warm-up period, run length and the number of replications are implemented.

The optimization algorithm optimizes the settings for the SS and SLT. The algorithm exists of two phases. The first phase generates points for different settings of SS and SLT. Based on these points and applying logarithmic regression, we estimate the optimal settings for SS and SLT if we define a target fill rate for each C&S. In the second phase, the knowledge about optimal SS's and SLTs is used and we optimize via a local search to the point in which we reach our target fill rate for the whole portfolio.

Although phase 2 is a local search algorithm, together with phase 1, we have been able to optimise the inventory of C&S. We have been able to increase the fill rate of the portfolio from 95,8% to 99,1%, while decreasing the average inventory value with 15,72%. The current situation at Ben & Jerry's is very inefficient because the SS's and SLT's have not been analysed and optimized when the fixed production cycles were implemented and therefore the settings are not in line with the new planning. In addition to this, Ben & Jerry's prefers high safety lead times to become more flexible. With this, Ben & Jerry's could use inventory that was intended to be used in a later production cycle.

By analysing the results, we can draw the following conclusion:

- The main changes in the optimization process have resulted in lowering the SLT's and increasing the SS.
- For 48 of the 50 C&S we have reached a cost-saving and/or an increase in the fill rate.
- Analysing the results on the level of a single chunk or swirl, shows that SLT's are more in line with the variability in the lead time. SS can cover the uncertainty in both, the lead time and the production quantity.
- If a chunk is underforecasted in one product and overforecasted in another finished product, the underforecast can be compensated by the overforecast of the other finished product.
- It is important to understand how overforecasting affects the performance of Ben & Jerry's. Overforecasting leads to an increase in the average inventory and as soon as the structural overforecasting becomes more accurate or structurally underforecasting, the fill rate for the portfolio decreases immediately. This means, if there are major changes in the forecasts, the optimised settings are not sufficient and this research should be repeated to optimise the settings. If the forecast changes from overforecasting to no structural over or underforecasting, the main increase in the inventory occur in the SLT. If we move from no structural forecast error to underforecasting, the biggest increase happen in the SS.

We recommend Ben & Jerry's to implement the new settings for SS and SLT. This is not only a matter of implementing the parameters in SAP, but Ben & Jerry's should make sure that suppliers and vendor-managed inventories are also implementing the new situation. If important parameters are changing, the inventory of C&S should be optimised again to find the right balance between SS and SLT. Important parameters that affect the availability of C&S are the uncertainty in the lead time, production quantity, changes in the production cycles, new products in the portfolio and the forecast accuracy. Finally, we recommend Ben & Jerry's and Unilever to improve the forecasts, because the structural overforecast mainly result in higher stocks. At last, we believe Ben & Jerry's could improve in actively measuring and using the variability in the lead time, the demand uncertainty, line efficiencies and other parameters that affect the planning and logistics department.

With this research, we have optimised the inventory of chunks and swirl in such a way that Ben & Jerry's is more able to stick to the optimal production cycle, while the costs for inventory are decreased. As a result, both the line efficiency and the service level will increase. In other words, we have aligned the inventory of C&S's with the new production planning.

Acknowledgement

My master thesis called "Align the inventory of ingredients with the new production planning at Ben & Jerry's" is the final test for the Master Industrial Engineering and Management at the University of Twente. The research is executed at Ben & Jerry's in Hellendoorn.

In this acknowledgement, I would like to thank all employees of Ben & Jerry's and the people in the graduation committee. Wilco, you have been a great supervisor, who could always help with finding the right data and contacts within Ben & Jerry's. Matthieu, I appreciate your critical view on my work. Our conversations always resulted in new insights and inspirations. Engin, thank you for making the time to provide me with valuable feedback. I have learned a lot from all of you. Finally, I want to thank my friends, family and housemates for their support and help.

I hope you enjoy reading my thesis.

Frank Buisman
Enschede, November 2021

Abbreviations

Abbreviation	Explenation
BOM	Bill of Materials
C&S’s	Chunks and Swirls
ERP	Enterprise Resource Planning
KPI	Key Performance Indicator
LCG	Linear Congruential Generators
MRP	Material Requirements Planning
RSM	Responsive Surface Methodology
SKU	Stock Keeping Unit
SLT	Safety Lead Time
SS	Safety Stock
ZUN	Special measuring unit in SAP

Table of Contents

1 Introduction

This thesis is the result of the research performed at Ben & Jerry's in Hellendoorn. Section 1.1 provides the relevant background information about Ben & Jerry's. Section 1.2 shows the motivation for the research. Section 1.3 sets the goal of the research. Section 1.4 states the scope and limitations of the research and section 1.5 provides the research questions and the research plan to reach the objective of this research.

1.1 Ben & Jerry's

Ben & Jerry's was founded in 1978 as a scoop shop by two friends, Ben Cohen and Jerry Greenfield. Their ice cream is gaining popularity fast and the first franchise scoop shops were opened in 1981. Soon Ben & Jerry's ice cream was sold in all kinds of places, like (franchise) scoop shops, restaurants and supermarkets. The ice creams are distributed in mini cups (100 mL), multipacks (4 mini cups), pints (500 mL) and a large version of 4,5 litre. The most famous ice cream is the Chocolate Chip Cookie Dough, which was 'born' in 1991.

The strategy of Ben & Jerry's is to make the best ice cream, to have sustainable growth and to take social responsibility. 7.5% of its profit is reserved for charities and most flavours serve a social purpose. For example, 'Cone Together' promotes an inclusive society in which all people come together – no matter the distance and differences – and that all people have the same rights, including refugees (Ben & Jerry's, 2019).

Ben & Jerry's has 4 plants, which are located in the United States, the United Kingdom and the Netherlands. The plant in the Netherlands is located in Hellendoorn and produces the whole product portfolio for the Netherlands, Europe and some countries in Asia, like Singapore, New Zealand and Australia. In total, there are 257 Stock Keeping Units, SKU's, and these are produced on 5 production lines.

In 2000 Ben & Jerry's became a wholly-owned subsidiary of Unilever for 326 million dollars (Smit, 2019, p. 51). By this acquisition, Unilever is taking over the distribution and marketing of products. Furthermore, Unilever is taking care of IT systems and specific analyses like demand forecasts.

1.2 Research motivation

During 2020, the average demand was 20% higher, compared to 2019. In the first quarter of 2021, the average demand was more than 40% higher than in 2020. This increase has several causes. First of all, Ben & Jerry's has a double-digit growth and the plant in Hellendoorn is producing for a big market. In addition to this, the plant in Hellendoorn is one of the two plants that produces the whole product portfolio and therefore a lot of growth happens in Hellendoorn. At last, the corona pandemic has increased the demand for Ben & Jerry's products. Based on these causes, the average demand at the factory in Hellendoorn has increased significantly and this is putting a lot of pressure on the production lines. Especially since fulfilling demand is important for Ben & Jerry's. The target service level equals 98,5%. In addition to this, the demand pattern is getting more robust and unpredictable, due to the upcoming trend of ecommerce.

Based on these arguments, Ben & Jerry's should make full use of its production capacity. In other words, Ben & Jerry's should make sure that the line efficiency is as high as possible. To do this, a production cycle is created in which the total changeover time is minimized. The implementation of this new production cycle started in January 2021 and the expected change over reduction per line is 35%, 22%, 45%, 11% for lines 1 to 4 respectively. However, the planners are not able to stick to the optimal production cycle and therefore the initial motivation of the research is to investigate how this can be facilitated.

By zooming out the whole picture of the problem becomes clear. This is depicted in the problem cluster in figure 1 on page 4. During the first 3 months of 2021, there are too many backorders. This can be derived from the fact that the service level is 93,2% and thus below the target service level. These backorders are caused by the high demand pattern and low line efficiency, which was only 60,8% during the first 10 weeks of 2021. The low line efficiency is caused by not being able to produce the planned amount and by a higher total changeover time because the planners could not stick to the optimal production cycle. Both problems are partly caused by not having the right amount of ingredients that should be used in production. When this happens, the production of the SKU with the lacking ingredients is either shortened or postponed. The lack of ingredients is caused by unpredictability in supply quantity and supply timing. So, suppliers are sometimes not able to deliver or not able to deliver the right quantity and the supply could be delivered late. The other reason why the planners disrupt the optimal production cycle is when a product is becoming out of stock and the planners want to fulfil demand. This can happen because the actual demand deviation from the forecast and the safety stock could not respond to these differences. At last, there are other technical reasons why the actual production output differs from the planned output. The production process of Ben & Jerry's is very automated and there are all kinds of wastes, such as shortages of operators, minor stoppages and idle time, certain defects and rework. The biggest loss in this category is the breakdown and equipment failure. All these errors are tackled by the employees within the production area and are therefore out of scope for this research.

The problem of products becoming out of stock could be solved by optimizing the safety stocks of finished products to anticipate the differences between forecast and actual demand. At this moment, this should not be the priority, since Ben & Jerry's is having too many backorders and a safety stock will immediately be sold, due to the high demand. In addition to this, Ben & Jerry's prefers to have stock at an ingredient level instead of at the level of finished products. This means that when a product is becoming out of stock, the production run of that product is increased in the next run. Therefore, the inventory of ingredients should make sure that Ben & Jerry's could react to the differences between the forecast and the actual demand pattern.

Tackling the problem of the lack of ingredients should enable the optimal production cycle and it will increase the production output since ingredients are not affecting the produced amount anymore. Therefore, solving the problem of the ingredients results in a higher line efficiency, a better output and thus fewer backorders. For this reason, the motivation of the research is to optimize the inventory of ingredients. In section 1.4, the types of ingredients are explained and the scope is determined.

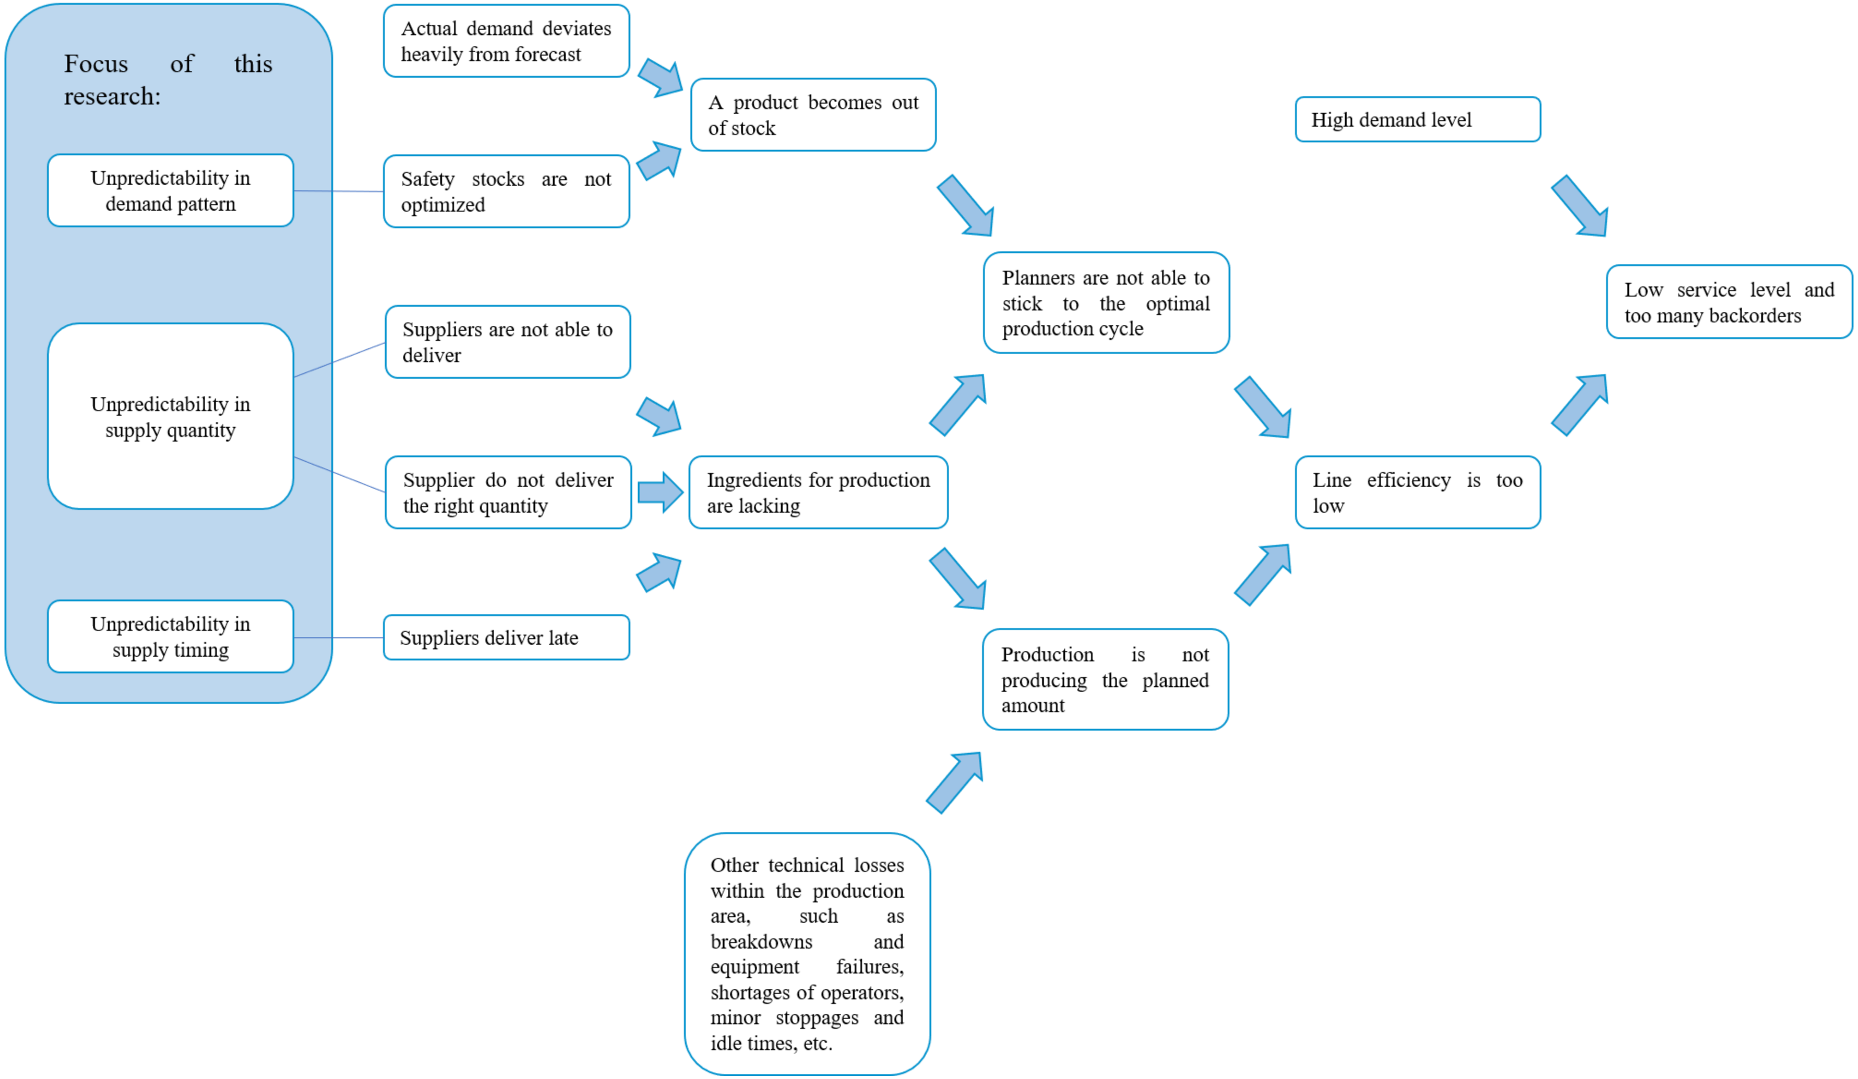

Figure 1: Problem cluster

1.3 Research objective

The goal of the research is to optimize the inventory of the ingredients. As explained in section 1.2, the problem is caused by the unreliability in supply quantity and supply timing. In addition to this, the inventory of ingredients should be able to cover the uncertainty in the demand pattern, such that the next production run can be increased when a product is out of stock.

It could be that Ben & Jerry's should change its inventory policy and that parameters like the lead time, order quantity, safety stock should be adapted to solve the problem. One of the questions that arise is whether the lead times in SAP and MRP are in line with the actual lead times. The goal of this research is to optimize the inventory policies of ingredients and being able to deal with the unpredictability in supply quantity, supply timing and the demand pattern. There should be a right balance between inventory investment and the added value that the inventory of ingredients contributes to the production output.

1.4 Scope and limitations

As described in the former sections, this research is focused on the optimization of the inventory of ingredients. However, there are different kinds of ingredients that are used in production and some are irrelevant in this research. The materials used in production are as follows:

- Packages and plastics: Mini cups, pints, plastic sealing, etc.
- Raw materials: Milk, cream, sugar, water, etc.
- Chunks: Cookie Dough, brownies, chocolate pieces, etc.
- Swirls: All kinds of sauces, like strawberry, caramel, etc.

This research will focus on chunks and swirls (C&S) since these ingredients provide most of the out-of-stocks. These ingredients have a high lead time and an extra limitation is the shelf life that could expire, which mainly happens with the swirls. The raw materials are out of scope because these materials are rarely out of stock. Although these materials are used continuously, the supply is done by local suppliers, which have close contact with Ben & Jerry's. Besides, there are limited possibilities to keep the raw materials as a safety stock. Therefore, the raw materials are out of scope. The packages are also out of scope since those materials already have a safety stock.

Ben & Jerry's has 5 production lines. Four of these lines are producing Ben & Jerry's products and the fifth line is producing the product Magnum After Dinner (MAD). Although this product is not using any C&S, it is good to mention that in the analysis we will not take the statistics of this line into account. Therefore the analysis will only contain the data from Ben & Jerry's line 1 to 4 and not the MAD line.

The demand forecast is provided by Unilever and is considered as given. Unilever is making a sales forecast for each of the countries, which is based on several factors, like historical demand, trendlines, sales, etc. Based on all these input factors, Unilever provides a forecast of each product to Ben & Jerry's. Since Unilever is making these analyses, we consider the demand forecast as given.

The production cycle is recently optimized and therefore not analyzed in this research. The inventory models in this research will be based on the optimal production cycle and nothing will be changed in this cycle.

1.5 Research Question & research plan

The described problem and research objective lead to the following research question:

What improvements to the inventory policies of chunks and swirls should Ben & Jerry's make and implement to be able to stick to the optimal production cycle?

To solve this research question, several sub-questions are formulated to structure the research. All these sub-question cover a different part of the research and require a different kind of analysis and resources.

1.5.1 Context analysis – chapter 2

1. What is the current situation at Ben & Jerry's regarding the inventory of C&S's and how does Ben & Jerry's plan its production?

   In this chapter, we will quantify the problem of ingredients not being available for production. Furthermore, it is important to analyse the product portfolio of Ben & Jerry's. We investigate how the product portfolio of Ben & Jerry's looks like and what C&S's are used. Besides, it is important to know the lead times, shelf life and costs of the C&S's. After that, it is important to understand how the production is planned and thus when the C&S's are needed. We need to understand how the inventory is organized, what policies and parameters influence the inventory position of the C&S's. At last, we analyse the uncertainties in the supply quantity, supply timing and demand pattern. This provides the context of the research. This research question is answered by conducting interviews with different stakeholders and via data from systems like SAP, MRP and Amis.

1.5.2 Literature research – chapter 3

2. What methods exist in literature to cope with the unpredictability in the supply quantity, supply timing and demand pattern?

   After the analysis of the current situation, it is important to explore what methods exist in literature to cope with the unpredictability in the supply quantity, supply timing and demand pattern. The literature study is focused on different types of inventory policies and how these policies can take the unpredictability in supply quantity, supply timing and demand pattern into account. In this way, a theoretical background is created, which can be used in the development of the inventory models for the C&S's.

1.5.3 Solution design – chapter 4

3. Which inventory policies are most suitable for the C&S's, taking the costs and material availability into account?

   In this chapter, the literature is used to develop models that are suitable for the inventory of C&S's. The inventory should be aligned with the optimal production cycle and the shelf life and lead time should be taken into account. The safety stock of the inventory policies should be able to deal with the unpredictability in the supply quantity, supply timing and demand pattern.

1.5.4 Solution test – chapter 5

4. What are the effects and improvements of the proposed inventory policy?

   After setting the right policies, the outcomes of the policies are evaluated. The new policies are compared with the old situation, based on the service level, material availability and the inventory value of C&S's.

1.5.5 Implementation plan – Chapter 6

Based on the found theory in literature, a model is developed and tested. When the results are positive for Ben & Jerry's, the question arises how Ben & Jerry's should implement the findings of the research. It might be that Ben & Jerry's should change their way of working and it could be that certain parameters or settings in IT systems, such as SAP, should be changed. Therefore it is important to provide Ben & Jerry's advice on how to implement the findings.

1.5.6 Conclusion and Recommendation – Chapter 7

Finally, the answer to the research question should be answered. We conclude which inventory policies should Ben & Jerry's implement and what are the corresponding effects and benefits. It is important to mention how Ben & Jerry's should implement the findings and there could be some recommendations for future research.

2 Current situation

The goal of this chapter is to understand the way of working at Ben & Jerry's and to quantify the problem. Section 2.1 describes the portfolio of Ben & Jerry's. Section 2.2 and 2.3 explore the way of working regarding the planning and inventory. Section 2.4 and 2.5 provide an analysis of the scope of the problem and the relevant uncertainties. Section 2.6 provides a conclusion.

2.1 Product portfolio

We will first explore the portfolio on the level of finished products. Thereafter we will zoom in into the C&S's.

2.1.1 Finished products

To analyse the current situation, we first have to understand the product portfolio of Ben & Jerry's. In section 1.1 is stated that Ben & Jerry's is having 257 SKU's. These 257 SKU's consist of 36 different flavours and these flavours are served in 4 different packagings. In addition to this, there are 7 different clusters and every cluster represents a package that contains several languages. For example, one of the clusters is a package specifically for Australia and New Zealand. In addition to this, the finished products for the market in England are transported on a different kind of pallet, the CHEP pallet, and this provides an extra SKU per flavour. The 36 flavours also contain non-dairy products. Therefore, there are 36 different flavours and these are served in different sizes, in different clusters and on different pallets. Therefore, the 257 SKU's can be reduced to 36 flavours.

Of these 36 SKU's, there are some SKU's that are equal to another SKU, but with a specific addition. For example, there are SKU's that represent the normal Cookie Dough while some SKU's represent a Cookie Dough with a chocolate swirl on top.

2.1.2 Chunks & swirls

The chunks and the swirls are additives to the ice creams. In the whole product portfolio, there are 28 chunks and 35 swirls. Every ice cream contains multiple chunks or swirls. That is quite a lot and we have to structure them, based on volume, lead time and shelf life. The costs for all C&S's are in the same range and therefore there is no distinction made based on the costs.

Volume

First of all, there are differences in the volume of the C&S's. Figure 2 shows the cumulative volume per chunk or swirl if the C&S's are sorted from high to low volume. It shows that there are big differences in the volume that a chunk or swirl represents. We assume that the volume of each chunk and swirl is considered high, medium or low if the volume of the chunk or swirl is approximately 10%, between 1 to 5% or below 1% respectively. This is shown in table 1. Furthermore, figure 2 and table 1 show that the two most popular chunks represent 25,7% of the total volume. These two chunks are the cookie dough and the brownie, which are the ingredients for Ben & Jerry's most popular products. The lowest category contains 53,9% of all the C&S's, while it only produces 12,4% of the total volume.

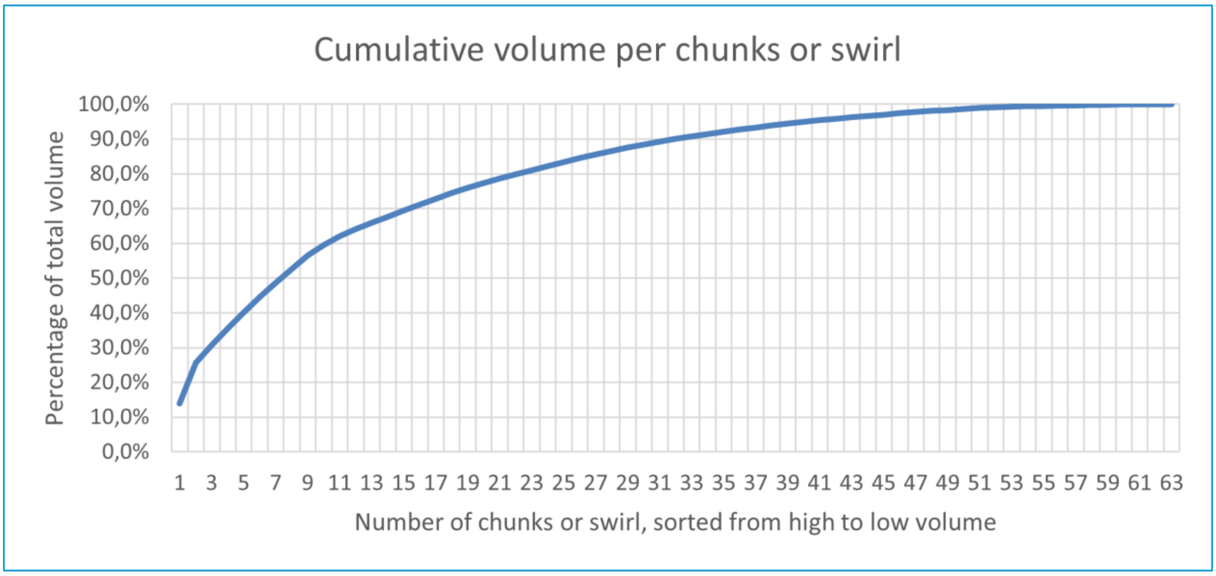

Figure 2: Cumulative volume per chunk and swirl, sorted from high to low volume.

Lead time

All C&S's have different lead times. C&S's are coming from different parts of the world and therefore the variation is large. C&S's produced in Europe have a relatively short lead time, while the C&S's produced in America have a lead time up to 14 weeks. In section 2.1.3, we divide the whole portfolio into different sets and we qualify the lead time to provide a better understanding of the portfolio. We distinguish low, medium and high lead times if the lead time is either less than a week, less than a month or more than a month respectively. The results are shown in Table 4.

Shelf life

Since we are dealing with food, we have to take into account that C&S's could expire. This means that every chunks and swirl has a shelf life. The shelf life differs between 6 days and 720 days. There are only 5 C&S's that have a shelf life which is less than a week. However, these ingredients are not the ingredients that become expired. Analyzing the depreciation costs of 2020 shows that 5 swirls have been expired and these swirls have a shelf life of at least 63 days. These 5 swirls all have a low or medium volume, while the lead time is medium or high. The expired volume is for 4 approximately 1-2% of the total volume. For one swirl 31,4% of the total purchased volume got expired. The reason for this is that the forecast for the end-product got heavily reduced in the second half of the year and the volume was already purchased. Therefore, Ben & Jerry's was having too much in inventory and therefore it got expired. Based on this analysis, we conclude that the shelf life is not a structural problem in the organisation of the inventory.

Note that in this analysis, only the costs are taken into account that are relevant for Ben & Jerry's. Some depreciations are paid by suppliers, for example when ingredients are delivered too close to the expiration date.

Safety stock and safety lead time

Every chunk and swirl contain a safety lead time, set by the material planner. The safety lead time varies between 4 and 46 days and the average safety lead time equals approximately 14 days. This means that Ben & Jerry's plan to have the ingredients on average 14 days before it is needed in production. The brownies currently have a safety stock in addition to the safety lead time, but this is a temporary setting, according to the material planner. Ben & Jerry's implements the safety stock because the brownies come from America, have a high lead time and therefore they want to be more certain about having the brownies on time. According to the material planner, the safety stock makes sure that orders are also placed when there is not a forecast yet. This should prevent shortages of brownies for production on the long term. As described in the paragraph about volume, the brownies are the second most used chunk and therefore important for production. The setting for the safety stocks and safety lead times are determined by the material planner, based on his experience and not on quantitative analysis. In general, we can state that the longer the lead time, the higher the safety lead time. This is the case since longer lead times involve higher risks.

Lead time versus Safety lead time

Since every product contains a lead time and safety lead time, it is interesting to see how these two parameters are related to each other. We expect a positive correlation between the safety lead time and the standard deviation of the lead time. However, as subsection 2.5.2 describes, we are not able to quantify the standard deviation in the lead time. Therefore we can not investigate this relationship. Despite that, we assume the following: As soon as lead time increases, there is more uncertainty in the quantity needed and the timing of supply and therefore we expect that the safety lead time would become larger. In other words, we expect a positive correlation between the lead time and safety lead time. Figure 3 shows the lead time and the safety lead time of all C&S's. The figure does show a positive correlation, but the correlation coefficient is only 0,57. This indicates a moderate correlation (Boston Univeristy, 2019). In addition to this, there are 14 C&S's, from which the safety lead time is bigger than the lead time. All these cases are shown above the line 'safety lead time = lead time' in figure 3. This means that these products arrive earlier at Ben & Jerry's than it takes to order these ingredients. For example, one of the caramel pieces has a safety lead time of 14 days, while the lead time only takes 3 days. This is quite strange. This could imply that the safety lead time and inventory policies are not optimal. There is no explanation for these settings within Ben & Jerry's.

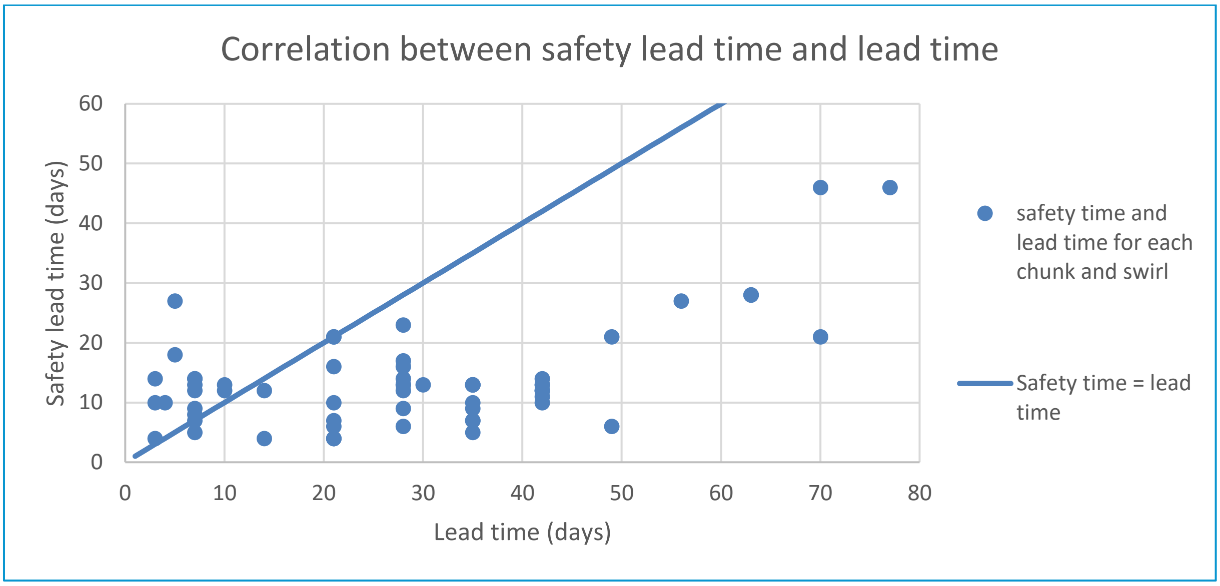

Figure 3: correlation between lead time and safety lead time.

Order quantity

All C&S's have restrictions on the amount allowed to be ordered. The restrictions are imposed by the suppliers and are documented in the contracts. All of the C&S's have a rounding value. The rounding value either states that the number of kilograms should be rounded to a whole pallet or a whole container. In addition to this, most C&S's could have extra restrictions on the height of the order. The restrictions are either a minimum and/or maximum order quantity or a fixed order quantity. However, if the required quantity that should be ordered exceeds the maximum order quantity, SAP creates an extra order. Therefore, the maximum order quantity only restricts the value of a single order. Only a few C&S's do not have any restriction on the number of pallets.

At last, there are no rules that an order should not be placed if e.g. only 10% of the minimum order quantity is needed. If Ben & Jerry's needs the chunk or swirl for the production, it is ordered because the chunk or swirl will eventually be used.

2.1.3 Defining sets – The relationships between C&S's and finished products

To search for the right inventory models, we have to understand the relationship between each chunk, swirl and finished product. For each finished product, we want to know which C&S's are related to the finished product and vice versa we want to know for each chunk and swirl in which finished product the chunk or swirl is needed. In this way, we can explore which inventory systems are related to each other. Therefore, we seek a method to understand the complexity of the portfolio.

The relationship between chunks, swirls and finished products can be found in the Bill of Material (BOM). The BOM is a comprehensive list of items required to create a product (Reedy, 2021). It could be seen as the recipe and shopping list of a finished product.

To do the analysis a VBA code has been written that analyses the whole BOM. The VBA code loops over the whole BOM and it creates separate sets with flavours, C&S's. The flavours, C&S's in each set are mutually exclusive. This means that each flavour, chunk and swirl is in only one set. In other words, the algorithm puts a finished product and the corresponding chunk and swirls in the same set. When one of the ingredients is also needed in another finished product, the finished product and all its corresponding C&S's are added to the set. In this way, each set contains all the chunks, swirls and finished products that are in some way related to each other.

Table 1 shows the created sets and the corresponding characteristics. For example, set 4 contains 1 finished product and the finished product contains 3 C&S's, so 3 items. The chunks, swirls and the finished product could not be found in another set. Set 5 and 6 have the same structure in which multiple items correspond to only one product. The only difference between set 5 and 6 is that there are 2 items needed in the production of the single product. Therefore, table 4 shows that set 4 to 6 contain a single-product and multi-item problem. Set 7 and 8 contain a single product and single item relationship. Set 1 to 3 contain multi-product and multi-item problems. This means that there are more combinations between chunks, swirls and finished products. Therefore, these sets are more complex to solve.

Set	Product (number of products in the set)	Item (number of items in the set)	Volume of total portfolio	Lead time chunks & swirls*
1	Multi-product (31)	Multi-item (32)	59,0 %	Low-high
2	Multi-product (2)	Multi-item (5)	33,4 %	Medium-high
3	Multi-product (2)	Multi-item (4)	3,1 %	Medium-High
4	Single-product (1)	Multi-item (3)	0,9 %	Medium-high
5	Single-product (1)	Multi-item (2)	1,1 %	Medium
6	Single-product (1)	Multi-item (2)	1,1 %	High
7	Single-product (1)	Single-item (1)	1,1 %	Medium
8	Single-product (1)	Single-item (1)	0,3 %	High

Table 1: Set and corresponding set-characteristics

*As explained in subsection 2.1.2: **Low** < week, week < **medium** < month, **high** > month

2.2 Planning

The section regarding the planning is divided into two subsections. The first subsection, called 'planning of the optimal production cycle', elaborates on the planning regarding the optimal production cycle. The subsection, called 'Planning strategy', explains the new strategy of Unilever that influences the way of planning for Ben & Jerry's.

2.2.1 Planning of optimal production cycle

Since January 2021 Ben & Jerry's is implementing its new run strategy. The run strategy has minimized the total changeover time and it provides for every production line the optimal sequence in which the finished goods should be produced. This results in a three-weekly cycle for B&J line 1 and a four-weekly cycle for B&J line 2 to 4. The production is producing 6 days a week and 24 hours a day. On Sunday there is no production and the lines are cleaned on Saturday in the afternoon. Since the production stops at the end of the week, the run strategy also provides a clear structure in which products should be produced in every week of the cycle.

The planners plan every week 8 weeks ahead, which is also depicted in figure 4 for a three-weekly planning. The optimal production cycle provides the input for which products should be produced in each week of the cycle. This means that at the beginning of week 14, the planners plan for week 22. In week 15, the planner plan for week 23, and so on. The other parameters the planners use are the forecast and line efficiency. The forecast is provided by Unilever. The forecast is determined for the upcoming 1,5 years. This forecast is provided on a weekly level. Every 4 weeks and sometimes within the four weeks, the forecast is updated. This happens when supermarkets decide to have Ben & Jerry's in sales and when Unilever is improving the forecast based on new data. The short-term forecast is provided on a daily basis. Based on the forecast and line efficiency, the planners compute how long the production run for each product takes. The production should make sure that the inventory of finished products is increased to the expected demand for the upcoming 9 weeks. This is called the coverage ratio. Because the products from production arrive on average after one week at the different locations of Unilever, the coverage ratio is increased in the week after the production run, see the example in figure 4. Unilever demands a coverage ratio of 9 weeks in this way, two whole cycles could be skipped without Unilever having backorders to their customers.

Thus, the optimal planning is determined with the optimal planning sequence, the forecast and line efficiency.

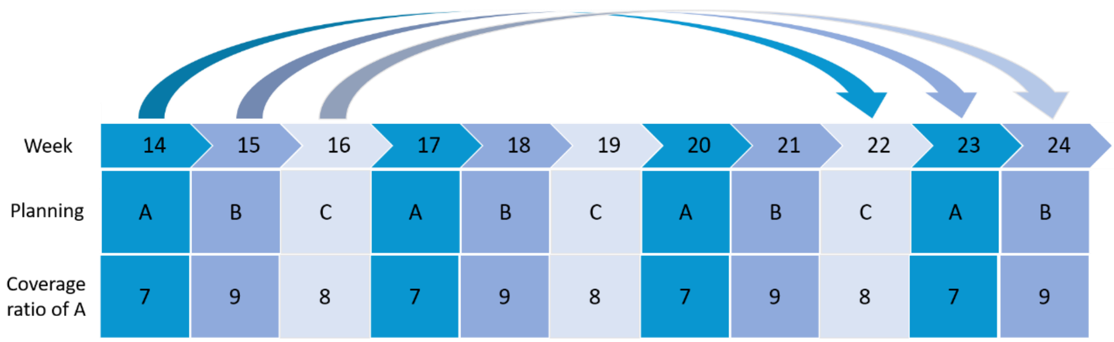

Figure 4: Example of a three-weekly planning cycle

However, the optimal planning cycle is creating an unconstrained planning. It does not take capacity constraints and material availability into account, but in reality, these factors affect the optimal planning. Therefore, the next step is to take these factors into account. The planners try to optimize the planning including these constraints manually. The main solution for the planners lies in the fact that the sequence is repeated every three or four weeks, depending on the line. If week 23 is difficult to plan, it is easy to make switches in the production lengths between week 23, 20 and 17 because in these weeks the same products are produced. Therefore making switches between those weeks enables planners to make full use of capacity and to balance the stocks of different finished products.

Since Unilever has a relatively high coverage ratio, the switches that planners make to balance inventories and to make full use of capacity, will most of the times not lead to backorders from Unilever to supermarkets. However, these switches could affect the product service level of Ben & Jerry’s to Unilever, which should be 98,5%. The product service level, in this case also the fill rate, is measured as the fraction of demand that is served without delays or lost sales (Schalit & Vermorel, 2014). The service level in the past year equals 93,0 % and the service level during the first quarter of 2021 equals 94,5%. Both are below the target service level.

Another solution of lacking materials for production is to switch the sequence of products within a week. What changes the planners are making depends on the situation, such as the number of backorders, the coverage ratio, the importance of the product, the amount of lacking ingredients and for example when the lacking ingredients will arrive.

The way of planning shows that the planners plan 8 weeks ahead, but the planning does not become fixed and is continuously changed. On Thursday, the planning for the next week becomes official and that plan is called ‘plan A’. From that point, the changes in the planning are officially documented, including the reason for the changes. Every change after this moment results in a ‘plan B’, ‘plan C', etc. There is no week that has only a plan A. Most changes in the last week are about changing how much is produced of one or the other product or the change redistributes how much of every cluster within a flavour is produced. All the changes that are made in the weeks before are not documented.

To complete the planning, there is also planning for the long term, which is used for Unilever and suppliers. The planning is mainly focused on the volumes of every product and thus how many ingredients are needed in the upcoming years.

2.2.2 Planning strategy

Unilever has formulated a strategy regarding production planning, currently being rolled out over all plants. The strategy states the ideal situation in which a plant is able to produce the whole portfolio of a plant within one week. In other words, Unilever wishes to produce the products that are sold a lot in the last week. The flexibility in the production should enable Unilever to decrease its inventory positions. To reach the ideal situation, many steps have to be taken, such as decreasing changeover times and downtime and increasing the line efficiency. Another step is that the inventory of ingredients should enable last-minute changes in the production planning. In this way, planners are able to increase the production run of the product that is sold a lot in the last week. This should be the case in the ideal situation, which is a situation that should be achieved in the long run. According to the planning and logistics manager, it is important to take this into account in the model. Parameters, like the coverage ratio and cycle length should be adaptable. The development is part of the so-called ‘master plan’, which is currently being developed at Ben & Jerry’s.

With the current production cycle, it is also important to increase flexibility in the production, and thus the inventory of ingredients should enable this. Since a production run should increase the inventory of finished products to 9 weeks, the production amount equals the difference between the inventory position and the expected demand for the upcoming 9 weeks. Currently, there is already flexibility required with the current philosophy of the coverage ratio of 9 weeks. Currently, the planning of the production is not that flexible due to the limitation of the ingredients, most of the times due to C&S’s. Therefore, the reality is not that each production run increases the coverage ratio to 9 weeks. In addition to this, the inventory of ingredients is getting more important when steps towards the new strategy are taken, such as decreasing the coverage ratio and shortening the production cycles.

2.3 Inventory

The planning for production provides the basis for material planning. Every Wednesday, there is a Material Requirements Planning run. This is based on both the production plan for the upcoming 8 weeks and a forecast for the upcoming year. It provides the number of ingredients that should be available at every moment and it computes how much should be ordered at the suppliers of Ben & Jerry’s. Since the amount of required ingredients is dependent on the production planning and bill of material, the demand for ingredients could be characterized as dependent demand.

The MRP system calculates when the ingredients are needed and when the ingredients should be ordered. The amount and moment of ordering are in line with the planning for production, the restrictions on order quantity, lead time and safety lead time. In addition to this, most C&S’s have a rounded value. This rounds the amount to order to e.g. a pallet or container. There is no optimal order quantity computed that balances the purchase and inventory costs.

To make sure that Ben & Jerry’s could produce, a safety lead time is implemented on the ingredients, as discussed before. This means that there is extra time planned between the expected arrival and when the ingredients are needed in production.

Ben & Jerry’s has limited capacity in Hellendoorn for inventory and this inventory is only used for the ingredients that are needed in the upcoming week. Besides, there are inventories in Bergen op Zoom and Holten. The inventory in Bergen op Zoom is used for frozen products and ingredients, like cookie dough. The inventory in Holten is used for cooling products, like swirls, and for products that should be stored at room temperature. Most inventories are stated on the balance sheet of Ben & Jerry’s. A few inventories are vendor-managed. This means that suppliers are authorised to manage the inventory of the ingredients for Ben & Jerry’s (Yao et al., 2007). It is a way for Unilever to decrease the inventory levels of C&S’s.

2.4 Quantifying the lack of materials

The motivation for this research from the perspective of the planning and logistics department is described in chapter 1. The production department is also familiar with the problem of lacking ingredients. The Operations Manager explained that the lacking ingredients are leading to nonoptimal changeovers and it is causing that certain production runs are shortened because ingredients are lacking. Therefore, experts from both the logistical and manufacturing area are familiar with the problem.

To quantify the problem, we have to take a look at the data. As the problem cluster on page 4 shows, the lack of ingredients is causing two problems. First of all, planners have to deviate from the optimal production planning. Secondly, the production is not able to produce the planned amount. Both problems are analysed with different data and discussed in the following two subsections. You could state that the first method analyses how the planners deal with the lacking materials and the second method analyses how the production department is affected by the lacking materials.

2.4.1 Changes in the optimal planning

The first way to explore the scope of the problem is by analysing the changes in the production plan. As described in section 2.2 about the planning, only the changes in the planning of the last week are officially documented. Therefore, it is only possible to quantify a small part of all the changes. In addition to this, the reasons for the changes are often not stated and if it is stated, the reason is not very specific. One of the possible reasons for changes is that a production amount is decreased due to material problems. It does not say which material is lacking or what kind of material, like C&S's. This makes it hard to quantify the problem this way.

The other way to quantify the problem is via comparing the optimal planning cycle with the realised planning. This can only be based on the planning data of 2021 since this is the period in which the optimal planning is implemented. In this analysis, it seems very easy to compare what should be planned and what is planned. However, as soon as the planning deviates it is barely possible to understand all the reasons behind the changes. This is the case because only in the last week the changes are documented and planners can't make a reconstruction of all the changes. This is caused by the way the planners are dealing with the capacity and material availability constraints. As soon as these parameters are affecting the optimal cycle, the planners are dealing with the issue manually. As described in section 2.2 about the planning, the planners try to solve this issue by transferring production hours to another cycle. In the worst case, certain production runs are transferred to another week and/or to another production line.

The main problem of both methods could easily be explained with an example. Assume a three weekly production cycle, as shown in figure 2 on page 9. Besides we assume that a problem occurs in the supply quantity for one of the products that is produced in week 20, let's say product A. The planners adjust the planning for this week by decreasing the planned amount of product A and they will increase the production hours of the other products in that week, to make full use of the capacity in that week. Besides, the decreased production hours for product A in week 17 are compensated in for example week 20. This means that the coverage ratio of product A will decrease in week 17 to 20 and vice versa for the other products in that week. Most of the times these changes will not happen shortly before production, so in the last week for production, and therefore these changes are not documented. As you can see, the planners try to stick to the optimal cycle and they try to solve issues by transferring production hours to the same week in another production cycle.

In addition to this, the provided example is very basic and we only look at one product. Analysing all the products, weeks and production lines make the puzzle very complicated to analyse. Besides, the reasoning is complicated since the planners take different measures into account in choosing how to make the adjustments in planning. So there are no clear rules in how to solve certain situations. Therefore, the degree of complexity and the lack of documentation makes it too complex to quantify the problem with this data.

One of the production leaders has analysed the production planning of the specific line, and the conclusion was that in the first 21 weeks of 2021, in at least 12 weeks the planning was suboptimal. Specific reasons and analysis lacking. According to the planners, the optimal planning is affected by the availability of materials almost every week. According to the material planner, the materials that are most often out of stock are the C&S's with a long team time. We can conclude that the optimal planning is not being followed, based on the opinion and findings of several stakeholders within Ben & Jerry's.

2.4.2 Losses in production

The other way to explore the scope of the problem is via analysing the number of production hours that are lost within production due to material losses. This means that the production is not able to produce the planned amount. We will first explore the material losses in 2020 and thereafter a comparison with the first three months of 2021.

During 2020 79,2% of the planned operating time provided a valuable output. Of all the manufacturing performance losses, the technical errors within production provide the biggest losses. The material availability is the second-biggest loss, which is accountable for losing 330 production hours. Which is accountable for 2,5% of the total planned operating time. This is a waste of 4.6 million ice creams that could have been produced. However, two different errors are labelled as a material availability loss within production. First of all, if Ben & Jerry's is not having the ingredients on stock, while the production department is needing the ingredients. The other reason is that Ben & Jerry's is having the ingredients on stock, but there are problems in bringing the ingredients to the production line. An example of this is that the swirls were frozen and could therefore not be pumped into the production process. These situations are also labelled as a loss due to the material availability at the lines. According to planners, the distribution between both reasons is fifty-fifty. Therefore, Ben & Jerry's could have produced approximately 2.3 million pines more in 2020 if the material was not lacking.

Since the optimal cycle was not implemented in 2020, it is interesting to see if the problem becomes bigger or smaller in 2021. Comparing the first three months of 2020 with the first three months of 2021, we see an increase in the losses due to material availability of more than 80%. If we compare the first quarter of 2021 with the whole year of 2020, we see that the material losses in the first three months of 2021 are almost 50% of all the material losses in 2020. This shows that the problem of lacking materials becomes a bigger issue in 2021. However, the production department does not recognize this increasing trend. Therefore it could be that the fifty-fifty ratio is changing. This would mean that production is getting more affected by the logistical process of bringing the materials from the warehouse to the line than not having the ingredients available at Ben & Jerry's. However, this distinction could not be made based on the data.

2.5 Uncertainty in supply and forecast

As described in the first chapter, three uncertainties should be analysed. These uncertainties are the uncertainty in supply quantity, supply timing and demand.

2.5.1 Supply quantity

Suppliers not being able to deliver the right quantity is not a structural problem. It happened that a supplier of non-dairy products was not able to deliver the products fully in line with the allergy guidelines or it recently happened that a supplier temporarily delivered chunks with pieces of plastics in them. Both examples are quite extreme and Ben & Jerry's is not taking any risks regarding the quality of its products. In both cases, it happened that the supply was temporarily stopped. In general, there are no suppliers that are structurally underdeliver.

2.5.2 Supply timing

The supply timing is currently a problem at Ben & Jerry's. According to the material planner, the C&S's that are most often arriving late are the C&S's coming from America. These are the C&S's with the longest lead time. Reasons for the delay vary from bad weather conditions to strikes in the harbour and problems with containers.

We have tried to measure and quantify the actual lead times to be able to compare them with the agreed lead times and the lead times in SAP. The goal was to find the distribution and standard deviation of the lead time. Although we are able to find all changes of a single purchase order in a table, like changes in the arrival date, quantity and place of arrival, we were not able to distract data out of the system to make the analysis. Data specialists of Unilever in Poland and India and a consultancy bureau have tried to distract the data. All attempts have failed because the system does not recognize the information in the table. I could not distract the data, nor get a big dump of all the raw data.

Despite that, the procurement operations managers of Unilever have provided very useful data. With their data, we have compared different lead times:

1 Lead time in SAP: This lead time is used on an operational level to determine when and how much to order.
2 Lead time in the contract: This lead time is the official lead time in the contract. The contract is between Unilever and the supplier.
3 Extra provided information: The total time it takes for a supplier to produce the chunk or swirl and to transport it to Ben & Jerry's. It indicates the time the demand from Ben & Jerry's to the supplier should be known to be able to produce the chunk or swirl and

to transport it without arriving late. We call this the tactical lead time since it is not the lead time that is used on an operational level.

Comparing the lead time in SAP (1) and the contract (2), we see many differences. From the 50 C&S's, only 21 C&S's have the same lead time in both SAP and the contract. There are 14 C&S's of which the lead time in SAP is bigger than the contract lead time. The suppliers who struggle to keep up with the growth of the volumes at Ben & Jerry's have this. In most cases, it is a way to make sure those suppliers have an extra 1 to 3 weeks to meet the demand before Ben & Jerry's uses the order. It is kind of a compensation from Ben & Jerry's to the suppliers. There are 14 C&S's of which the lead time in SAP is smaller than the contract lead time. This happens with suppliers producing outside the Netherlands and which have stock located in the Netherlands. In this way, the operational lead time is smaller than the contract lead time. So the stock in the Netherlands is not reflected in the contract.

Besides, it is interesting to compare extra information (3) with (1) and (2). We see that the tactical lead time is in most cases equal to the contract lead time or bigger than the contract lead time. The most interesting is that there are 8 cases in which the tactical lead time is lower than both the lead time in SAP (1) and the contract (2). This implies that the suppliers promises and operates with a lead time that is longer than the time it takes to make and transport the chunk, which equals lead time (3).

The question is which lead time we should use in the model on a tactical level to determine the safety stock and safety lead time. The most obvious is to use the lead time that is agreed upon with the supplier. However, the model is on a tactical level and it is important to take the uncertainties into account during the whole lead time. For example, a chunk takes 97 days to arrive from start of production to arrival at Ben & Jerry's. However, the supplier agrees on a lead time of 21 days because the supplier has an inventory closer to the factory of Ben & Jerry's for resupply. On an operational level, we would use the 21 days, but on a tactical level, it would not be fair to determine the safety stocks and safety time on those 21 days. Those 21 days would hold if the safety inventory is organised well by the supplier. The point is, this research should determine what the right safety stock or safety lead time should be on a tactical level. After the optimization of this research, Unilever and Ben & Jerry's can discuss the safety stock with the supplier and start implementing it.

Furthermore, they have been able to qualify variations for different lead times. Some C&S's coming from America have a lead time of approximately 10 weeks in SAP, while the agreed lead time is 12 weeks. In addition to this, the procurement operations managers state that the actual lead time varies between 12 and 14 weeks. C&S's with a shorter lead time also have some variation in the lead time. C&S's coming from England also have some variation in the lead time, due to Brexit. For this reason, we have agreed with the procurement managers and the planning and logistics manager to approximate the variability in the lead time to become 20% larger than the 'normal' lead time.

2.5.3 Demand uncertainty:

The forecast is specified for all 257 SKU's and not on a flavour level. This is the case since the ice creams are put in the right package for each country. For my analysis, I aggregate the forecasts and demand patterns to a flavour level. In general, we can state that almost every flavour is being overforecasted on short term. On average, the forecast is 7% higher than the actual demand. This is a conscious choice to prevent shortages.

In contrast to this, the long term forecast, the forecast for the next few months, is underforecasting demand. If look at the development in the forecast, we see that the forecast is constantly increasing up to the point that is overforecasting the demand shortly before the demand occurs. For example, 6 months before demand occurs, the forecast is 82,5% of the level the forecast is shortly before the demand occurs. Three months before, the forecast has increased to 92,8% compared to the forecast shortly before demand occurs. Therefore, we see that the forecasts increase and on the short term almost every flavour is overforecasted.

The long term forecast does not influence the expected production, since the production is planned 8 weeks before the production run. However, the forecast does influence the amount that is ordered for ingredients with a long lead time.

2.6 Conclusion

In this chapter, the current situation within Ben & Jerry's is investigated. The goal was to know the portfolio of Ben & Jerry's (2.1), understand the process regarding the planning (2.2) and inventories (2.3), quantify the scope of the problem (2.4) and understand the uncertainties (2.5).

This chapter shows that the portfolio of Ben & Jerry's is complex in terms of the relationships between different chunks, swirls and finished products. The inventory of C&S's is limiting the planners to a) follow the optimal production cycle and b) make last-minute changes in the amount to produce. In addition to this, the flexibility in the production is getting more important, due to the steps that should be taken to reach the new strategy of Unilever.

The main problem with the inventory of C&S's are the C&S's with a lead time longer than a month. The short term forecast is on average overforecasting demand. In addition to this, the long term forecast is far below the short term forecast, while the short term forecast determines the actual production quantities. The variation in the lead time can increase the lead time to become a maximum of 20% larger than the 'normal' lead time.

The model developed in chapter 4, should make sure the target service level of 98,5% for the whole portfolio is reached. To this, the uncertainty in the demand pattern and the uncertainty in the lead time should be taken into account. Another important aspect to implement in the model is the overforecasting on the short term and a long term forecast that is much lower than the short term forecast. At last, the restriction on the order quantity, like MOQ's, should be taken into account.

At last, the shelf life and the uncertainty in the supply quantity are considered out of scope. The uncertainty in the supply quantity is not a structural problem and the shelf life is not a major problem in the current inventory policies. After building the model, it will be manually checked if the shelf life remains not a problem.

3 Literature research

This chapter reviews the relevant literature that explores what inventory policies exist and how these policies can take the unpredictability in supply quantity, supply timing and demand pattern into account. Firstly, in section 3.1 we explore the working and weaknesses of MRP and ERP. Thereafter, section 3.2 provides an introduction to inventory control policies. Section 3.3 provides literature about safety stocks and safety lead times. Section 3.4 explores the possibilities for simulation models. The last section, 3.5, is about simulation optimization. 3.6 provides the conclusion of this chapter.

3.1 MRP & ERP

The material requirements planning (MRP) converts the master production planning into the planning of all materials used for production. To do so, it needs the Bill of Materials which shows all immediate components and their numbers of units of the parents (Silver et al., 2016). With these input parameters, a detailed schedule for all components and raw materials is given. It determines how much is needed and the moment it is needed to produce the finished products. The demand of components is characterised as dependent demand since it is dictated directly from the master production schedule. The MRP computes when to order and how much to order. To do this, several other input parameters are required, such as the lead time, planning horizon, inventory status of each item and forecasts.

Enterprise resource planning (ERP) could be viewed as a direct extension of the MRP. Whereas the MRP is a primary tool for the production department, ERP is used for the entire firm. Since ERP employs standard MRP logic, the production planning and inventory control in ERP is identical to the functions of MRP. ERP facilitates communication between departments because the entire firm is dealing with the same data via the ERP system. Large ERP systems such as SAP, which is used by Ben & Jerry's, and Oracle provide extra tools for scheduling. However, there are no tools or methods within the system that could optimize the inventory policies. According to Silver et al. (2016), there are several weaknesses in using MRP systems. Four of them are relevant for this research. First of all, within MRP there is less incentive for improvement. People assume the input numbers in MRP as given and therefore parameters, such as the lead times, are not analysed and improved. Secondly, the procedures for safety stocks, as will be described in section 3.3, are based on smooth demand. However, the demand of components is dependent on the Master Production Schedule and it is arithmetic. Thirdly, safety stocks could trigger the MRP to place new orders, but this might not be optimal since there are periods without activity. In other words, the MRP could decide to place orders due to the settings of the safety inventory, but because the optimal production cycle dictates when a product is produced and thus when the ingredients are needed, it could be more optimal to order the ingredients a week later. At last, when multiple components are needed for production, the individual components should not be treated in isolation. If one ingredient is lacking, the production is cancelled.

3.2 Inventory control policies

We search for a mathematical model which represents the actual system of Ben & Jerry's. We search for inventory models in which the demand of a particular component, in this case, chunk or swirl, is dictated by the production schedule of the finished products. The production schedule of the finished products should follow a fixed and repeating cycle. Since almost every finished product contains multiple C&S's, we search for a multi-item inventory system. The complexity increases when more components are required to assemble a finished product (Silver et al., 2016). In addition to this, there are parts of the portfolio in which there are multiple links between different components and different finished products. Therefore, the complexity is magnified. We call this a multi-item, multi-product system. The literature we aim for should take the uncertainty in both the demand and lead time into account. Parameters like safety stock and safety lead time should be optimised to reach a target service level of 98,5% for the finished products.

Most literature provides information about single-item inventory systems for finished products and independent demand. Besides, problem statements about multi-item systems mostly explore the possibilities for decreasing the purchase costs when different items are ordered at the same supplier or via the same transportation mode. Furthermore, we find literature about multi-items in MRP situations, either focussing on the planning or the inventory. It is hard to find a model that takes all requirements into account or that could be adapted to be suitable for the situation at Ben & Jerry's. According to Law (2014), if systems are highly complex, so that valid mathematical models of them are themselves complex, precluding any possibility of an analytical solution, the model must be studied through Simulation.

Simulation is the process of designing a model of a system and conducting experiments with this model for the purpose either of understanding the behaviour of the system or of evaluating various strategies (Shannon, 1975). It is a method to model reality, experiment with different settings and to support decision making. In subsection 3.4 we will elaborate on how simulation techniques could be used. Before we elaborate on the different types of simulation models, we need to understand what kind of inventory policy Ben & jerry's implements and what theories could be used in the simulation to cover the uncertainty in demand and lead time.

Based on the method of working at Ben & Jerry's, described in section 2.3, we can conclude that Ben & Jerry's has a (R, s, Q) policy. The R states for the review period, which is one week since the MRP run is every Tuesday. The reorder point s equals the 'forecast during the review period, lead time and safety lead time' + the safety stock – 1. The minus 1 should make sure that if there is not enough inventory to cover the expected demand for the chunk or swirl plus the safety stock, an order should be placed. The amount that is ordered, Q, should take the restrictions into account, such as the rounding value, minimum and maximum order quantity and the fixed order quantity.

3.3 Safety stock and safety lead time

In general, it is necessary to consider safety lead time (SLT) or safety stock (SS) when uncertainty is involved. Safety inventory buffers against uncertainty in both demand and supply (Chopra & Meindl, 2015). SLT means scheduling orders for completion slightly ahead of the required time (Silver et al., 2016). SS is the amount of inventory kept on hand, on average, to allow for the uncertainty of demand and supply in the short run. In other words, the SS is the net stock just before the replenishment arrives. Both safety inventories result in higher inventories, thus higher holding costs, and it reduces the risk of shortages (Dolgui & Prodhon, 2007).

Many industries using MRP argue that SS are not appropriate in a dependent demand situation. It would be better to avoid shortages and excess inventories by adjusting lead times and shifting priorities of orders (Silver et al., 2016). The most common approach is to provide SS for finished products and for items that are at a bottleneck and to set SLT for raw materials (Silver et al., 2016). According to Buzacott and Shanthikumar (1994), SLT is usually only preferable to SS when it is possible to make accurate forecasts of future required shipments over lead time. Otherwise, SS is more robust in coping with changes in customer requirements or with fluctuations in the forecast of lead time demand. The next two sections zoom in to the computation for SS and SLT.

3.3.1 Safety Stock calculation

There are different theories about how to compute the right amount of SS. The most famous formula to compute SS that takes the uncertainty of demand and the variation in the lead time into account is shown below with formulas (1) and (2) (Chopra & Meindl, 2015). However, the formula's are applicable for continuous and independent demand. In addition to this, it assumes lead time and demand are uncorrelated.

$$SS = F^{-1}(Z) * \sigma_{L+R} \quad (1)$$

$$\sigma_{L+R} = \sqrt{E(L+R)\, var(D) + [E(D)]^2\, var(L)} \quad (2)$$

In this formula, we take the square root of multiplying the variation in the demand with the lead time plus multiplying the variation in the lead time with the expected demand.

Jodlbauer and Reitner (2012) describe a method to calculate the optimal safety stocks for multi-item stochastic production models with common production cycles. The research makes a trade-off between the service level and the relevant costs, such as the holding and setup costs of production. In addition to this, it does not only determine the optimal amount of SS but is also determines the length of the optimal production cycles. The research does not propose a basic method to calculate the optimal SS, a complex model is created that should be solved by certain algorithms. In addition to this, the variation in the lead time is not taken into account.

There is no research found that fulfils all criteria at Ben & Jerry's, so a multi-item inventory system in which finished products are produced in fixed production cycles. The uncertainties that should be covered should be the uncertainty in the production quantity and the lead time variability.

3.3.2 Safety lead time calculation

Suppliers are expected to deliver most of the time according to the agreed lead time. However, the actual lead time could deviate from the agreed lead time, based on how busy suppliers and transportation companies are. Sometimes the lead time will be larger and sometimes it will be shorter. Taking the uncertainty into account, supply lead times could be stated as a stochastic variable. According to Das and Abdel-Malek (2003), it could be assumed that lead times are normally distributed. There is no clear formula to determine the optimal safety lead times. According to Dolgui and Prodhon (2007), safety lead time is usually k times the standard deviation of the lead time. In contrast with this literature, subsection 2.5.2 about the supply timing states state that the lead time at Ben & Jerry's does not seem to be normally distributed.

3.4 Simulation study

Simulation is a powerful modelling and optimising technique for testing new processes without carrying out actual experiments (Geetha et al., 2020). It is an ideal way to test multiple safety lead times and safety stocks to find the right balance between inventory and not lacking materials used in production. In this way, a service level of 98,5% could be reached.

There are all kinds of simulations. Discrete event simulation is suitable for problems in which variables change in discrete times and by discrete steps (ÖZgün & Barlas, 2009). Continuous simulation state variables are changing continuously in time. The state changes are described as a function of time. Another technique is Monte Carlo Simulation, which uses repeated random sampling of certain input distributions to come up with a set of possible outcomes (2001). The output can be analysed and used to support decision making.

Simulation models are classified along three different dimensions (Law, 2014):

- Static vs. Dynamic Simulation Models: A static simulation model represents the system at a specific time and time plays no role, while in a dynamic simulation the system evolves over time.
- Deterministic vs. Stochastic Simulation Models: A deterministic simulation model does not contain any probabilistic components. Stochastic simulation models have at least some random input component. The outcome of the simulation should therefore be treated as an estimate of the true characteristics of the model.
- Continuous vs. Discrete Simulation Models: In discrete-event simulation, the state variable changes instantaneously at separate points in time. These points are the ones where an 'event' occurs. The event is an occurrence that may change the state of the system. In continuous simulations, the state variable changes continuously with respect to time.

Based on the information above, the simulation model of this research should be a dynamic, stochastic and discrete event simulation. We can simulate multiple years of demand, production quantities, inventory levels and orders. The demand patterns and lead times are stochastic variables. The events, such as the occurrence of demand, production runs and order generations (might) change the state of the simulation. Therefore, the simulation model that is suitable for this research is a dynamic, stochastic simulation with discrete events.

The most difficult step in the process to create the model is to determine which features should be included in the model (International Conference on Simulation in Engineering Education et al., 1994). Missing an essential element would make the simulation useless, while too much detail could also affect the results. Therefore, it is important to start with a basic model and to expand that model. This is taken into account in the development of the model.

3.5 Simulation Optimization algorithm

The goal of the optimization algorithm is to make runs of the simulation model, where each run uses certain settings, in an intelligent manner and to determine eventually a combination of the decision variables that produce an optimal or near-optimal solution (Law, 2014).

Within the simulation optimization approach, two important terms are the factors and responses. The responses are the output performance measures of a simulation run and the factors represent the input parameters and the structural assumptions.

Running different configurations unsystematically is very inefficient and therefore experiments should be carefully designed. Factor screening or sensitivity analysis determines which factors have the greatest effect on a response. After the analysis which factors are important and how they affect the responses, Response Surface Methodology (RSM) could be applied to optimise the simulation model. We will now zoom in on each facete of this optimization approach, so we zoom in to the experimental design, RSM and logistic regression.

3.5.1 Experimental design

Experimental design is used to determine which factors have the greatest effects on a response, and preferably in the least amount of simulation time. There are different ways to define the experimental design. One is to apply one-factor-at-a-time (OFAT). With this, the effect of one factor is measured while the other factors are fixed. To measure the effect at least two experiments should be executed and therefore we need at least 2n experiments for n variables. This approach is not efficient in terms of run time and it does not provide insights in the interactions between different factors.

The 2^k factorial design is applicable in an early phase of the simulation and is more efficient than OFAT. It requires to define two levels for each factor and then to run each of the 2^k possible factor combinations. By using two levels of each factor, the responses are assumed to be linear over the range of the factor. By comparing the outcomes of all 2^k experiments, not only the relationship between factors and responses can be found but also the interaction between factors can be measured.

The problem with 2^k factorial design is that the design becomes unmanageable very quickly if the number of factors increases. The solution might be to apply fractional factorial designs. It states that only a subset of size 2^{k-p} is considered in the experimental design. For every variable left out of the subset ($p = 1, 2, \ldots$), the number of experiments is divided by two. With this approach, the focus is only on the important factors and not on all the possibilities ($2^{k-p} < 2^k$).

3.5.2 Response Surface Methodology & Regression

A way to optimize a simulation model is to apply the Responsive Surface Methodology (RSM). RSM consists of the following two parts:

1. Predict the model response for system configurations that were not simulated, since the execution time for the simulation is large.
2. Find a combination of input-factor values that optimizes a response.

To relate the model responses to the input factors, some form of regression should be applied. For example, let E[FR(SLT, SS)] denote the expected fill rate of a chunk for particular values for the SLT and SS. For a 2^2 factorial design, the regression model would look as follows:

$$E[FR(SS, SLT)] = \beta_0 + \beta_{SLT} * x_{SLT} + \beta_{SS} * x_{SS} + \beta_{SLT,SS} * x_{SLT} * x_{SS}$$

In this formula, β is the linear coefficient and it describes the linear transformation of the effect of the factors to the response. The different x's represent the settings for the different factors. This is an example of a linear relationship. Other possibilities are regression models that apply a quadratic, logistic, logarithmic, or any other relationship.

3.5.3 Metaheuristics

There are simulation software programs that provide optimization packages. These packages use metaheuristics like genetic algorithms, simulated annealing and tabu search. Although the packages can be adapted to the situation, there are several problems with these kinds of heuristics. First of all, if the number of decision variables k increases, it is harder to find the optimal point in the k-dimensional space. Besides, it is not possible to evaluate the objective formula by plugging the values of the decision variables into a simple formula. Indeed, an entire simulation run is needed to find an estimate of the outcome. All the simulation runs increase the running time of the simulation optimization approach.

3.6 Verification and validation of the simulation

Verification is the process to ensure a conceptual model is correctly transformed into a correct model with the agreed-upon specifications and assumptions. Validation assures that a model represents the real system to a sufficient level of accuracy (Carson, 2002). Both verification and validation are important for the credibility of the results. Multiple techniques should be applied in different phases of the study. As described by Law (2015), to improve the validity and credibility the following should be executed:

1. Collect high-quality information and data of system
2. Interact with the manager on a regular basis
3. Maintain a written assumptions document and perform a structured walk-trough
4. Validate components of the model by using quantitative techniques
5. Validate output from the overall simulation

To validate components of the model by using quantitative techniques (4), there are three different methods, which are described by Law (2015):

1. Comparison with the existing system. Sargent (2010) describes this as using historical data to check whether the model behaves as the real system.
2. Comparison with expert opinion. According to Sargent (2010), individuals knowledgeable about the system are asked whether the model and its behaviour are reasonable.
3. Comparison with another model. This means that the results of the simulation model are compared with other valid models, such as other validated simulation models or other validated analytic models Sargent (2010).

3.7 Performance measure in a Simulation study

Section 3.4 describes that simulations using stochastic variables provide an outcome that should be considered as an estimate. Therefore, it is important to make sure the results are statistically significant. To make sure the results of the simulation study are reliable, a warm-up period, run length and the number of replications should be determined. We use the procedures described by Law (2014).

3.7.1 Warm-up period

According to Law (2014), if one is trying to determine the long-term or steady-state behaviour of a system, then it is generally advisable to specify a warm-up period for the simulation, that is, a point in simulated time when the statistical counters (but not the state of the system) are reset. This means that, when evaluating the performance of the system, the performance during the warm-up period is not measured. To determine the warm-up period, we make use of the approach by Welch (1983), which is described in the book by Law (2014). The conclusion of the approach can be found in 4.3.1 and the analysis can be found in Appendix A.

3.7.2 Run Length

The simulation for Ben & Jerry's is non-terminating since there are no natural events that specify the length of the run. In addition to this, we are dealing with seasonal demand for the finished products, and therefore we have steady-state cycles. It is important to round the run length to an integer number of years, to make sure there is a balance between low and high seasons.

3.7.3 Replications

To determine the number of replications, welch's approach is applied, described in the book by Law (2014). We apply a fixed number of replications in the simulation. Although the replications are independent and have a different set of random numbers, each replication number in different experiments has the same set of random numbers.

3.8 Conclusion

This chapter provided a review of the relevant literature for this research. Due to the uncertainty in both the demand pattern and lead time, together with the fixed production cycle and the correlations between different C&S's, the situation at Ben & Jerry's is very complex. Therefore, we will use a simulation study to model the situation at Ben & Jerry's and experiment with different settings. The model will be a dynamic, stochastic and discrete event simulation. The settings that we can change are the safety stock and the safety lead time. SLT is usually only preferable to SS when it is possible to make accurate forecasts of future required shipments over lead time. Otherwise, SS is more robust in coping with changes in customer requirements or with fluctuations in the forecast of lead time demand. There is no existing method that proposes the amount of SLT or SS for a manufacturing system with the same conditions as Ben & Jerry's. At last, a stochastic simulation model results in estimates of the 'true' outcome. Therefore it is important to implement a warm-up period, a run length that provides statistical results and the right number of replications. To optimize the simulation, RSM could be applied. This means that output values are estimated via regression and thereafter a combination of input factors should be found that reproduce a certain outcome, in this case, the desired fill rate for the portfolio.

4 Simulation model

This chapter explains the working of the simulation model that is developed for Ben & Jerry’s. Section 4.1 introduces the simulation model and explains the working of one simulation run. Section 4.2 provides the technical description of the simulation model. Section 4.3 states the settings of the simulation, such as the warm-up period, rung length, number of replications and the generation of random numbers. Section 4.4 validates the model for a simulation run with the current settings for safety stock and safety lead time. Section 4.5 describes how a simulation run is analysed and how parameters are changed via a simulation optimization technique.

4.1 Model explanation

The purpose of the simulation study is to optimise the amount of SS and SLT of C&S’s to make sure Ben & Jerry’s is able to reach its target service level of 98,5%. The model is built on a tactical level and it should take the uncertainty of supply timing and the demand into account. To cover those uncertainties, we should optimise the SS and SLT. Eventually, we balance the service level with the average inventory value. In this way, we can make distinctions between various solutions. We do not compare the service levels with the operational costs, like the purchase and inventory costs, because for Unilever and Ben & Jerry’s the tradeoff is between the investment in inventory and the service level.

In general, the following steps will be simulated. We generate on the level of finished products a short-term and long-term forecast and the demand patterns. Based on those two, we can determine what the expected production quantity is, which equals the forecast for C&S’s. In addition to this, the actual production planning is determined and this provides the actual demand for the C&S’s. However, if a chunk or swirl is lacking, the production quantities are adapted. We will keep track of the inventory levels of C&S’s and if needed we can place orders.

In this research, the service level of Ben & Jerry’s in this research is described as a fill rate. The fill rate specifies the percentage of demand which is directly satisfied from shelf (Silver et al., 2016). In our model, we measure a fill rate on three different levels:

1. Fill rate of a single chunk or swirl: The fraction of demand of a chunk or swirl that could directly be satisfied from shelf. The total demand of a chunk or swirl equals the total amount that should have been used in production. The fulfilled demand equals the amount of the chunk or swirl that was not lacking when it should have been used in production. Both the demand and the satisfied demand are measured in kg.
2. Fill rate of a finished product: The fraction of demand from Unilever to Ben & Jerry’s of a single finished product that could be directly satisfied from shelf. The demand and fulfilled demand are measured in ZUN.
3. The fill rate of the portfolio: This is a weighted average fill rate of the fill rates of all finished products. The weight for each finished product is based on the yearly sales, which is measured in ZUN. For Ben & Jerry’s the fill rate of the portfolio equals the service level.

Since Ben & Jerry’s wants its service level to 98,5%, we target the fill rate of the portfolio to be at least 99%. This is the case since other losses affect the service level of Ben & Jerry’s and those losses are not taken into account in our model. It is not necessary that all finished products have a fill rate of at least 99% and therefore the fill rate of the portfolio is a weighted average. Therefore, we aim to reach a fill rate of 99% for the portfolio while minimizing the investment in the inventory.

The stochastic variables in the simulation are the forecasts, demand patterns and the lead times that are realized. These stochastic variables reflect the uncertainty in the demand pattern and the supply timing. Other input parameters are the fixed and repeating production cycles (see subsection 2.2.1), an increasing trend in the forecast when demand is getting closer (see section 2.5.3), the restriction on the order quantity (see subsection 2.1.2). Besides we take the correlations between chunks and swirl into account. The complexity of this is explained in section 2.1.3. In chapter 3, we described the inventory policy of C&S's as an (R, s, Q) policy class. We stick to this inventory policy.

We do not take into account the production capacity of the different lines and other ingredients than the C&S's. This process is explained in detail in subsection 2.2.2, paragraph C. Initially, we do no change the forecast. However, after the optimization, we execute a sensitivity analysis by changing the forecast biases. At last, the model is created for all finished products and C&S's to make sure the correlations and complexity of the portfolio is modelled.

4.1.1 General description

Since the model is on a tactical level, we simulate the demand, production and inventory on a weekly level. Section 4.3 explains that in one simulation run the weeks of in total 5 years are simulated.

Figure 5 shows the different steps that are simulated in a single run, or in other words one replication. A single run consists of three phases: (i) initialisation of the simulation by generating the necessary variables, (ii) the execution of the simulation by looping over all weeks and (iii) evaluating the run.

In phase one, we start by generating the forecast patterns of each finished product. There are two forecasts, the short term and long term forecast. The short term forecast predicts the demand for the upcoming three weeks and the long term forecast predicts demand for 4 and more weeks ahead. Based on one of the forecasts a demand pattern is generated, which is explained in subsection 4.2 in more detail. After setting the forecasts and demand for each finished product, we create the production cycle for each finished product that states when the products should be produced. With this, we can determine the expected production quantity for each finished product and we can translate the expected production to a forecast for each C&S. The long term forecasts for C&S is needed in the process of placing orders. For example, for a Chunk with a lead time of 10 weeks, we need to know what the long-term need is to be able to decide if an order should be placed.

In phase 2, we start looping over the week and in each week the following steps are executed. These steps are also depicted in figure 5. We start by computing for each finished product (given it is a production week for that finished product), what the actual planned production quantity is. The actual planned production quantity is not only dependent on the forecast but also on which product was over or under forecasted in the past production cycle. The actual planned production quantities are translated to the demand for each C&S. The next step is to determine the begin on hand inventory for each C&S, which indicate the material that is available for a production run in that week of the simulation. The begin OH inventory is compared with the demand for the C&S in that week. If chunks or swirls are lacking, the production quantity, and thus the demand for each C&S, is adapted. When all production quantities are adapted in a way that no C&S are lacking, the ending OH inventory is determined and we decide for each C&S if an order should be placed. For the order, we first determine if the order arrives late or not, which is based on a list that states for each supplier what percentage of the orders arrive late. If the order should arrive late, we determine the delay of the arrival of the order. The process of placing orders takes the settings for SS and SLT into account. The planned production that could not be fulfilled is postponed to the next production cycle.

Phase 3 starts when all the weeks of the simulation run are executed. We analyse the fill rate of each chunk & swirl, finished product and the whole portfolio. This e.g. means we continually measure when a C&S affects the production planning. Also, the average inventory values are computed. These are the KPI's we use to adapt the parameters for SS and SLT.

One additional process that happens in phase 2, is that forecasts are updated. Section 2.5.3 stated that forecasts on the short term are structurally overforecasting demand, while the forecast on the longer term is much lower. In other words, when demand is getting closer, the forecasts is increasing until the point we are structurally overforecasting almost every finished product. Therefore, at the beginning of the week, some forecasts are updated and this affects some of the expected production quantities and thus the expected demand for C&S. The expected demand for C&S is used in the process of deciding if an order should be placed. Therefore, this process is important for C&S with a long lead time.

The decision variables, so the SS and SLT, influence when an order is placed and how much is ordered. Based on this, the inventory levels of C&S's are different with different settings. This influences the percentage of the production that is fully fulfilled or partly postponed to the next cycle.

Figure 5: Chart flow of single simulation run

4.1.2 Requirements

First of all, there are some general requirements for the simulation model. First of all, the simulation is created in Excel with Visual Basic of Applications (VBA). The employees at Ben & Jerry's have experience with VBA and therefore they could adjust the model if certain processes change in the future. Secondly, it should be easy to insert new data as an input for the simulation. This is important because many input parameters might change. For example, Ben & Jerry's might adjust its portfolio by introducing new flavours. In addition to this, it is uncertain whether the height and unpredictability of demand for finished products in the current market will remain after Corona. If this is not the case, Ben & Jerry's could decide to run the simulation model with new data, new market conditions and with the changes in the portfolio. Thirdly, the model should be easy to run and the output should be created automatically via a smart algorithm. At last, it is important to transfer the knowledge to the employees of Ben & Jerry's.

Next to the general requirements, there are some technical requirements regarding the model. First of all, most C&S's have a rounding value. That means that the amount that is ordered should be rounded to e.g. pallets or containers. In addition to this, some C&S's have a fixed, minimum or maximum order quantity, which should be taken into account.

4.1.3 Decision variables & KPI's

The decision variables are the safety stock and safety lead time for each individual C&S. Both SS and SLT are discussed in chapter 3. The safety stock will be measured in ZUN, a special measuring unit in SAP which is defined as a bundle of ice cream packages. For the concrete definition see table 2. The safety lead time will be measured in the number of weeks. Both parameters will be optimized in the simulation. The right SS and SLT should make sure that the right quantities of C&S are available for production while having the lowest possible investment in inventory. Therefore, the KPI's are the fill rate of chunk or swirl and the investment in inventory.

Package	100 ml	438 ml	465 ml	4,5 l
1 ZUN equals	12 packages	8 packages	8 packages	2 packages

Table 2: Definition of 1 ZUN for each type of package.

4.1.4 Input variables

Forecast of finished products

There are two types of forecasts available. Both forecasts predict the demand for finished products from supermarkets and scoop shops to Unilever. The first forecast is determined 4 weeks before demand arrives and the other forecast is determined one week ahead. It is provided on a weekly level and the measuring unit is in ZUN. The forecast is used to determine the expected production quantity and is used to generate realistic demand patterns. This will be described in section 4.2, the technical description of the model.

The demand for finished products

The demand for finished products states the demand at Unilever coming from the supermarkets and scoop shops. The demand is provided per week and is measured in ZUN. It is generated based upon one of the forecasts and the generation happens before the model simulates all simulation weeks. The actual demand patterns influence the actual production quantities. This is in line with the strategy of Unilever that plants like Ben & Jerry's should increase the production of finished products that are sold a lot in past weeks. In this way, the coverage ratio of 9 weeks after the production run should be guaranteed.

Bill of Material (BOM)

The Bill of Material translates the production quantities of finished products, measured in ZUN, to the demand for C&S's, measured in kilograms.

Supply lead times

The supply lead times provide information about when an order will arrive. In addition to this, some swirls need time to de-freeze. We have added this time to the supply lead time. Otherwise, swirls could be present on time, but not being able to be used. Therefore, we have to take this into account when ordering the C&S's. As described in section 2.5.2, the lead time can become 20% larger than the 'normal' lead time. Therefore, we apply a uniform distribution between 100% and 120% of the lead time. Orders either arrive on time or late. We do not simulate orders to arrive early. The computation is based on the normal lead time in days, and we will round it up to the corresponding number of weeks.

On Time in Full (OTIF)

For every supplier, we have the percentage of orders that actually arrives on time on average in full. Since we have stated before that there are no structural issues with the suppliers not being able to deliver the right quantity, we assume that the OTIF only describes the percentage of orders that do not arrive late. The OTIF values vary between 51% and 71%. If an order arrives one day later than expected, it is already counted as not on time, although it might not have affected the planning.

Review period

The review period equals one week. As explained in chapter 2, the MRP run is every Tuesday.

4.1.5 Out of scope and assumptions

In our model we have the following assumptions:

- The forecast error is normally distributed. This is also motivated in section 4.3.
- Forecasts are biased. Most finished products are structurally overforecasted. This is also described in section 4.3.
- Lead time is uniformly distributed between 100% and 120% of the normal lead time of a chunk or swirl.
- The forecasts in the long term are lower than the forecast in the short term. This is an input from Unilever and we are not able to change their way of creating the forecasts. This is described in section 2.5.3. This means we increase the forecast when demand is getting closer.
- If a chunk or swirl is lacking for multiple production runs in a week, the backorders are allocated according to a weighted average of the demand of the finished product. For example, we miss 90 kg of a chunk, while we needed 1000 and 2000 kg for the production run of finished products A and B. This means that 1/3 of the demand of the chunk was needed in the production of product A and 2/3 in product B. Therefore, we decrease the production run of A and B by 30 and 60 kg respectively, which corresponds to the weighted average demand (1/3 and 2/3).
- The production that could not be fulfilled is postponed to the next production cycle.
- Each production run increases the inventory of a finished product to the level on which the coverage equals the demand for the upcoming 9 weeks.
- The inventory policy of C&S follows the (R, s, Q) policy class, which Ben & Jerry's is currently using.

The following parts are not included in the model and therefore out of scope:

- Changes in the forecast are initially out of scope. The forecasts are not determined by Ben & Jerry's, but they are developed by the marketing and sales departments of Unilever in different countries. In the sensitivity analysis, we experiment with different forecast biases.
- We will not experiment with different production cycles. However, settings that are likely to be changed in the future, such as the cycle length and coverage ratio, are very easy to change in the model.
- Changes to the inventory of Finished Products are out of scope. We keep track of the inventory of Finished every week because a production run should make sure the coverage ratio is increased to 9 weeks. Since these stocks are located at Unilever, we do taken the inventory into account in the computation of the average inventory value and we do not experiment with different settings for the inventory of Finished Products.
- We only focus on the availability of C&S. Raw materials that are continuously being used, like milk and sugar, are considered out of scope.
- Production capacity is not taken into account.
- Line efficiency is not affecting the output of a production run.

4.2 Technical description

Before simulating all the weeks, four factors are determined for the entire simulation run. These factors are (a) forecast for finished products, (b) the demand for finished products. (c) the production cycle, so the week number in which a product is produced, and (d) the (long term) expected production quantity. Those four factors are described in section 4.2.1. In section 4.2.2 we provide the technical description of the simulation of a week. In the week itself, the following happens: (i) The forecasts are updated, and thus the expected production quantities, (ii) the actual planned production quantity is determined, (iii) the inventory policies of C&S's are executed, (iv) the fulfilled production quantity is determined, which is based on whether or not C&S's are lacking, and (v) orders are placed if necessary.

4.2.1 The initialisation of a simulation run

1. Generate forecast of finished products

First of all, we determine the short term and long term forecast of finished products for the entire simulation run. First, we explain the way we generate the forecasts. After that, we explain how we have created the input forecasts from which the random forecast is generated. At last, we explain how we create the long term forecast.

The forecast data is available for 2020 and 2021 up to and until week 30. Since the Simulation is doing multiple replications per simulation, we do not want to have the same forecast every run. Therefore, we pick a random forecast from the same period, with a random range of 5 weeks. In this way, we can keep track of the seasonal patterns. Since we use the two different forecasts, we will pick the same random week of both forecasts, to keep the correlation between the two different forecast inputs. Figure 6 shows the two different forecast patterns (blue and red) and the actual demand pattern of one of the finished products (green). The figure shows that both forecasts are correlated and therefore it would not be realistic to combine two forecasts from a different week in time.

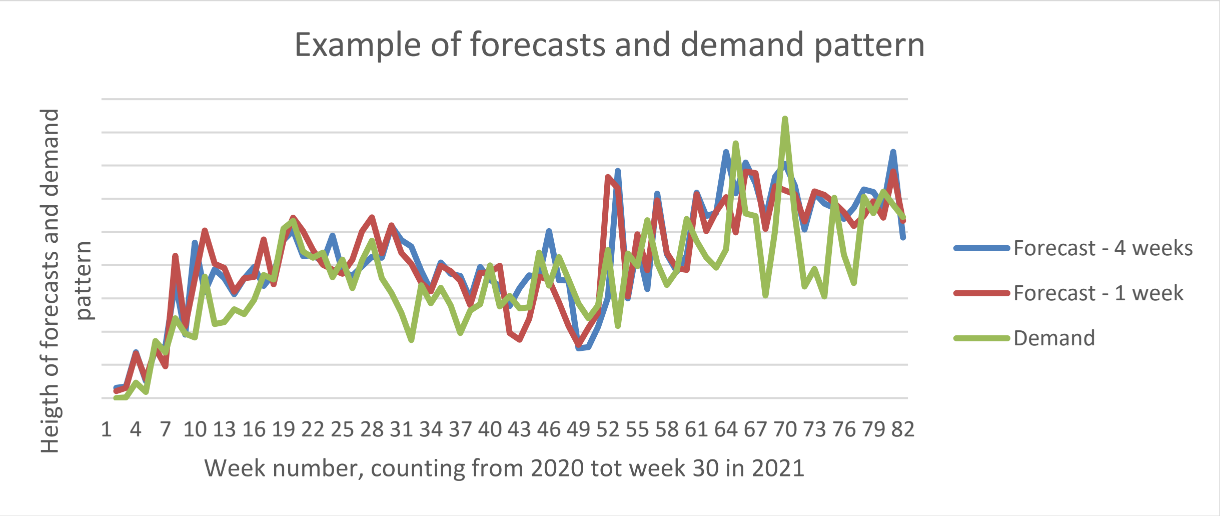

Figure 6: Example of forecast and demand pattern (Y-axis is confidential)

The random forecast is not picked from the demand pattern that is shown in figure 6. In the first few weeks, the demand and forecasts are very low because this is a new finished product. We do not want to use those weeks in our model because it is not representative for the forecast and demand pattern of the last year. Therefore, we use the first 30 weeks of 2021 and week 31 to 51 from 2020. In addition to this, we see that the pattern in 2021 is higher than the pattern of 2020. Therefore, we reduce the data of 2021 with the yearly growth factor the make sure both patterns are in proportion.

The 'forecast - 4 weeks' states the forecast of the demand exactly 4 weeks before the demand occurs. The same holds for the 'forecast - 1 week'. There is no forecast available for a longer period. However, many C&S's have a longer forecast than those 4 weeks and we need to know how much we should order for C&S's with a long lead time. In addition to this, we know that in most cases the forecasts in the short term is overforecasting demand and the forecasts in the long term is much lower than the short term forecast. This is explained in section 2.5.3. Since we had aggregated forecasts in every month for every month, we could analyse how the forecast develops over time. It shows that the forecast is increasing when the demand is less far in the future. For the forecast in the long term, we use the '4 weekly forecast', multiplied with a percentage. Table 3 shows to which forecasts we will look if we are in week n:

Week	n=0	n = 1 to 3	n = 4 to 12	n = 13 to 16	n ≥17
Forecasts	-	'1- week'	'4 – weeks' * 1,0	'4 – weeks' * 0,928	'4 – weeks' * 0,865

Table 3: Reference point for the forecasts if week n is simulated

At last, this means that every week, week n + 12 and n + 16 are updated, including the expected production of finished products and the forecast for the affected C&S's. In this way, we can take a lower forecast in the long term into account. Since we only had the aggregated long term forecasts, it was not possible to analyse flavour. Therefore, all finished products have this trendline in the forecast.

2: Generate demand pattern of finished products

The available data for the demand for finished products only include 2020 and 2021. We want to generate different and realistic demand patterns for every replication. Therefore a method is created to simulate realistic demand patterns. As one might expect, the forecast error is normally distributed. So it is expected that small deviations between the forecast and actual demand are more likely than the bigger deviations. This is a reasonable assumption and is commonly used in the computation of safety stocks (Barrow & Kourentzes, 2016).

To derive a realistic demand pattern from the forecast, we have to find the average and standard deviation of the forecast error. In addition to this, we have statistically motivated that the forecast error is normally distributed with the computed mean and standard deviation. This analysis can be found in appendix A.

The forecast error average and standard deviation are used to generate a demand pattern that is based on the forecast.

3: Generate the production cycle

Since each finished product is produced in a specific week of a three or four weekly production cycle, we have to make sure that these cycles are included in the model. We take the cycle length and the exact production week of the cycle into account. We state a production week for a finished product with the number 1 and zero if the product is not produced in that week of the simulation. To generate the result, the following general formula is developed:

$$Production\ week\ = \begin{cases} 1, if\ ((n\ +\ (CL\ -\ PW)\ )\ mod\ CL)\ =\ 0 \\ 0, else \end{cases}$$

Where,
n = week number of the simulation
CL = Cycle length {3, 4}
PW = In which week of the cycle the production should take place

Applying the formula results in the following example:

Example: Cycle Length = 4, the production week is in week 3 of the cycle. Assume the cycle starts in week 1. See table 4 to determine if a week is a production week (1) or not (0).

Week	*1*	*2*	***3***	*4*	*5*	*6*	***7***	*8*
(n + (CL – PW))	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*
(n + (CL – PW)) mod CL	*2*	*3*	*0*	*1*	*2*	*3*	*0*	*1*
Conclusion: Production week (1) or not (0):	*0*	*0*	*1*	*0*	*0*	*0*	*1*	*0*

Table 4: Example to determine the production cycle for an ingredient with a cycle length of 4 weeks and productions happens in the third week of the cycle.

4. Determine the expected production quantity

For a three-weekly cycle, the expected production quantity of week n sums the forecast for week n + 7 to n + 9. For a four-weekly cycle, the expected production quantity loops over the forecast for week n + 6 to n + 9. This is the case since every production run should cover the demand for the upcoming 9 weeks. Figure 7 shows that the production week 3 covers the demand for weeks 4 to 12. The production in week 6 should cover the demand for weeks 7 to 12. The figure shows that the demand for weeks 7 to 12 is already covered. Therefore, we expect that the production run in week 6 should produce the demand for weeks 13 to 15. In other words, the production run of week n equals the forecast for week n + 7 to n + 9. The same analysis holds for the four-weekly cycle.
[Note that the actual planned production quantity is determined when the simulation is running the week of the production. This is the case since the actual planned production quantity could be adjusted if the product is sold a lot in the past weeks, or in case of undersales].

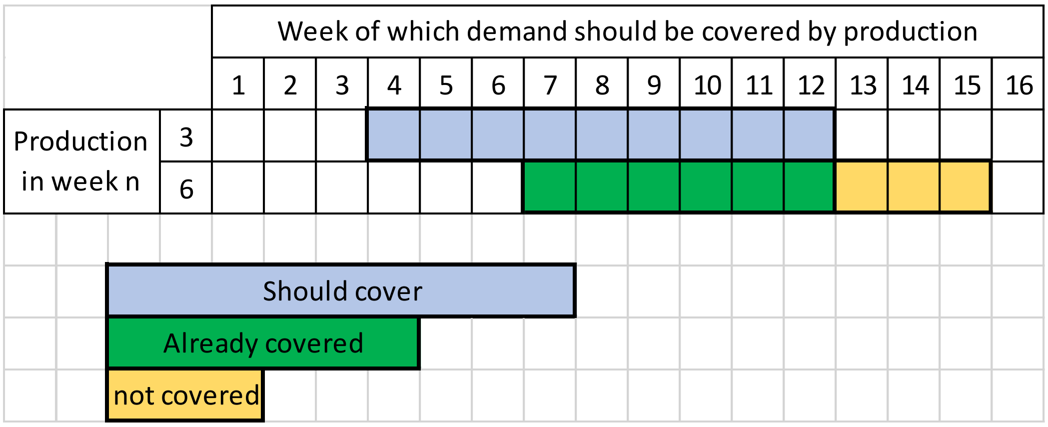

Figure 7: Example to determine the forecasted demand

4.2.2 Simulation of the week

If the steps of 4.2.1 are executed it is time to start simulating per week. The following steps are executed.

A: Update forecasts and expected production quantity of Finished Products

Before simulating all weeks, the forecast and all expected production quantities are already determined. Table 5 shows how the forecast increases over time. Besides, it states we should update week n + 12 and week n + 16. The update is as follows:

$$F_{12} = F_{12} * \frac{1{,}000}{0{,}928}$$
$$F_{16} = F_{16} * \frac{0928}{0{,}865}$$

Where,
F_n = Four-weekly forecast of week n

The next step is to adjust the expected production quantity that is affected by the new knowledge based on the updated forecast. The expected production quantity is updated for both the finished products and the C&S's. It is important to update the expected forecasts for C&S in the long term, because this affects the expected amount of a Chunk or Swirl that is needed. Updating how much we need affects if and when an order should be placed. Therefore, this process is important for C&S with a long lead time.

B: Determine the planned production quantity of Finished Products.

The planned production is the production quantity that you actually want to produce. The actual production quantity is affected by the fluctuations of the demand pattern in the past weeks. This is in line with the new strategy of Unilever and Ben & Jerry's, in which they want to work towards the situation that you increase the production of the products that are sold a lot in the last week(s). To find the planned production, the following steps are executed, <u>all for finished products:</u>

1 Determine the begin on hand inventory:
$Begin\ OH\ inventory_n = Ending\ OH\ inventory_{n-1}$
2 Check if finished products arrive from production:

$Received\ production_n\ =\ Fulfilled\ production_{n-1}$

3 Determine the demand fulfilled from stock:

$Fullfilled\ demand_n\ =\ min[demand, max\{0, Begin\ OH\ inventory\ _n + \ Received\ production\ _n -$
$Backorders_{n-1}\}]$

4 Determine backorders:

$Backorders_n = max\{0, Backorders_{n-1} + Demand_n -\ Begin\ OH\ inventory_n\ -\ Received\ production_n\}$

5 Determine Ending on hand inventory:

$Ending\ OH\ inventory_n = max\{0, Begin\ OH\ inventory_n + Received\ production_n\ -\ Demand_n\ -$
$Backorders_{n-1}$

6 Determine inventory position:

$Inventory\ position_n = Ending\ OH\ inventory_n\ -\ Backorders_n$

7 Determine the required amount to reach the 9 weeks coverage ratio:

$Required\ amount_n\ =\ \sum_{n=1}^{3}('Forecast\ -\ 1\ week'_n) + \sum_{n=4}^{9}('Forecast\ -\ 4\ weeks'_n)$

8 Determine the planned production:

$Planned\ production_n = Required\ amount_n\ -\ Inventory\ position_n$

C: Execute the begin of the inventory policies of chunks and swirls

The forecast for C&S's is determined before all weeks are simulated and every week a forecast and expected production quantity is updated, the affected forecast for C&S's is updated. The actual demand for C&S's is translated from the planned production via the BOM to the demand for C&S's. The following steps are all executed to run the inventory policy of chunks and swirls:

1 Determine the begin on hand inventory:

$Begin\ OH\ inventory_n = \ Ending\ OH\ inventory_{n-1}$

2 Determine demand fulfilled from stock

$Fullfilled\ demand_n\ =\ min[demand, max\{0, Begin\ OH\ inventory\ _n + \ Received\ order\ _n\}]$

3 Determine backorders

$Backorders_n = max\{0,\ Demand_n -\ Begin\ OH\ inventory_n\ -\ Received\ order_n\}$

Formula 2 and 3 are different compared to the formulas of the finished products since we do take the backorders from a former period into consideration. This is the case because the next step is to make sure that the backorders affect the amount of the planned production that is fulfilled. The amount that is not fulfilled results in the backorders on the level of finished products. Those backorders are added to the next production run. This means that backorders on the level of C&S's are reflected in the actual production quantity, and thus received production.

If we would take the backorders of former periods into account on the level of C&S's, it would not reflect reality. Let's take a look at an example. Assume we are not simulating a production week and we have 10 backorders from the former period. In this week, an order of 15 arrives. Formula 3, which is used for finished products, would create that the backorders are resolved. This is not possible because we are not in a production week and therefore, backorders can not disappear. Therefore, we have chosen to take care of the backorders on the level of finished products.

D: Create feedback loop from material availability to fulfilled production quantity

The backorders of chunks and swirl lead to backorders in the production. In other words, a part of the production is postponed to the next cycle. The backorders of C&S's are assigned to the weighted average demand of the products. This becomes clear with an example. Assume that there are 100 backorders of a chunk and the chunk should be used in the production for 40 finished products of A and 60 finished products of B. The 100 backorders are for 40% assigned to product A and for 60% assigned to product B. This is a loop that keeps continuing until all backorders are processed in the fulfilled production and the back ordered production. The process is shown in figure 8 on the next page. At last, all the postponed production should be added to the expected production of the next cycle.

E: Adapt the amount of used materials and the process of placing orders

We first have to repeat the three formulas of step B, but not with the planned production but with the fulfilled production. In this way, we make sure that if one chunk was only for 80% available, the other C&S's are also used for 80%. This shows that if one chunk or swirl is lacking, it could increase the inventory level of the other C&S's that are used in the same product. After the three formulas, the following steps are executed <u>for each chunk and swirl:</u>

1 Determine ending On Hand inventory:
$Ending\ OH\ inventory_n = max\{0, Begin\ OH\ inventory_n + received\ replenishment\ order_n\}$

2 Determine inventory position:
$Inventory\ position_n = Ending\ OH\ inventory_n + Pipeline_n$

3 Determine forecast during the lead time and safety lead time:
$Forecast_n = \sum_{n=1}^{N} 'Forecast\ for\ chunk\ or\ swirl'_n$
N = Nominal lead time + Safety lead time

4 Determine if an order should be placed, in what quantity and when it arrives:
The order amount is dependent on the restrictions on the order quantity. Figure 9 shows the process to determine the amount to order. The reorder point equals the 'forecast during the nominal lead time and safety lead time' + Safety Stock (SS).
For the actual lead time: First, a random number is generated and compared with the percentage of orders that is on time in full (OTIF). If the order is on time, we use the normal lead time and otherwise we get another random number to determine how much later the order arrives. The longest possible lead time is 20% longer than the normal lead time of the chunk or swirl.

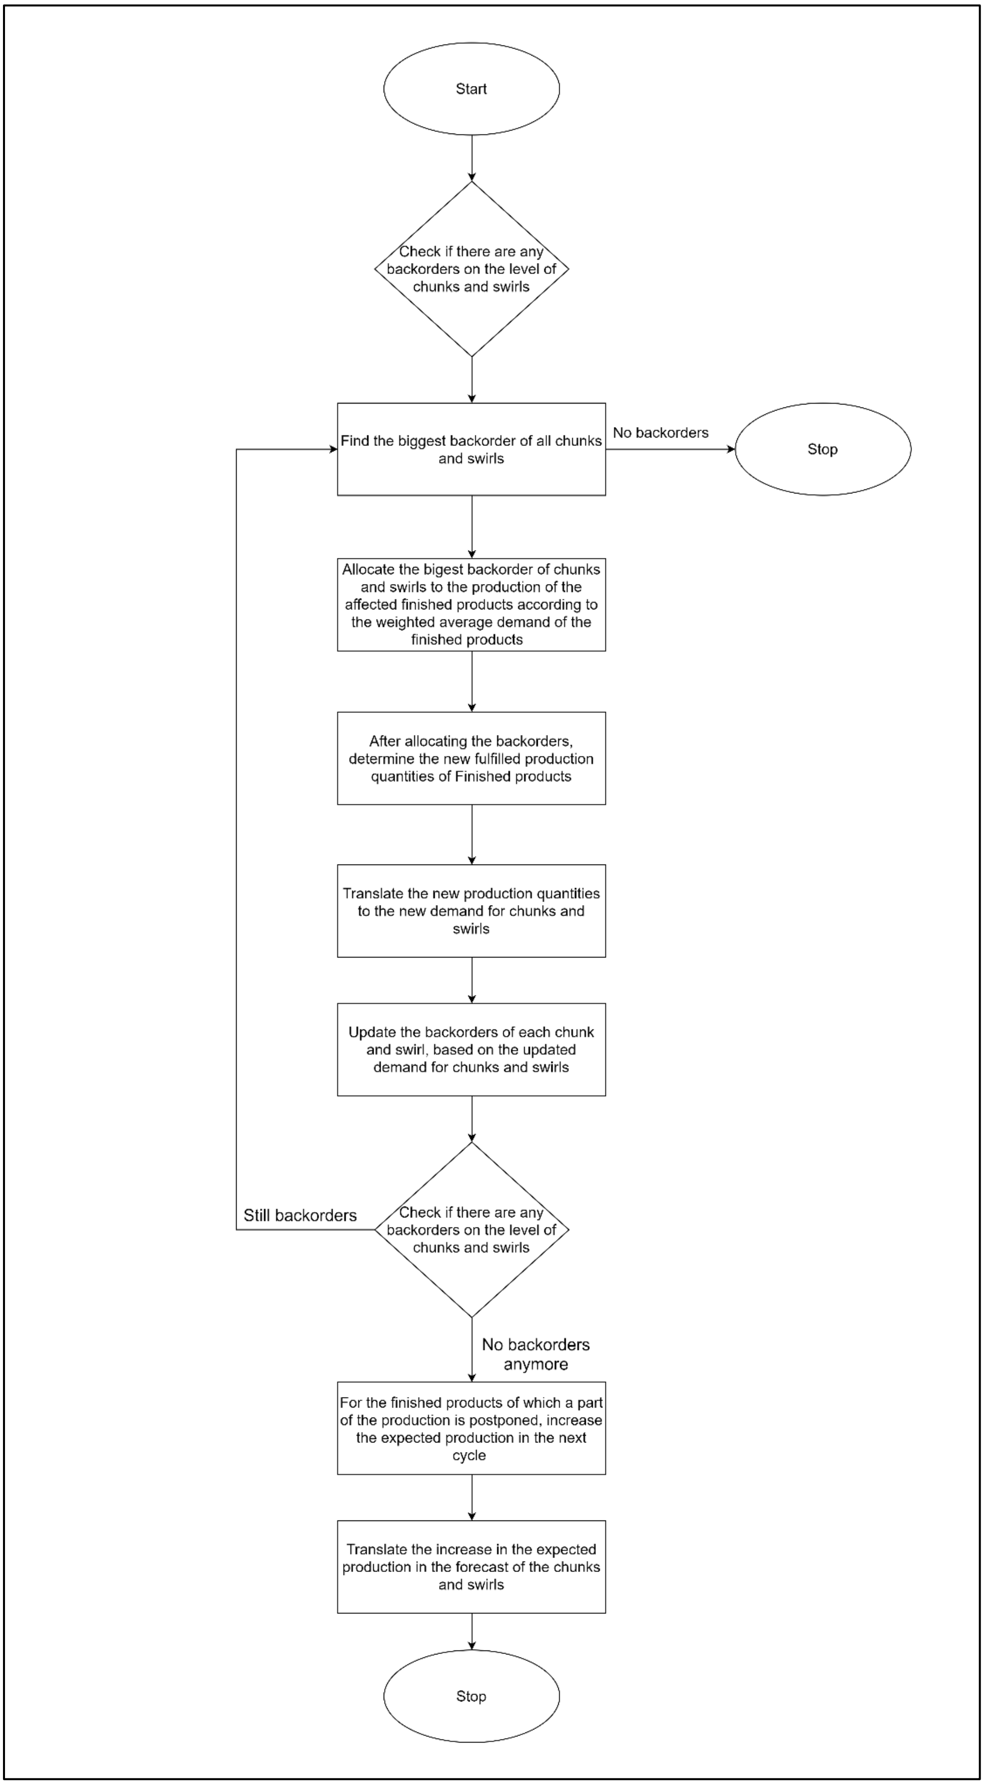

Figure 8: Process flow in case of backorders in the level of chunks and swirls

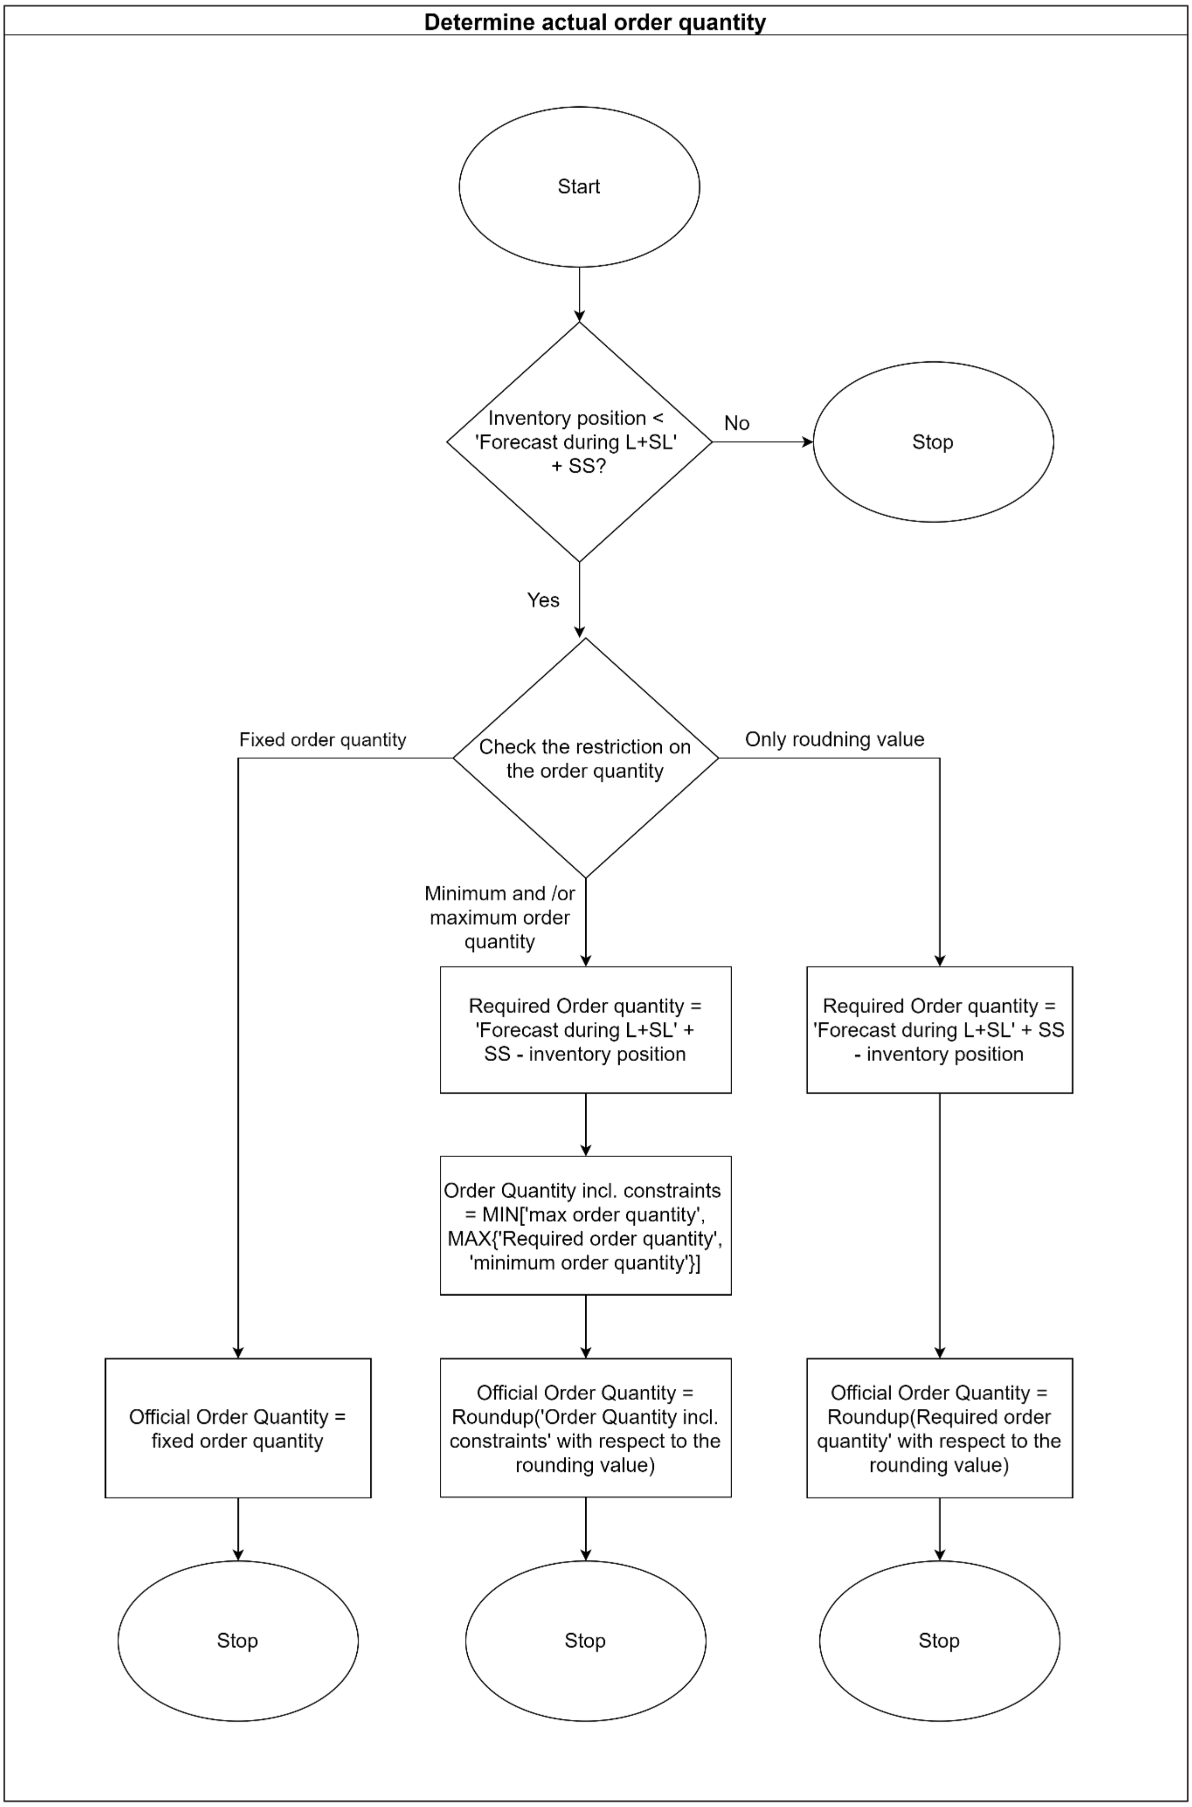

Figure 9A: Process to determine the actual order quantity

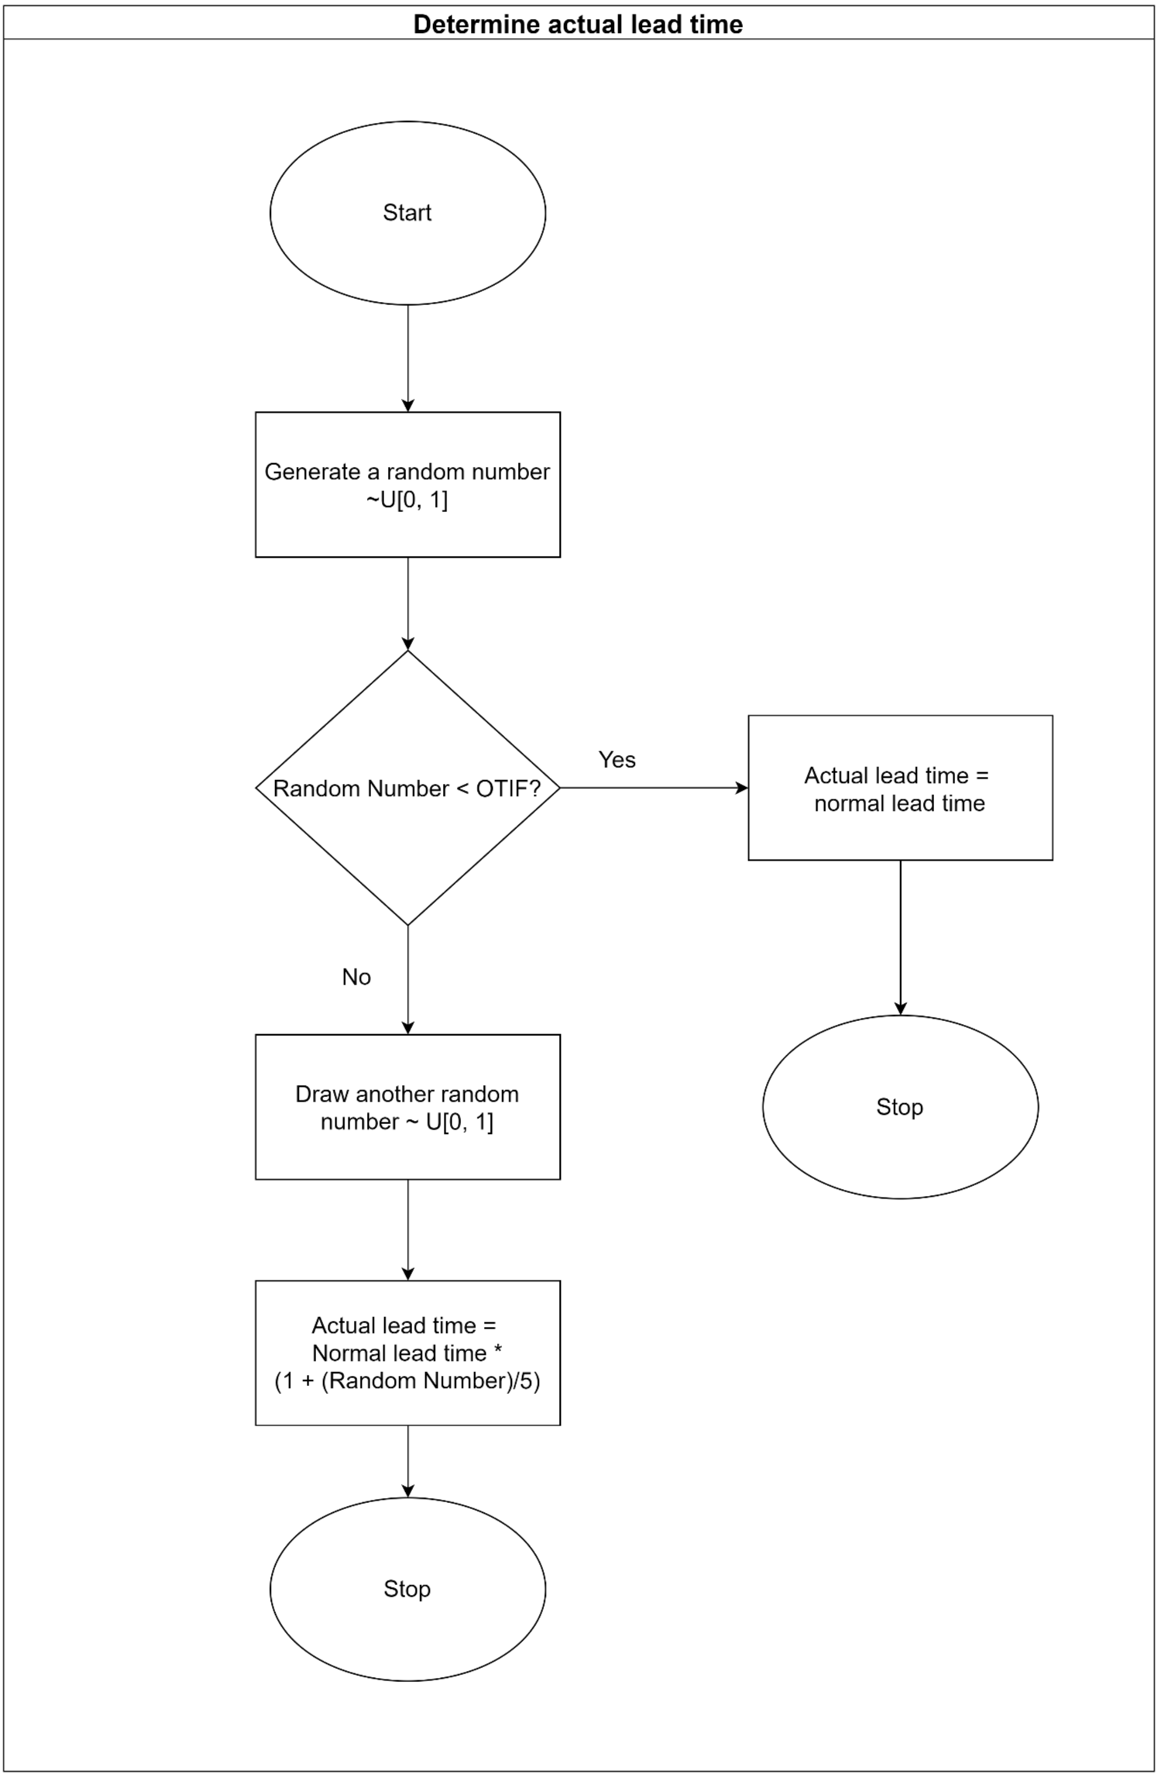

Figure 9B: Process to determine the actual lead time

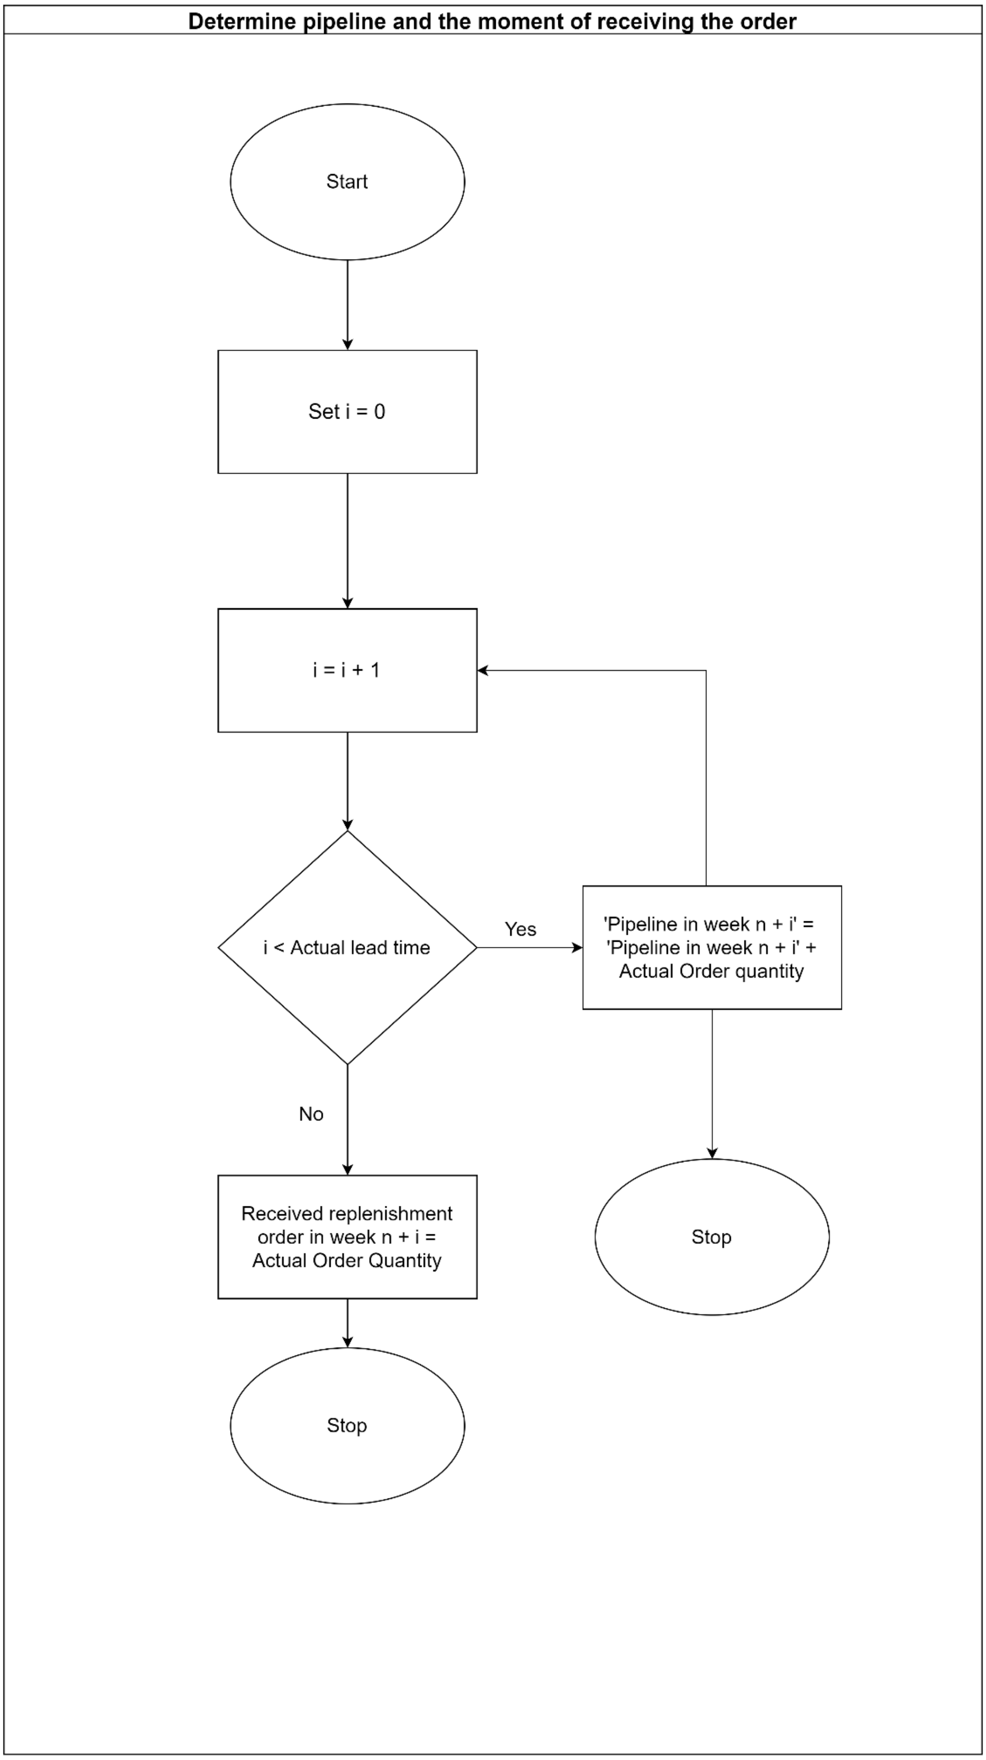

Figure 9C: Process to fill the pipeline

4.3 Simulation settings

Section 3.4 explained that a stochastic simulation model provides estimates of the 'true' outcome. Subsection 3.6 explained that measuring the performance of a simulation can have statistically more value if the right values for the warm-up period, run length and the number of replications are implemented. This section provides the conclusion of the analysis. The analysis can be found in appendix A. For the analysis, we have applied the approach by Welch (1983), which is described in the book by Law (2014).

Warm-up period:	0,5 year
Total run Length	5 years
Total number of replications:	3 replications

4.4 Random number generator

Linear Congruential Generators (LCG) are used to generate the random numbers in this simulation study. Appendix C describes the process of creating a good set of random numbers. Appendix D explains how synchronicity is built into the model.

4.5 Model verification and validation

Verification is concerned with whether the 'assumptions document' has been correctly translated into a computed 'program'. We have done this, by debugging all codes in every phase of the development phase. Besides the simulation runs have been made graphically to see and understand what is happening in a simulation run.

The validation determines whether a simulation model is an accurate representation of the system. There are different ways to validate the model. The simulation inputs can be compared with the real input parameters. The output parameter of the simulation can be compared with the actual performance of the system.

In order to create a realistic model, the simulation is developed together with multiple stakeholders from Ben & Jerry's. The part to determine the production quantities is checked with the production planners and the planning and logistics manager. The part of the simulation that manages the inventory of C&S's is checked with the material planners.

In addition to this, we have tried to make the input parameters as close to reality as possible. We could not only assume the forecast error to be normally distributed, but we have motivated it for almost every finished product. The only products that did not have statistical values were the finished products in the package of 4,5 liters. This package is sold to the scoop shops, but their opening times and demand patterns have been very unpredictable due to covid. Other input parameters are retrieved from systems, such as SAP and other internal documents.

Furthermore, we can make an in-depth analysis of a simulation run. We have manually compared the production quantities in the model with the production quantities in reality by opening the actual planning sheets. We have seen that the production quantities are all approximately on the same level. It is very time-consuming to make an overview of all production quantities to compare it with the simulation because we have to open each planning manually. The same can be done with the inventory values of C&S's and we see that the inventory levels in the mode have the same range as the inventory level in reality. At last, we compare output parameters. Chapter 5 shows that simulating with the current settings for SS and SLT, also called our ‘zero measurement’, results in a measured fill rate of the portfolio of 95,77%. Section 1.2 stated that the actual fill rate of Ben & Jerry’s equals 93,2%. Although the service level of the simulation and the actual service level of Ben & Jerry’s are not the same, we still consider the output as reliable because not all losses that affect the service level are included in the simulation model. For example, the production is not always as efficient as it should be and there could be all kinds of breakdowns and failures that affect the production. Therefore, together with the employees of Ben & Jerry’s, the service level of the zero measurements is considered to be in line with the service level in reality.

At last, we have compared the total demand of each chunk and swirl as a percentage of the total demand in the simulation, with the relative volume percentages in reality. The results are shown in figure 10 The 45-degree line indicates a perfect fit between the relative volume of a chunk in the simulation and reality. Most C&S’s are close to this line. However, we see two chunks that are far below the 45-degree line. Those two chunks are the cookie dough and the non-diary brownie. Both chunks are produced in America and are not only used in Hellendoorn, but also in a small Ben & Jerry plant in the UK. Since the transport from America to Europe is the same for both fabrics, we take that demand into account. The relative volume data of those chunks for the UK are not included in the dataset, and therefore we observe that deviation.

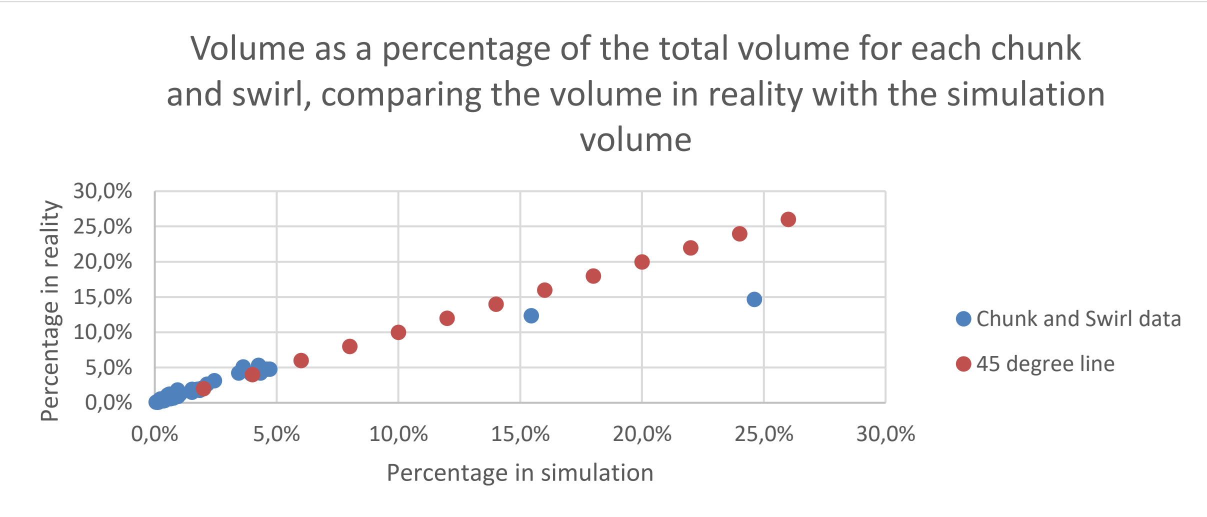

Figure 10: Compare the relative volume of each chunk and swirl in the simulation with reality

Therefore, together with the stakeholders of Ben & Jerry’s, we believe that the model reflects reality in all aspects, from the production quantities to inventory values and the service level.

4.6 Simulation optimization

The goal of the simulation optimization algorithm is to find a(n) (near) optimal solution as fast as possible. This means we want to reach a specific target service level of the whole portfolio of Ben & Jerry’s while minimizing the average investment costs in inventory. The optimization algorithm optimizes the settings for SS and SLT.

Since we have 50 C&S and for each C&S we can set a different SLT and SS, we have 100 decision variables. For each SLT and SS, we can apply different settings and therefore the solution space becomes too large to execute all possibilities. Section 2.1.3 states we are able to split the portfolio into different parts. It would be easier to optimize the less complex parts of the portfolio, but there is one set that contains 31 of the 40 C&S. Optimising that part of the portfolio is the big challenge. Adding the remaining C&S and finished products to this big part, will not affect the way of optimizing the portfolio because the big part is already the bottleneck of the simulation algorithm. In addition to this, by optimizing the whole portfolio at once, we can focus on creating one suitable algorithm that can handle the enormous complexity.

Due to the long run time of a single simulation run, we are not able to apply existing algorithms, such as Simulated Annealing. Therefore, we develop an algorithm that suits this simulation study. Section 3.4 explains the theory about Responsive Surface Methodology (RSM) and this consists of two parts: (i) Predicting the model responses and (ii) finding a combination of input factors that optimize a certain response. We are inspired by this methodology and we also apply 2 different phases in our model. The first part should provide information about the behaviour of the simulation and the second phase uses the knowledge to reach the target fill rate for the whole portfolio (99%). Our algorithm in the second phase is best described by a local search algorithm.

Before we explain the algorithm in detail, we explain some key characteristics of the simulation model. Based on these characteristics, we build the first part of our optimization algorithm. The first part of the optimization algorithm will teach us what optimal SLT and SS settings are for every single chunk or swirl if aim for different fill rates. In the second part, we will set a target fill rate for each C&S and we use the information that we have found in the first part of the simulation. The goal in the second phase is to reach the target fill rate for the whole portfolio.

4.6.1 Part one

Characteristic 1: Relationship between fill rate of a C&S and the required inventory to reach the fill rate

Figure 10 shows the relationship between the fill rate of a single chunk and the average inventory value. Due to confidentiality, we have not specified the values on the horizontal axis. The blue curve describes how the inventory increases to be able to reach a higher fill rate for a chunk or swirl. Let's define function f(x), which describes the relationship between the amount of inventory (the average inventory value) and the fill rate of a single chunk or swirl. Function f is increasing, while the rate of increase itself is decreasing. Therefore f(x) is called concave down. This means if we want to increase the fill rate step by step, the increase in the inventory is increasing. This is the case because if we want to increase the fill rate, fewer risks are allowed to be taken and therefore the inventory, and thus the average inventory value, should increase.

The blue Curve is created by not implementing any SLT and by adding different amounts of SS. The same curves can be created if we implement a SLT of 1 week and add different amounts of SS, see the red curve in figure 11. Implementing a SLT of 2 weeks and adding different amounts of SS results in the green curve. Why the curves behave in this manner, will be explained in chapter 5 when we discuss the results. In this chapter, we accept that curves with different SLT's behave differently. Looking at the table, we can say that if we aim for a fill rate of 98%, the green curve is cheaper than the red and blue curves. Therefore, a SLT of 2 weeks would be dominant over a SLT of 0 and 1 week.

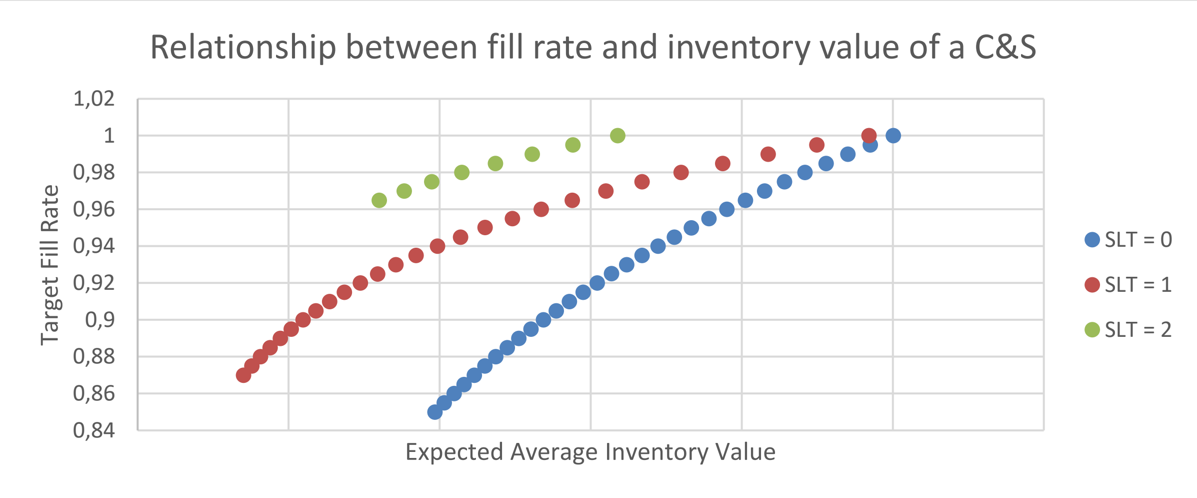

Figure 11 Relationship between the fill rate of a C&S and its (safety) inventory

Characteristic 2: Relationship between the fill rate of a C&S and the required SS to reach the fill rate

The same analysis as characteristic 1 can be done if we compare the fill rate with the necessary SS that should be added to reach a specific fill rate. This means we change the horizontal axis from the average inventory value to the amount of SS that is added, measured in kg. These curves are also concave down because if we want a higher fill rate for a C&S, more safety inventory is needed.

Where characteristic 1 explores which kind of SLT would be optimal or dominant for a specific fill rate, this characteristic states we have the same concave down effect if we compare the fill rate with the SS. If we know how to estimate the concave down curves, we could estimate the amount of SS that should be added to reach the target fill rate for this chunk or swirl.

Characteristic 3: The concave down relationship curves can be approximated by using a logarithmic function

The relationship between the inventory and the fill rate of a chunk or swirl could be described by applying logistic regression. However, we are only interested in the fill rates above 95% and we are able to recreate the curves by applying logarithmic regression. A logarithmic curve is described by $y = A + B * \text{Ln}(x)$. If we want to estimate the curves for each single C&S between a specific fill rate and what it would cost to reach the fill rate, we apply the following formula:

$$Fill\ rate = A + B * LN(Average\ inventory\ value) \tag{i}$$

Appendix B explains how parameters A and B are retrieved. Rewriting the function above, we can estimate the average inventory value for a target fill rate as follows:

$$Estimated\ inventory\ value = e^{\frac{Fill\ rate - A}{B}} \tag{ii}$$

To estimate the logarithmic curves, we need at least 3 different points and those points are generated by executing at least three experiments per SLT. In this way, we can create the logarithmic curves for each SLT. Formula (ii) enables us to determine which curve, and thus which SLT, is optimal.

Applying the same theory, not with the average inventory value but with the amount of SS, we can estimate how much SS should be added in addition to the optimal SLT to make sure a specific fill rate is reached. This results in the following two formula's:

$$Fill\ rate = A + B * LN(SS) \tag{iii}$$

$$SS = e^{\frac{Fill\ rate - A}{B}} \tag{iv}$$

Formula (iii) results in the logarithmic lines and with formula (iv) we can estimate how much SS we should add to reach the target fill rate.

General idea

The three characteristics state that the relationships between a fill rate and the average inventory value or the amount of SS are concave down. We can estimate those curves by applying logarithmic regression. The logarithmic curve between the fill rate and the average inventory value provides the information on which SLT would be dominant for a specific fill rate for a chunk or swirl. The logarithmic curve for the fill rate of a chunk or swirl and the SS helps us to compute how much SS should be added to reach the target fill rate for the chunk.

An overview of how all experiments are defined is shown in table 5. Note that the number of experiments of phase 1 is a fixed number of experiments. Since literature did not provide a method to set a good starting settings for the SLT and SS, we start with 5 experiments in which each experiment uses a different SLT. In those 5 experiments, we measure the average and the standard deviation of the backorder. Thereafter, we define for each SLT 4 experiments in which we increase the SS. In the first experiment, we add the average measured backorders and in the three experiments thereafter the add the measured standard deviation of the backorders to the SS. In theory, we only need three points to estimate the logarithmic curves, but we create more datapoints because our datapoints are created by the simulation, and thus randomness is involved. To get reliable logarithmic curves, we have created this design in phase 1 of the algorithm.

Experiment	Setting		Purpose
	SLT	SS	
1	0	0	Measure average (π_{SLT}) and standard deviation of the backorders (σ_{SLT})
2	1	0	
3	2	0	
4	3	0	
5	4	0	
6	0	π_0	Create logarithmic curves with SLT = 0
7	0	$\pi_0 + 1\sigma_0$	
8	0	$\pi_0 + 2\sigma_0$	
9	0	$\pi_0 + 3\sigma_0$	
10	1	π_1	Create logarithmic curves with SLT = 1
11	1	$\pi_1 + 1\sigma_1$	
12	1	$\pi_1 + 2\sigma_1$	
13	1	$\pi_1 + 3\sigma_1$	
14	2	π_2	Create logarithmic curves with SLT = 2
15	2	$\pi_2 + 1\sigma_2$	
16	2	$\pi_2 + 2\sigma_2$	
17	2	$\pi_2 + 3\sigma_2$	
18	3	π_3	Create logarithmic curves with SLT = 3
19	3	$\pi_3 + 1\sigma_3$	
20	3	$\pi_3 + 2\sigma_3$	
21	3	$\pi_3 + 3\sigma_3$	
22	4	π_4	Create logarithmic curves with SLT = 4
23	4	$\pi_4 + 1\sigma_4$	
24	4	$\pi_4 + 2\sigma_4$	
25	4	$\pi_4 + 3\sigma_4$	
26 +	Phase 2 – variable number of experiments		

Table 5: Overview of different experiments in part 1 of the optimization algorithm

4.6.2 Phase 2

In phase 2, we need to use the insights of phase 1. In phase 1, we have created different logarithmic curves that help us define the optimal settings for SLT and SS. The question is, which target fill rate we should aim for. The chart flow of phase 2 is depicted in figure 12. In general, phase 2 consist of the following steps: First, we define a simulation target. In the first experiment, we set the simulation target equal to the target fill rate of the whole portfolio (99%). With this, we compute for each C&S the individual target fill rate. Based on the fill rate, we check with formula (ii) of phase 1 what the optimal SLT is for the target fill rate of the chunk or swirl. With formula (iv) we compute how much SS should be added to be able to reach the target fill rate for the chunk or swirl. With the found settings, an experiment is executed and we measure the realised fill rate of the simulation. If the relative error between the realised fill rate in the simulation and the desired fill rate of the portfolio (99%) is smaller than 0,1%, we accept the outcome. Otherwise, we adjust the simulation target and the individual fill rates to run another experiment. We first explain how we set individual fill rates for each C&S based upon an overall simulation target. Thereafter, we explain how we adjust the simulation target after each experiment to get closer to the reach the target fill rate for the portfolio (99%).

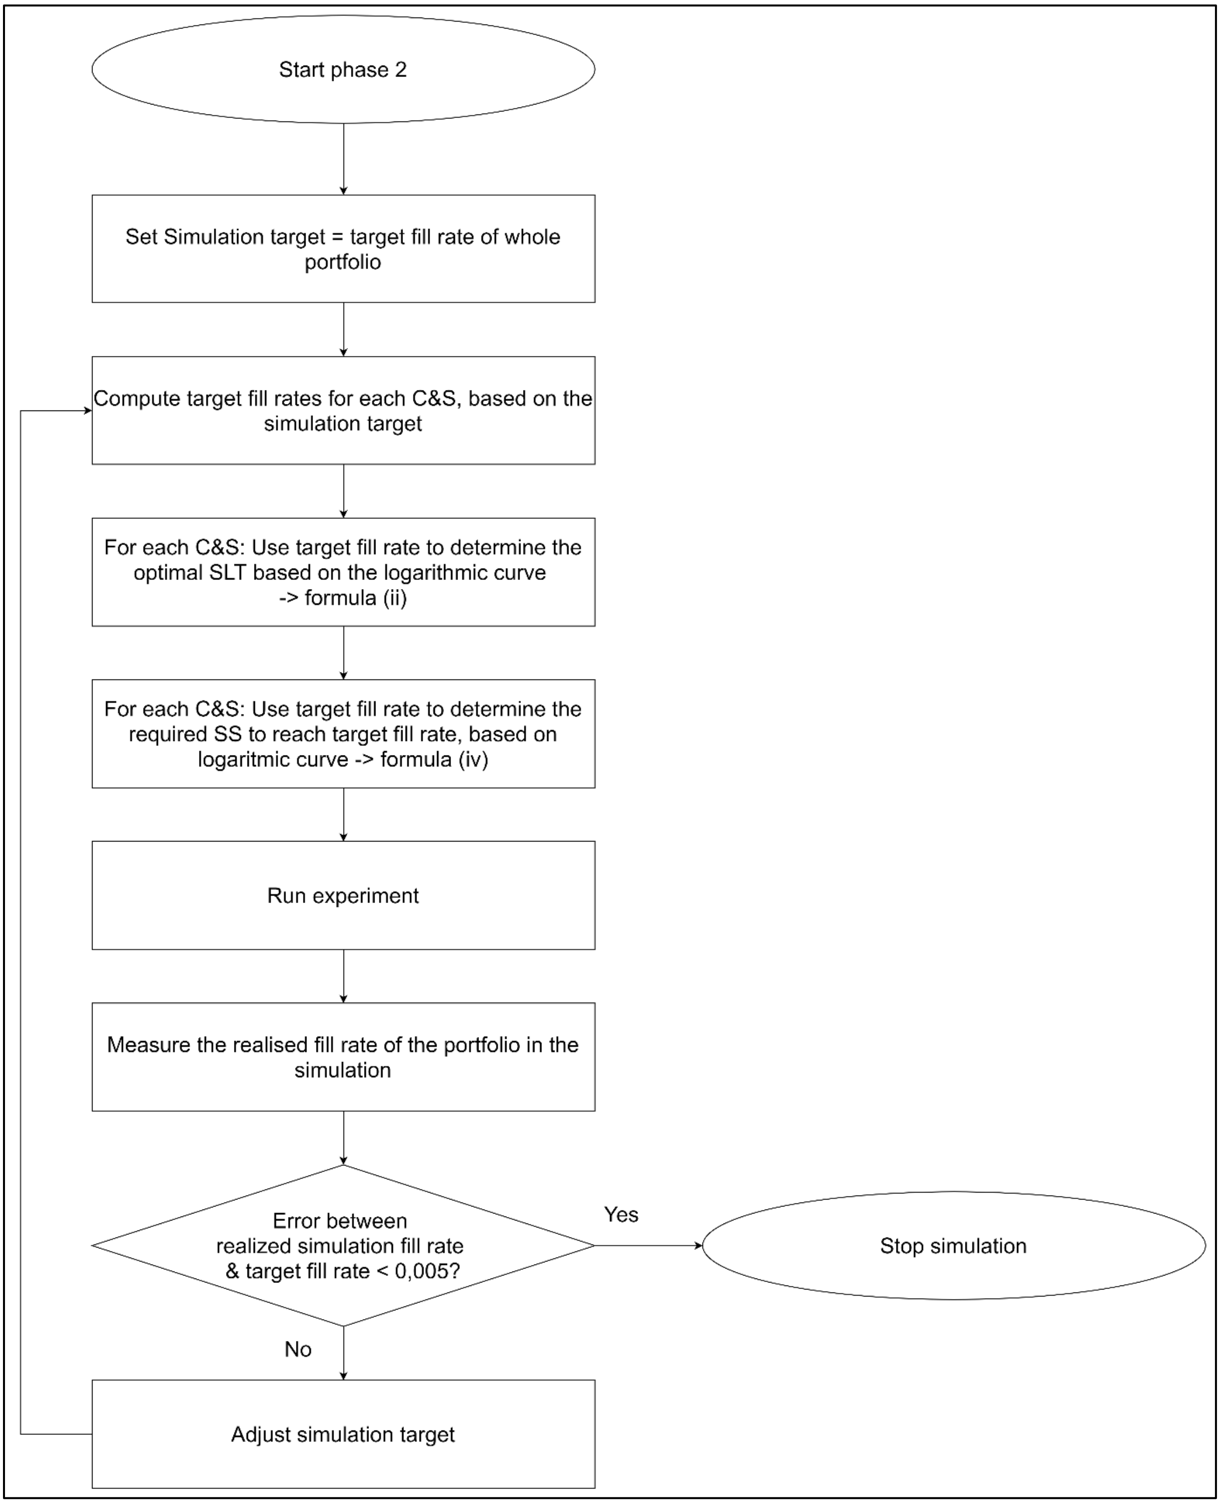

Figure 12: Chart flow of phase 2 of the simulation optimization algorithm.

Settings individual fill rates for each C&S based on the simulation target

Settings the target fill rates is done in a quite general way and we do not make any distinction in the importance, costs and volume of chunks, swirls or finished products. Therefore, we aim to have all finished products around the same fill rate. Not every finished product needs to have a fill rate of 99%, but the weighted average should be 99%. To translate the target fill rate of the portfolio to a target fill rate for each chunk and swirl, we do ignore correlations between C&S's for simplicity, but we do take into account the number of C&S's that are needed to produce each finished product.

We start with defining an overall simulation target. The individual fill rates are retrieved from the simulation target. To do this, we use a general method that only takes the complexity into account. As an example, we use a simulation target of 99% (=0,99).

Option 1: We have 1 chunk in 1 finished product.

This is the least complex option. We set the target fill rate of the chunk equal to the simulation target.

$$Target\ fill\ rate\ of\ chunk\ =\ simulation\ target = 0{,}99$$

Option 2: We have multiple C&S in one finished product

Assume the amount of C&S in the finished product equals X, the target fill rate for each of the X C&S equals:

$$Target\ fill\ rate\ for\ each\ chunk\ =\ Simulation\ target^{\frac{1}{X}} = 0{,}99^{\frac{1}{X}}$$

Option 3: We have multiple C&S in multiple finished products

Assume we have X C&S in Y different products. We do the same as option 2, but we search for each C&S the maximum amount of other C&S that are included in the finished product. For example, chunk A is used in finished product Y and finished product Z. Product Y and Z contains in total 4 and 6 different C&S's respectively. We compute the target fill rate of chunk A as follows:

$$Target\ fill\ rate\ chunk\ A\ =\ \max\{0{,}99^{\frac{1}{4}}, 0{,}99^{\frac{1}{6}}\} = 0{,}9983$$

In this way, we ignore correlation, but we take into account that the amount of different chunks that are needed in each finished product. If we would target each individual fill rate on 0,99, a finished product that used 6 different C&S would get a fill rate that is much lower than 0,99.

Adjusting the simulation target

In the first experiment of phase 2, we set the simulation target equal to the target fill rate of the portfolio. However, we expect that after running the first experiment, the measured fill rate of the portfolio in the simulation is expected to be bigger than the target fill rate due to two reasons. First of all, we ignore correlations between different C&S. Secondly, we take the highest target fill rate if multiple chunks are used in multiple products. However, after each experiment, we adjust the simulation target, and thus the target fill rates for each chunk and swirl.

The adjustment of the simulation target is done in the following way: In the first simulation, we defined the simulation target to be 0,990, which equals our fill rate target for the portfolio. Running the simulation resulted for example in a realized fill rate of the portfolio of 0,993. However, the desired outcome of the simulation equals 0,990. The simulation target for the next simulation is adjusted by cross-multiplying the values in the table below. This means, the new simulation target $Z = \frac{0{,}990 * 0{,}990}{0{,}993}$ = 0,987. With the new simulation target, we redefine the individual target fill rates for each C&S and we again measure the realized fill rate of the portfolio in the simulation.

Adjustment 1	Simulation target	Realized fill rate of the portfolio in simulation
Last experiment	0,990	0,993
New experiment	Z = 0,987	Desired = 0,990

Table 6: First adjustment in simulation target

In the second experiment, the simulation target was 0,987 and after running the simulation, the realized fill rate of the portfolio equals 0,9886. The desired fill rate of the portfolio remains 0,990 and by doing a cross-multiplying of the table, our new target equals $Z = \frac{0{,}987 * 0{,}990}{0{,}9886} =$ 0,988. Based on the new simulation target, we recompute the new individual target fill rates for each C&S.

Adjustment 2	Simulation target	Realized fill rate of the portfolio in simulation
Last experiment	0,987	0,9886
New experiment	Z = 0,9888	Desired = 0,99

Table 7: Second adjustment in simulation target

This local procedure is continued until the relative error between the simulation fill rate and desired fill rate is smaller than 0,1%. Although phase 2 is a local search algorithm, together with the insights of phase 1, big improvements are realized. At last, in the simulation model, multiple target fill rates for the portfolio could be defined that are being executed one after the other.

4.7 Conclusion

The simulation model is developed on a tactical level. The goal is to reach a specific target service level for the whole portfolio of Ben & Jerry's while minimizing the average investment in inventory. The simulation takes all important aspects into account, like production cycles, uncertainty in the demand patterns, uncertainty in the lead time, overforecasting on the short term while the long term forecasts are much lower, restriction on the order quantity and the correlations between C&S.

For the optimization part, we generate different points for the fill rate, SS and average inventory value. With this, we use logarithmic regression to determine the relationship between the average inventory value of a C&S with respect to its target fill rate. The regression curves for each SLT determine which SLT is dominant on a specific fill rate interval. Thereafter, another logarithmic regression curve determines which amount of SS should be added to the dominant SLT to make sure the target fill rate is reached. The Simulation stops when the target service level is reached and the inventory value is minimized.

5 Results and analysis

Section 5.1 provides the results of a simulation run with the current settings of the SS and SLT and how these settings should be optimised under the condition of the current overforecasting of demand. In section 5.2 a sensitivity analysis is executed that explore how results change if the current way overforecasting becomes underforecasting or if there is no structural forecast error. Section 5.3 provides the conclusion of the results.

5.1 Current situation - overforecasting

In this section, we will first show the results on a portfolio level for both the zero measurement and the optimization. Thereafter, the results are compared on a portfolio level and we will zoom in to the level of C&S's and show the most notable changes. The values for the inventory value are expressed as an index, due to confidentiality.

5.1.1 Zero measurement

In the current situation, almost every finished product is structurally overforecasted. This and the current settings for SS and SLT are implemented in the zero measurement. As table 8 shows, the current way of managing the inventory results in a fill rate of the portfolio of 95,8%. We set the index of the average inventory value of the current situation to 100. All other indexes will be compared to this zero measurement.

Setting	Portfolio fill rate	Average inventory value
Overforecasting (Current situation)	95,8 %	100

Table 8: Portfolio fill rate in the current situation

To compute the 95% confidence interval, we use the formula:

$$\left[\bar{X} - z_{\alpha/2,} \frac{\sigma}{\sqrt{n}} \; ; \; \bar{X} + z_{\alpha/2,} \frac{\sigma}{\sqrt{n}}\right]$$

This results in the following confidence interval:

$$\left[95{,}8 - 1{,}96 * \frac{0{,}42}{\sqrt{3}} \; ; \; 95{,}8 + 1{,}96 \frac{0{,}42}{\sqrt{3}}\right] = [94{,}95 \; ; \; 96{,}88]$$

About the settings, Ben & Jerry's prefers to implement SLT instead of SS. The reason for this is that SLT is stock that is already assigned to one of the production runs, while SS is extra stock and not specified for one of the production runs. This preference is reflected in the current settings for SLT and SS. All C&S have SLT and there are only 4 C&S that have SS on top of the SLT.

5.1.2 Optimization

We have optimized the simulation for different Target Service levels. The results are shown in table 9. In addition to this, the results of both the optimization and the zero measurement are depicted in figure 13.

Target service level	Reached portfolio fill rate	Average inventory value
96,0	96,0	75,0
97,0	97,3	77,9
98,0	98,2	81,6
98,5	98,7	83,5
99,0	99,1	84,3
1,00	99,99	100,3

Table 9: Results of the current situation after optimization

The 95% confidence interval that corresponds with the target service level of 99% equals:

$$\left[99{,}1\ -\ 1{,}96\ *\frac{0{,}33}{\sqrt{3}}\ ;\ 99{,}1\ +\ 1{,}96\frac{0{,}33}{\sqrt{3}}\right]\ =\ [98{,}48\ ;\ 99{,}77]$$

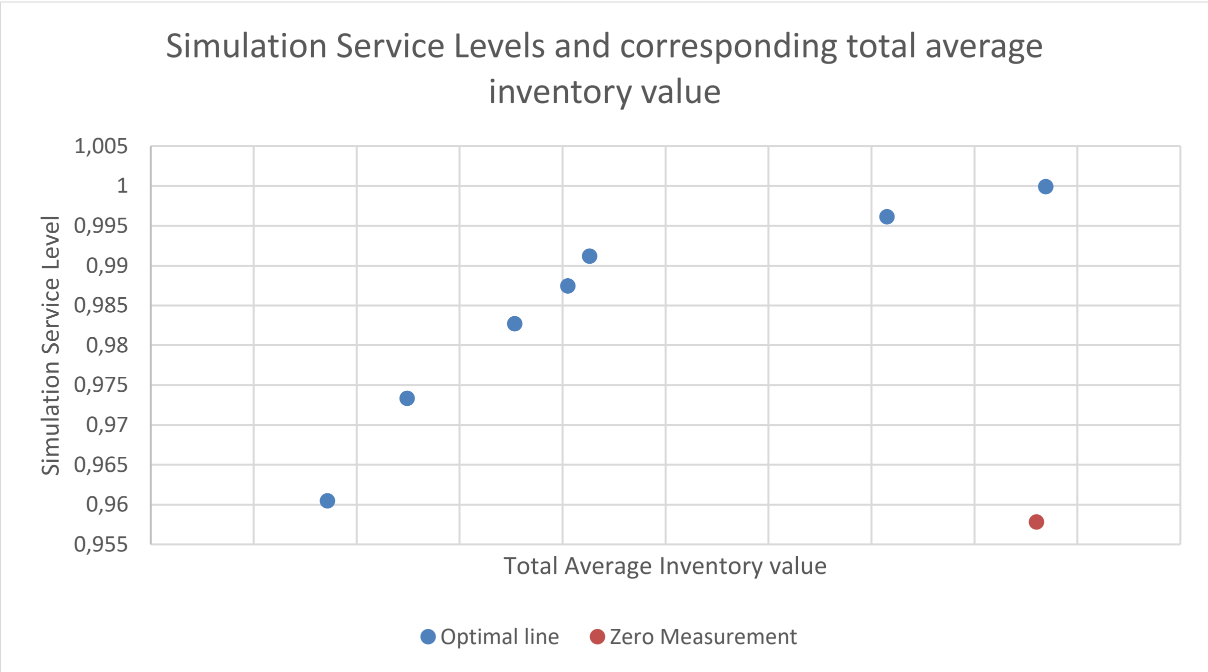

Figure 13: Results of the current situation of both the zero measurement and after optimization

5.1.3 Analysis

We analyse the results on three different levels. The first analyses are on a portfolio level, thereafter we discuss the results for two C&S to provide an example. At last, we provide an example on the level of a finished product.

Portfolio level

Figure 13 shows that the current way of managing the inventory is far from optimal. Compared to the optimal curve, the current situation does not make efficient use of its inventory and the result is a relatively high average inventory value and a low service level.

The inefficiency is the results of not having analysed and optimized the SS and SLT when the fixed production cycles were implemented. This is important, since creating fixed cycles provides a clear structure in when each C&S is needed. In addition to this, Ben & Jerry's created flexibility by implementing high SLT's. In this way, Ben & Jerry's could use parts of an replenishment order that was intended to be used in the next production cycle if there was not enough for the current production cycle. In other words, they use inventory that should be used in the next production cycle. The consequence is that the inventory increases during the whole period in between the two cycles, while the inventory is not used. The conclusion is that the inventory of C&S's was not in line with the new and fixed production cycles.

Since Ben & Jerry's in Hellendoorn has agreed with Unilever on a service level of 98,5% we have to aim for a slightly higher service level target because there are other losses that or not taken into account in this research. Therefore, we aim for a service level of 99%.

Based on these results, changing the SS's and SLT's, the service level is able to increase from 95,8% to 99,1% while the average inventory value decreases with 15,72 % from an index of 100 to 84,3.

Table 10 shows that 10 C&S's have both better performances and a decrease in the average inventory value. Table Y shows that for those 10 C&S the inventory value decreases with 5,6% and that the fill rates was on average 2,5% higher. The second row shows that 14 C&S had a cost-saving of 21,2% and as a consequence, the fill rate of the C&S was on average -0,3% lower. The third row shows that 24 C&S had an increase in the fill rate and as a consequence, the costs were 10,9% higher. The last row is a change that is not desired because both the inventory value increases and the fill rate decreases. However, the impact of both C&S is minor.

Inventory	Fill rate	Counted C&S (total = 50)	Inventory increase / decrease	Average Fill rate increase / decrease per C&S
Decrease	Increase	10	5,6% saved	+ 2,5 %
Decrease	Decrease	14	21,2% saved	- 0,3 %
Increase	Increase	24	10,9% extra	+ 4,9 %
Increase	Decrease	2	0,3% extra	- 0,2 %

Table 10: Number of positive and negative changes for the average inventory value and fill rate

Overall, we can state that the simulation has decreased a lot of SLT and SS was added. In the zero measurement, the average SLT was 2,32 and after the optimization, the SLT was on average 0,96. In the zero measurement, all C&S had SLT and in the optimization only 39 had SLT implemented. In addition to this, the average SS has increased from 901 kg in the zero measurement to 4.774 kg after optimization.

Chunks & Swirls

We will now discuss two examples to understand the changes on the level of C&S. The first example is the brownie. This chunk had the biggest improvement in terms of saving money.

The second example is the Peanut Butter Cup. This chunk had the biggest improvement in terms of the fill rate and the highest increase in the average inventory.

Example 1 – Brownie pieces

Brownie pieces are one of the most used C&S. It is used in 5 different finished products and it is used in production almost every week. With a lead time of 63 days, the lead time is relatively large. Due to the assumption that variation in the lead time could increase the lead time with 20%, equally to 2 weeks extra. The brownie pieces had a fill rate of 100% in the zero measurement. We set the average inventory value to an index of 100. The fill rate of 100% is achieved since Ben & Jerry's has implemented a SLT of 4 weeks and a SS of 22.400 kg, see table 11.

In phase 3 of the algorithm, the logarithmic lines are created that describe the relationship between the expected average inventory and a specific target fill rate. The logarithmic lines for the different SLT's are shown in figure 14. Comparing the different lines, we see that increasing the SLT from 0 to 2 weeks is beneficial. As soon as the SLT exceeds the maximum increase in the SLT, the results are getting worse. The figure suggests that implementing the SLT of 2 weeks would be dominant in the fill rate interval from 97% to 100%. In addition to this, implementing a SLT of 4 weeks and no SS would already result in a SS of 100%.

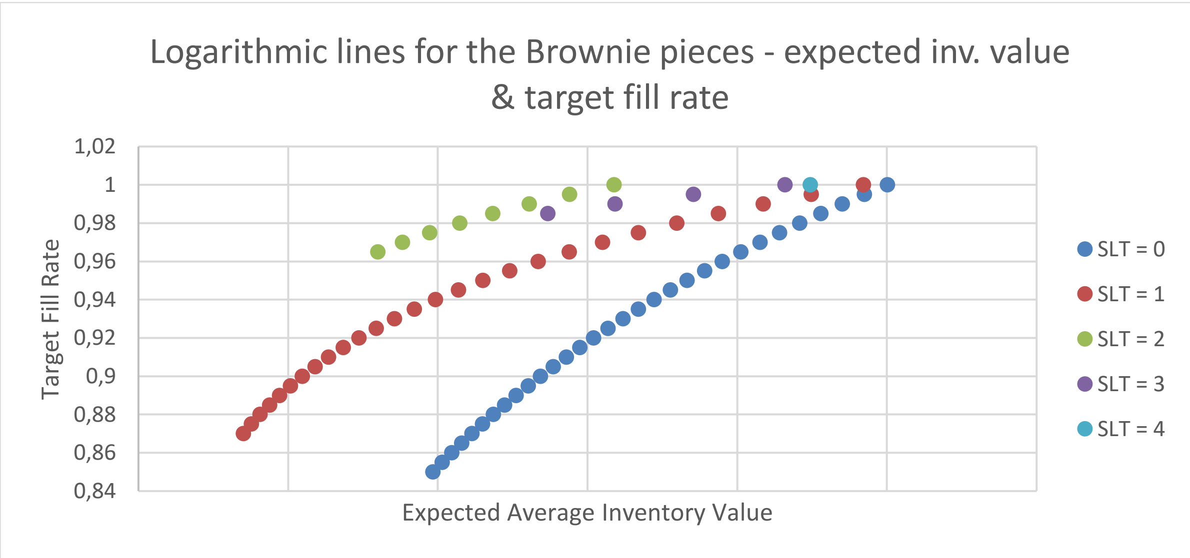

Figure 14: Logarithmic lines for the Brownie pieces, shown for different SLT's.

The optimization algorithm has decreased the SLT from 4 to 2 weeks and increased the SS from 22.400 kg to 31.303 kg. With this, the fill rate had only decreased from 100% to 99,4%. The consequence is that the average inventory value decreased with 41,3%.

Although extra SS would increase the average inventory value, the main costs savings are achieved by decreasing the SLT. If a product is produced almost every week and the production should already be on hand four weeks before production, the average inventory position increases a lot. This explains the big cost savings.

Brownie pieces	Zero measurement	Portfolio fill rate of 99%	Increase / decrease
Fill rate	100 %	99,4 %	- 0,6 %
Average inventory value	100	58,7	- 41,3%
SLT settings	4 weeks	2 weeks	- 2 weeks
SS setting	22.400 kg	31.300 kg	+ 8.900 kg

Table 11: Settings and performance of the brownie pieces in the zero measurement and after optimization

Example 2 – Peanut butter cups

Peanut butter cups are used in only 1 finished product and therefore it is used in production once every 4 weeks. The finished product of the peanut butter cups has very little structural overforecasting but the standard deviation of the forecast error is large, approximately 24% of the average weekly demand. The peanut butter cups have a lead time of 84 days and its variation in the lead time is also large. Currently, B&J implements 2 weeks of SLT and no SS, see table 12.

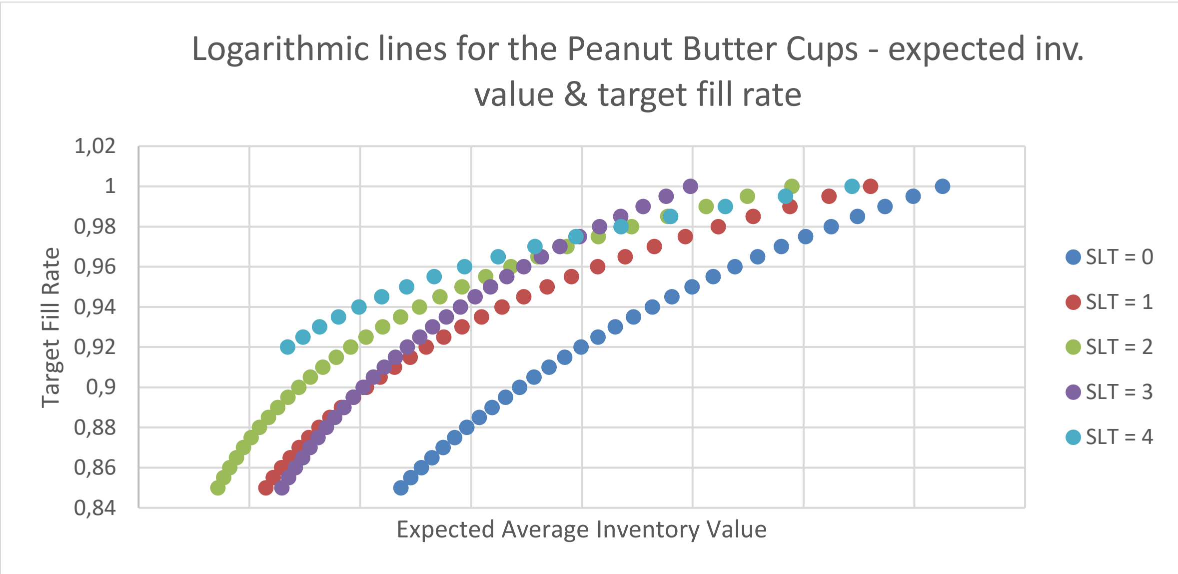

Figure 15: Logarithmic lines for the Peanut Butter Cups, shown for different SLT's.

After optimization, the simulation model has increased the SLT from 2 to 3 weeks and it has added 26.308 kg as SS. With this solution, the simulation can react to the fluctuation in the lead time, which is at maximum the following: $Roundup\left(\frac{84*1,2}{7}\right) - roundup(\frac{84}{7}) = 3$ weeks. The SS that is added to the SLT enables Ben & Jerry's to react to the big fluctuations in the demand pattern. According to the lines in figure 15, this solution is better than having 4 weeks of SLT and no SS. That solution would imply that Ben & Jerry's can use the inventory for the next production run if the Peanut Butter Cups are lacking. At last, with a SLT of less than three weeks, more SS should be added and that is less desirable compared to the currently found solution.

Peanut butter cups	Zero measurement	Portfolio fill rate of 99%	Increase / decrease
Fill rate	78,9%	99,7%	+ 20,8 %
Average inventory value	100	245.5	+ 145,5 %
SLT settings	2 weeks	3 weeks	+ 1 weeks
SS setting	0 kg	26.308 kg	+ 26.308 kg

Table 12: Settings and performance of the Peanut Butter Cups in the zero measurement and after optimization

Finished product

Analysing the result on the level of finished products, resulted in another insight. We zoom in on the finished product Phis Food. This is one of the few products that is being underforecasted. The fill rate of this finished product equals 88,2% and it uses three different C&S. Figure 13 shows the three C&S. Two C&S have a fill rate of 99,72% and 99,81% and the third had fill rate of 88,3%. The last one is the bottleneck for the production of the finished product. Since the finished product is being underforecasted, it is interesting to understand why only one chunk is the bottleneck.

The Phis Food is structurally underforecasted with 498 ZUN per week. Figure 16 also shows that the first and second C&S are also used in other finished products and those finished products are being overforecasted. The bottleneck, the third chunk, is not used in any other finished product. The analysis shows that the first and second chunk still have a high fill rate because the underforecast of the phis food is being compensated by the overforecasts of the other finished products. The last chunk does not have this, and therefore is not able to deal with the structural underforecast. In other words, the overforecast already lead to an increase in the inventory. This increase can be considered as some kind of SS for other products or future production cycles.

(1) Caramel swirl (99,72 %)

Phish Food 498 Zun

TopChocCharCookDough -11 Zun

CaramelChewChew (465 ml) -116 Zun

CaramelChewChew (100 ml) -572 Zun

(2) Marshmallow syrup (99,81 %)

Phish food: 498 Zun

Baked Alaska -966 Zun

(3) Dark comp. Pieces (88,31 %)

Phish Food 498 Zun

Figure 16: Three C&S of Phish Food and in which other finished products the C&S are used.

5.2 Sensitivity analysis

In this section, we execute a sensitivity analysis. It provides insights into the consequences when certain parameters change. We will compare different scenario's for the overcast, varying from overforecasting of demand to underforecasting of demand. We compare the different scenarios for the SLT and SS in the zero measurement and the optimized settings.

As explained in subsection 4.2.1. the demand for finished products is retrieved from the forecasts based upon a forecast error. The forecast error is normally distributed and consist of an average error and the standard deviation of the error. If we set the average error to zero and keep the standard deviation, we have the same fluctuation between demand and forecasts, but there is no structural forecast error. This equals the situation in which there is no structural over or underforecasting. The current average forecast errors, which are negative with overforecasting, are set to a positive number to create the situation with structural underforecasting.

5.2.1 Results

For all scenarios, we have executed a simulation run with the settings for SLT & SS of the zero measurement and the optimized settings. Table 13 shows the results of the current settings under the different forecast biases and table 14 for the optimized settings with different forecasts biases. Figure 17 shows the optimal lines for different target service levels and the corresponding inventory values of all three situations.

Setting	Service level	Average inventory value
Overforecasting (Current situation)	95,77	100,0
No structural over/underforecasting	94,05	86,4
Underforecasting	87,73	80,1

Table 13: Current settings with different settings for the forecasts

Setting	Service level	Average inventory value
Overforecasting (Current situation)	99,1 %	84,0
No structural over/underforecasting	96,1 %	71,0
Underforecasting	88,7 %	64,5

Table 14: Optimized settings if fill rate of portfolio equals 99,1% with different settings for the forecasts

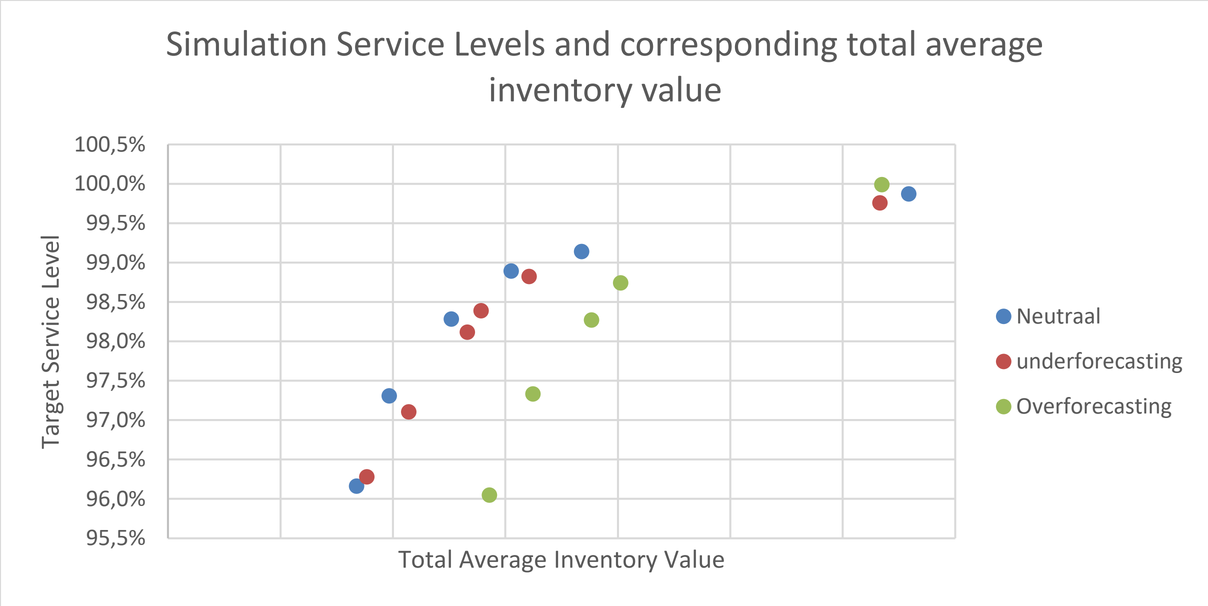

Figure 17: Target service level and the corresponding Average Inventory Value for different forecast situations

5.2.2 Analysis

Table 13, which shows the results of the zero measurements under different conditions, show that the service level is slightly higher than the situation when there is no structural overforecasting or underforecasting. Due to the overforecasting, not all C&S's are used at the moment that they were planned for. This increases the average inventory level. In addition to this, the increase in the average inventory leads to an increase in the service level. In a situation with structural underforecasting, the service level is much lower. Because Ben & Jerry's is not implementing SS, with few exceptions, it is not able to cover the uncertainty in the demand pattern. In addition to this, the inventory value will decrease because most often all materials will be used in the production. So, fewer materials are expected to be on hand after the production.

Table 14 shows the results if we implement the optimised settings for the SS and SLT if we aim for a fill rate of 99,1%. The table shows that the decrease in the fill rate of the portfolio is bigger with the new settings if the forecast is structurally less overforecasted.

Figure 17 shows that the situation of overforecasting is much more expensive than the situation of underforecasting or without structural over or underforecasting. Due to the overforecast, not all materials are used in the production as it was intended to and this increases the inventory positions. In addition to this, the situation of underforecasting is slightly more expensive than the situation without structural over or underforecasting. This is the case since we need more safety stock in the situation of underforecasting.

Both table 13 and 14 shows that if there is less being overforecasted, the results get worse if settings for SS and SLT are not changed. Therefore, the influence of the amount being over or underforecasted is significant.

At last, it is interesting to compare how the settings for SLT and SS change if we move from over forecasting to underforecasting. Table 15 shows that both the SS and the SLT decrease when we move from overforecasting to underforecasting.

Average of parameter	Overforecasting	Neutral	Underforecasting
SLT	0,96 weeks	1,36 weeks	1,42 weeks
SS	4775 kg	5300 kg	8030 kg

Table 15: Average SLT and SS under different forecasting situations

It is interesting to see that if we move from overforecasting to more accurate forecasts, so a neutral situation, the settings for SLT and SS will increase with 41,7% and 11% respectively. So, the main changes occur in the SLT. This can be explained with the fact that overforecasting already leads to an increase in the average inventory value because not all ingredients will be used in the production due to the overforecast. Therefore, there it is not needed to add much SS. The biggest changes, therefore, happen in the SLT. In the situation of underforecasting, SS is needed to cover the uncertainty in the demand pattern. Therefore, we see that moving from neutral to underforecasting, the biggest increase in the safety stock is within the SS. The SLT and the SS both increase with 4,4% and 51,5% respectively.

5.3 Conclusion

The current way of organising the inventory of C&S is not optimal. The service level is low while the average inventory costs are high. The solution is to move to the optimal line by either decreasing the average inventory and/or increasing the service level. Since the lack of availability of C&S's is not the only loss for B&J, we aim for a target fill rate of the portfolio of 99%, which is slightly higher than the service levels that Ben & Jerry's has agreed upon with Unilever. The target fill rate of the portfolio of 99% can be reached by changing the settings for SLT and SS. We have been able to increase the fill rate of the portfolio from 95,8% to 99,1% while decreasing the average inventory value with 15,72%.

The main changes in this optimization process are the result of decreasing the SLT's and compensating this with SS's. The optimization approach was successful for 48 out of the 50 C&S. These 48 C&S had a cost-saving in the average inventory value and/or an increase in the fill rate, which benefits the service level of the portfolio.

Analysing the results on the level of a single chunk or swirl shows that SLT's are more in line with the variability in the lead time. SS can cover the uncertainty in both, the lead time and the production quantity.

Analysing the results on the level of finished products shows that underforecasting finished products can be compensated with overforecasting finished products if both products are using the same chunk or swirl.

Comparing the optimal curves of the different forecasting conditions, we see that the overforecasting lead to a significant increase in the inventory value. This is the case since not all C&S are used in the production as they were planned to be.

Furthermore, if forecasts change from overforecasting demand to underforecasting demand, the performance in terms of the fill rate of the portfolio decreases. Therefore, it is important to be aware of structural changes in the forecast. If there are major changes in the forecasts, the current settings are not sufficient and the analysis should be repeated to optimise the settings.

At last, comparing the optimal solutions under the different forecast conditions, we conclude that moving from overforecasts to unbiased forecasts should mainly result in increasing the SLT's, while moving from unbiased forecasts to underforecasts should mainly result in increasing the SS's.

6 Implementation & Recommendation

This chapter describes an implementation plan and we discuss some recommendations. We will start with a plan of how to implement the results and how to use the model. Thereafter, we advise Ben & Jerry's to change the mindset about SS and SLT and at last, we have a recommendation about becoming more data-driven.

6.1 Implementation results

To implement the results, we propose to set a date in the future on which the new settings for SS and SLT are implemented. In the meantime, Ben & Jerry's, Unilever and the suppliers should enable to switch the parameters for SS and SLT. For Ben & Jerry's it is easy to change the parameters in SAP, but Ben & Jerry's should make sure that the suppliers and the vendor-managed inventories are changing to the new system at the same time. This requires some time.

For example, as described in chapter 1, some C&S are vendor-managed. The material planners and employees of the purchasing department should get in contact with those suppliers to renegotiate the safety inventories. Unilever has the power to influence vendor-managed inventories. To implement all the improvements, Ben & Jerry's should make sure that the safety inventories are also implemented by the vendor-managed inventories and its suppliers.

For the self-owned stocks, Ben & Jerry's could implement the new settings immediately. However, subsection 5.1.3 shows that there were 14 C&S of which the fill rate decreased (and the inventory costs decreased). It shows that not all C&S have an increase in the fill rate. To prevent a situation in which the C&S's with a decrease in fill rate are implemented in an earlier stage than the C&S with an increase in fill rate, Ben & Jerry's should work towards a predefined date in which the changes are made at once. In this way, both the safety inventories of Ben & Jerry's and from its vendors will be implemented at the same time.

Ben & Jerry's could implement the proposed changes for SS and SLT for all C&S's at once or gradually. In the last case, we advise to start with set 2, see section 2.1.3, since this set is a set with multiple C&S's and finished products and therefore it is a good test to check if the advice works out in reality.

Therefore, to implement the situation, we propose to set a date in the future on which all or a part of the new settings for SS and SLT are implemented. In this way, all stakeholders change to the new situation at the same time.

6.2 Implement the usage of the simulation model

First of all, the developed simulation model can be used by Ben & Jerry's if the settings for the SS and SLT should be optimised. We provide Ben & Jerry's a detailed overview in Dutch of how to use the model and how to make sure all the input factors are correctly put into the simulation. In this report, we generally explain when and how Ben & Jerry's should use the simulation model.

New optimizations are required if certain input parameters change or if certain uncertainties evolve. We advised Ben & Jerry's to use the model in the following situations:

- Forecast errors can change, for example, if Unilever decides to have a different way of creating the forecasts. This might affect the accuracy of the forecasts. This affects the uncertainty between the forecasts and demand patterns.
- If new products are added to the portfolio, Ben & Jerry's could run the model for the new finished products. If the new Finished Product is using one of the existing C&S, the affected C&S should be taken into account as well.
- If lead times or the deviations in the lead time change. This affects the uncertainty in the lead time and thus the availability of materials.
- If Unilever decides to use other suppliers, which affects the reliability of resupply.

In addition to this, we have used the model with all Finished Products and C&S included in one simulation model. However, it is easy to only simulate with a small part of the portfolio. This decreases the run time significantly. When running the model, correlations between C&S and Finished Products should always be taken into account.

6.3 Lead times

The purchase department has provided a lot of information about different lead times, as discussed in detail in chapter 2. Since there are many differences between the lead times in SAP, the contract and all other information, we advise the material planners to reconsider all lead times in SAP.

6.4 Mindset about SS and SLT

At this point, Ben & Jerry's prefers SLT over SS. As stated in chapter 3, this perception is broadly assumed within different industries with dependent demand. However, this research shows that high SLT's are suboptimal to having lower SLT's and adding some SS since all changes had a decrease in the SLT and some SS was added. This is a mindset that should change. SLT is not always better than SS.

6.5 Improvement of the forecast

The forecasts are created by Unilever and Ben & Jerry's is using the forecasts. With the overforecasts Unilever makes sure that finished products are not getting out of stock in supermarkets. However, the consequences of the overforecasts lead to an increase in the inventory of C&S for Ben & Jerry's. Therefore, Ben & Jerry's should suggest to Unilever to improve the Forecasts for finished products.

6.6 Data-driven

In the past half a year, we have seen that the employees of Ben & Jerry's are very good at reacting to situations that happen. If supply arrives late, the planners do everything that is in their power to make sure that the supply will still arrive on time. Production planners react to fluctuating demand, stockouts and lower production outputs than expected. Therefore, we see that there is much experience within the organisation.

However, we had difficulties in gathering certain datapoints and this is the result of the fact that Ben & Jerry's is not gathering, using and saving specific data. There is so much data that could help Ben & Jerry's to become more efficient on an operational level, tactical level and even on a strategic level. And the most important thing, every day new data is created that could help Ben & Jerry's improve on these different levels.

In this research, we have put a lot of effort into gathering information about the lead times. We wanted to analyse how the standard deviation of the lead time was distributed. This is a very important aspect for the material availability and thus when to order on an operational level. In addition to this, measuring the reliability of suppliers, based on the supply timing and the supply quantity, would help Ben & Jerry's and Unilever to identify which supplier is good or bad. This information could be used in the negotiation with the suppliers. In addition to this, another variable that could help Ben & Jerry's to improve the efficiency from the perspective of planning and logistics is measuring the line efficiency per finished product. This could e.g. make the planning for each production run more accurate. One of the drawbacks of ERP systems like SAP, as described in section 3.1, is that some parameters are considered to be fixed. However, measuring uncertainties like the variation in the lead time could improve the business. The point to be made is, being more proactive in collecting and using data could make your Ben & Jerry's become more efficient and effective.

7 Final conclusion

This research answers the following question:

What improvements to the inventory policies of chunks and swirls should Ben & Jerry's make and implement to be able to stick to the optimal production cycle?

The inventory of C&S is optimised with a simulation model. By using an algorithm we could predict the optimal SS and SLT for each C&S and with this, we could optimise the situation at Ben & Jerry's. We are able to increase the service level from 95,8% to 99,1% while the average investment in inventory can be reduced with 15,72%.

The current situation at Ben & Jerry's is very inefficient because the inventory of C&S was not in line with the new and fixed production cycles. When they have implemented the cycles, the SS and SLT's were not analysed and adjusted. In addition to this, they have implemented high SLT's to have the possibility to use inventory of C&S's that was intended to be used in a later production cycle.

Furthermore, Ben & Jerry's should take note of the following findings:

- SLT is not better than SS. There should be a balance between both to cover the uncertainty in both the lead time and the demand pattern.
- Overforecasting increases the inventory positions and thus the investments in inventory. Ben & Jerry's should suggest Unilever to improve the forecasts.
- If forecast errors structurally change, the fill rate and the inventory levels of the portfolio will immediately be affected. The analysis should be repeated to optimise the SS and SLT.
- Overforecasts can be compensated with underforecasts
- Moving from overforecasting to unbiased forecasts, will mainly increase the SLT's, while moving from unbiased forecasts to structurally underforecasting will mainly lead to higher SS's.

To implement the results, Ben & Jerry's should not only adjust the parameters in SAP but also the suppliers and the vendor-managed inventories should implement the changes. Therefore, Ben & Jerry's, material planners and the purchasing department of Unilever should make sure to implement the results.

Ben & Jerry's is advised to do the analysis when important parameters are changing, like the uncertainty in lead time or demand. Other parameters that influence the availability of materials are if new products are introduced or if production cycles are changed.

To make sure that Ben & Jerry's keeps track of changes in the forecasts biases or the uncertainty in the lead times, they should try to become more data-driven. There is much experience in reacting to certain situations, but we recommend to be continuously collecting and analysing these important data points. It could help Ben & Jerry's on each level, from the operational to tactical and strategical level.

Sources

Barrow, D. K., & Kourentzes, N. (2016). Distributions of forecasting errors of forecast combinations: Implications for inventory management. *International Journal of Production Economics*, *177*, 24–33. https://doi.org/10.1016/j.ijpe.2016.03.017

Ben & Jerry's. (2019). *Cone Together Ice Cream | Ben & Jerry's*. Https://Www.Benjerry.Co.Uk. https://www.benjerry.co.uk/flavours/cone-together-ice-cream/tubs

Bonate, P. L. (2001). A Brief Introduction to Monte Carlo Simulation. *Clinical Pharmacokinetics*, *40*(1), 15–22. https://doi.org/10.2165/00003088-200140010-00002

Boom, M. (2012, December 6). *Simulation Lecture 5: Random-number generators* [Slides]. Http://Www.Win.Tue.Nl/Courses/2WB05. http://www.win.tue.nl/courses/2WB05

Boston Univeristy. (2019). *The Correlation Coefficient (r)*. Https://Sphweb.Bumc.Bu.Edu/Otlt/MPH-Modules/PH717-QuantCore/PH717-Module9-Correlation-Regression/PH717-Module9-Correlation-Regression4.Html. Retrieved 13 April 2021, from https://sphweb.bumc.bu.edu/otlt/MPH-Modules/PH717-QuantCore/PH717-Module9-Correlation-Regression/PH717-Module9-Correlation-Regression4.html

Buzacott, J. A., & Shanthikumar, J. G. (1994). Safety Stock versus Safety Time in MRP Controlled Production Systems. *Management Science*, *40*(12), 1678–1689. https://doi.org/10.1287/mnsc.40.12.1678

Carson, J. S. (2002). *MODEL VERIFICATION AND VALIDATION*. 2002 Winter Simulation Conference.

Chopra, S., & Meindl, P. (2015). *Supply Chain Management* (6th edition). Pearson Education Limited.

Das, S. K., & Abdel-Malek, L. (2003). Modeling the flexibility of order quantities and lead-times in supply chains. *International Journal of Production Economics*, *85*(2), 171–181. https://doi.org/10.1016/s0925-5273(03)00108-7

Dolgui, A., & Prodhon, C. (2007). Supply planning under uncertainties in MRP environments: A state of the art. *Annual Reviews in Control*, *31*(2), 269–279. https://doi.org/10.1016/j.arcontrol.2007.02.007

Hopp, W. J., & Spearman, M. L. (1993). SETTING SAFETY LEADTIMES FOR PURCHASED COMPONENTS IN ASSEMBLY SYSTEMS. *IIE Transactions*, *25*(2), 2–11. https://doi.org/10.1080/07408179308964272

Ingalls, R. G. (2008). *INTRODUCTION TO SIMULATION*. 2008 Winter Simulation Conference.

International Conference on Simulation in Engineering Education, Knadler, C. E., Jr, Simulation, S. F. C., Tempe, S. A. W. M., & Vakilzadian, H. (1994). *International Conference on Simulation in Engineering Education: Proceedings of the 1994 Western Multiconference :*

January 24–26, 1994, Radisson Tempe . . . Hotel, Tempe, Arizona (Simulation Series). Society for Computer Simulation.

Jodlbauer, H., & Reitner, S. (2012). Optimizing service-level and relevant cost for a stochastic multi-item cyclic production system. *International Journal of Production Economics*, *136*(2), 306–317. https://doi.org/10.1016/j.ijpe.2011.12.015

Law, A. M. (2014). *Simulation Modeling and Analysis*. McGraw-Hill Education.

ÖZgün, O., & Barlas, Y. (2009). *Discrete vs. Continuous Simulation: When Does It Matter?*

Reedy, S. (2021, April 9). *What is a Bill of Materials (BOM) and How Do You Create One?* Arena. https://www.arenasolutions.com/resources/category/bom-management/creating-a-bill-of-materials/

Sargent, R. G. (2010, December). Verification and validation of simulation models. *Proceedings of the 2010 Winter Simulation Conference*. https://doi.org/10.1109/wsc.2010.5679166

Schalit, S., & Vermorel, J. (2014). *Service Level Definition (Supply Chain)*. Lokad - Quantitative Supply Chain. Retrieved 6 May 2021, from https://www.lokad.com/service-level-definition

Shannon, R. (1975). *Systems Simulation: The Art and Science by Shannon, Robert E. (1975) Hardcover*. 1975.

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). *Inventory and Production Management in Supply Chains*. Taylor & Francis.

Smit, J. (2019). Het grote gevecht. In *Het grote gevecht* (pp. 51–52). Prometheus.

Yao, Y., Evers, P. T., & Dresner, M. E. (2007). Supply chain integration in vendor-managed inventory. *Decision Support Systems*, *43*(2), 663–674. https://doi.org/10.1016/j.dss.2005.05.021

Appendix A – Proof of normally distributed forecast errors

To do this, we compute all forecast errors and measure the minimum, maximum, average and standard deviation. We create 9 intervals, so-called bins, that contain one part between the maximum and minimum forecast error. After that, we count for each bin how many forecast errors are observed in the corresponding interval and we compare it with the expected number of observations in the interval. This is shown in figure 18. The last step is to measure the error between the expected and observed numbers and by applying the chi distribution we most of the time can prove that that the forecast error is normally distributed. Sometimes a few forecast errors are deleted in this analysis if the forecast was deviation too much from actual demand.

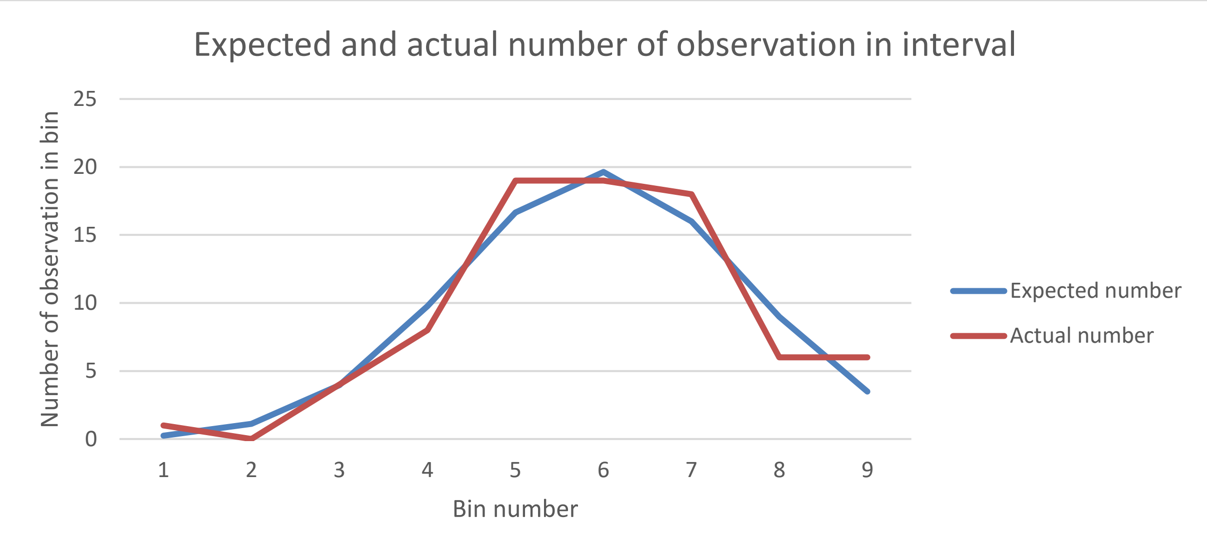

Figure A1: Expected number and observed number of forecast errors in bin

When it is motivated that a forecast error is normally distributed with a measured mean and standard deviation, we can use this to generate a realistic demand pattern. The demand pattern is created by adding a randomly generated forecast error to the forecast of that period. Figure 19 shows the actual demand (blue) and the generated actual demand (red). It shows that is possible to generate realistic demand patterns.

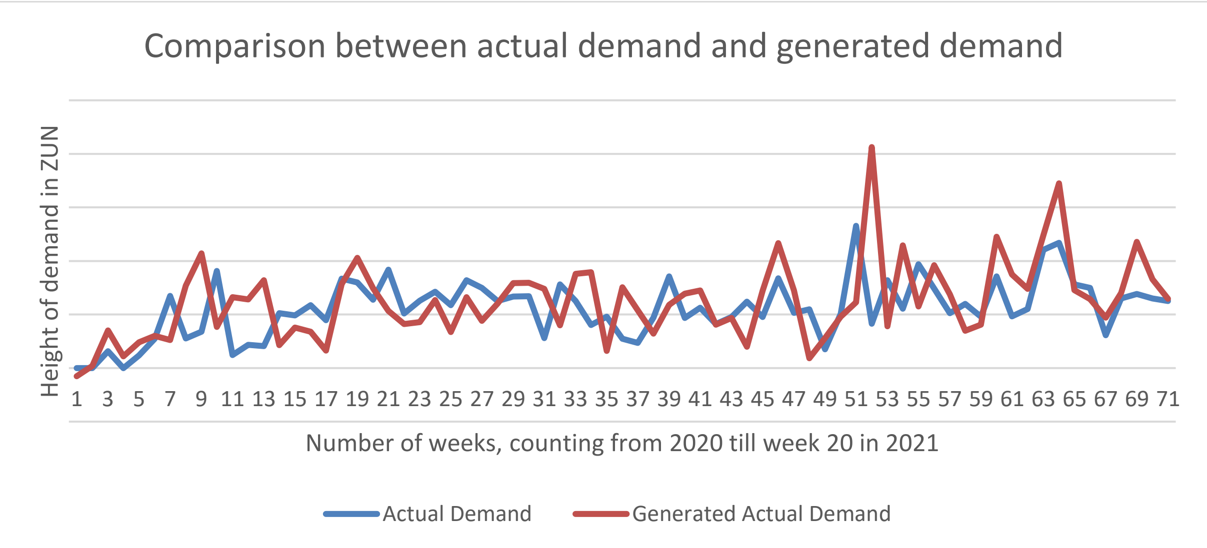

Figure A2: Comparison between actual demand and generated demand.

Since two different forecast patterns can be used for this analysis, both analyses are executed and one with the best fit is chosen.

At last, there is no correlation found on a weekly level that states that all ice creams have over- or under forecasting in the same period.

Appendix B – simulation settings

Appendix A describes the analysis of the warm-up period, run length and the number of replications. The approach that is applied is the method by Welch (1983), which is described in the book by Law (2014).

Warm-up period

To determine the warm-up period, we use data generated by the simulation model. We have chosen to make n = 5 replications and a run length of 10 years. The settings that are used for the SLT and SS are the settings that are currently implemented by Ben & Jerry's, equal to our zero measurement. As discussed before, the simulation is non-terminating and it contains a steady-state cycle. Each replication has the same initial conditions and the seed value of the random numbers continues in the next replication.

In the simulation model, the initial on hand stock of ingredients is initialised in a way that during the lead time and SLT of the chunk or swirl, the demand is fully fulfilled. We initialise the model in this way because no order can arrive before the week that equals the lead time + SLT. By initialising the model in this way, we aim to reach a steady state as soon as possible and to keep the run length of the model as low as possible. The on hand inventory of C&S's influence the amount of backorder on the level of C&S's and thus on whether there is a part of the production postponed. Therefore, we have applied the approach of Welch (1983) based on the ending on hand inventory. The results are shown in the figure below. The blue line is the average ending on hand inventory of the 5 replications. The red and green lines provide a 5 and 10 weekly average. The analysis is executed with the longest lead time and safety lead time. We believe that the simulation is not depending on the initial condition after half a year, so the first 26 weeks are used as a warm-up period.

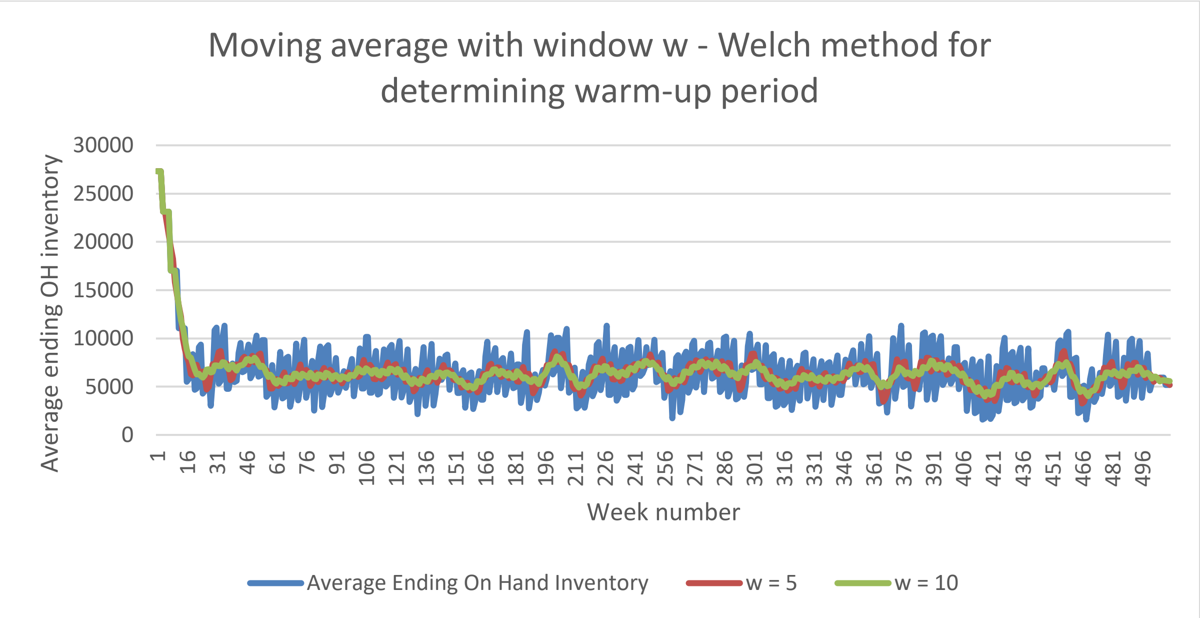

Figure B1: Welch's method to determine warm-up period

7.1.1 Run-length

The run length could be determined on two different levels. The most important KPI is to reach the target service level and therefore we can measure the weighted average service level over time. This analysis is executed for a combination of set 2 and 4, for the definition see section 2.1.3. The results are depicted in the figure below. It shows that the service level of the three finished products and the weighted average service level is stable after approximately 99 days.

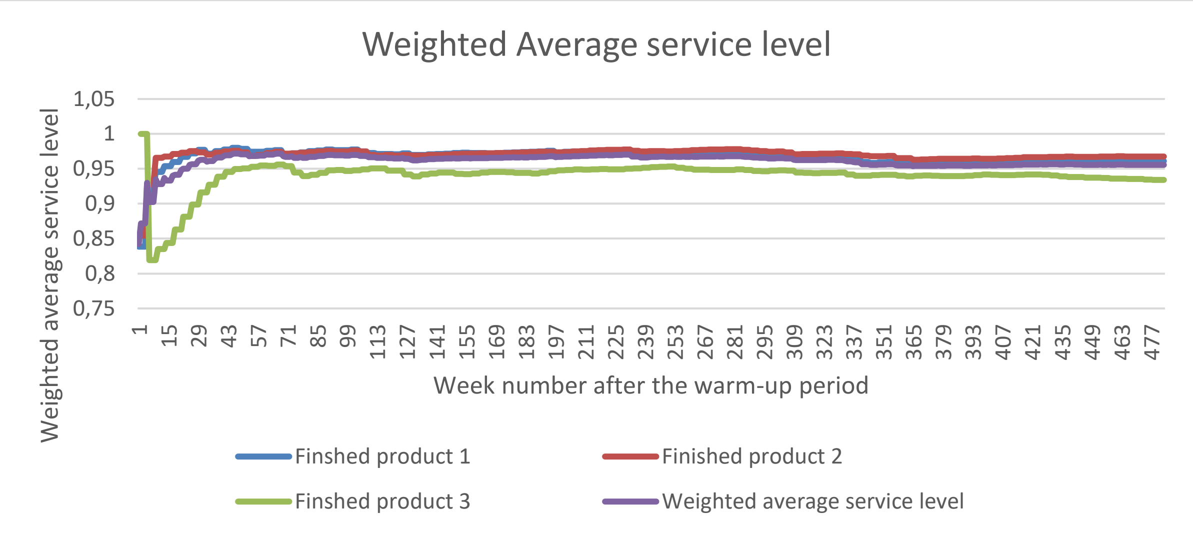

Figure B2:Analysis to determine Run Length

However, in the optimization phase, we are measuring the fill rate of each chunk and swirl to determine which one should have more SS or SLT. Therefore, we want the fill rate on a chunk level to be stable as well. The figure below shows that the fill rate of a chunk or swirl is stable after approximately 239 weeks. Adding those to the warm- up period of 26 weeks, we should implement a run length of 5 years. Due to seasonal patterns, we have rounded it to an integer number of years, so 5 years.

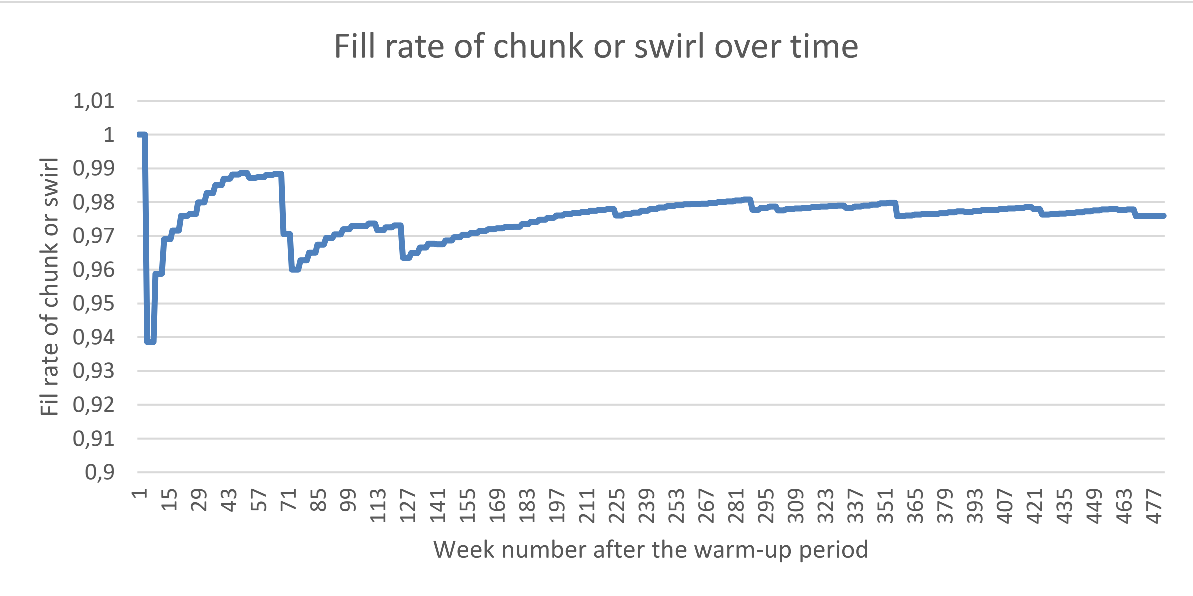

Figure B3:Analysis to determine Run Length

7.1.2 Number of replications

For the number of replications, we measure the performance of the weighted average service level after the warm-up period, until the end of the run length. We apply the following formula's:

$$\frac{t_{n-1,1-\alpha/2} * \sqrt{S^2/n}}{\bar{X}} < \gamma' \quad \text{(Law)}$$

$$\gamma' = \frac{\gamma}{\gamma + 1} \qquad \text{(Law)}$$

We have executed the analysis with the weighted average (WA) service level and the total inventory value of the whole portfolio.

Replication	WA service level	Average until n	Var	Tvalue	Error
n = 1	0,948628	0,948628	-	-	-
n = 2	0,962366	0,948628	9,43635E-05	12,7062	0,091343
n = 3	0,959597	0,957069	5,27861E-05	4,302653	0,018858
n = 4	0,957684	0,957764	3,5359E-05	3,182446	0,009879
n = 5	0,960547	0,957764	2,89384E-05	2,776445	0,006974

Tabel B1: Results Welch approach

Replication	Inventory value	Average until n	Var	Tvalue	Error
n = 1	€ 2.786.540,42	€ 2.786.540,42	-	-	-
n = 2	€ 2.752.221,08	€ 2.769.380,75	€ 588.908.297,50	12,7062	0,07873
n = 3	€ 2.731.415,02	€ 2.756.725,51	€ 774.919.751,96	4,302653	0,025085
n = 4	€ 2.774.490,97	€ 2.761.166,87	€ 595.516.105,70	3,182446	0,014063
n = 5	€ 2.738.062,74	€ 2.756.546,05	€ 553.397.285,54	2,776445	0,010596

Tabel B2: Results Welch approach

Using both KPI's, the weighted average service level and the total inventory value, the number of replications should equal n = 3.

Appendix C – LCG – random number generator

The LCG, Linear Congruential Generators, is introduced by Lehmer (1951) and it creates a sequence of integer values, Z_1, Z_2, …. The sequence is generated with the following formula:

$$Z_i = (a * Z_{i-1} + c) \ mod \ m$$

Where,

m	=	modulus,	nonnegative integer,	$m > 0$
a	=	multiplier	nonnegative integer,	$a > m$
c	=	increment	nonnegative integer,	$c > m$
Z_0	=	seed value,	nonnegative integer,	$0 \leq Z_i \leq m - 1$

The random number is generated via the following formula:

$U_i = \frac{Z_i}{m}$, which generates a random number in the interval [0,1]

The question is which values for these parameters provide a good set of random numbers. Although research provides such parameters, we have developed our own sequence. A good set of random numbers has the following properties (Boon, 2012):

- Uniformity: The numbers generated appear to be distributed uniformly on (0,1)
- Independence: The numbers generated show no correlation with each other
- Replication: The number should be replicable (e.g. for debugging or comparison)
- Cycle Length: It should take long before numbers start to repeat
- Speed: The generator should be fast
- memory usage: The generator should not require a lot of storage.

We have tested different parameters for the LCG and eventually we have come up with the following settings: Modulus = 2^24, multiplier = 1485, increment = 104729, Seed value = 15001. This sequence is motivated to be uniformly distributed for the first 100, 500, 5.000, 10.000, 25.000 and 50.000 random numbers. Furthermore, the average of the random numbers is 0,50 and it does not have any repeating value in at least the first 230.000 random numbers. At last, plotting the first 100 and 1000 random numbers does not show any structure.

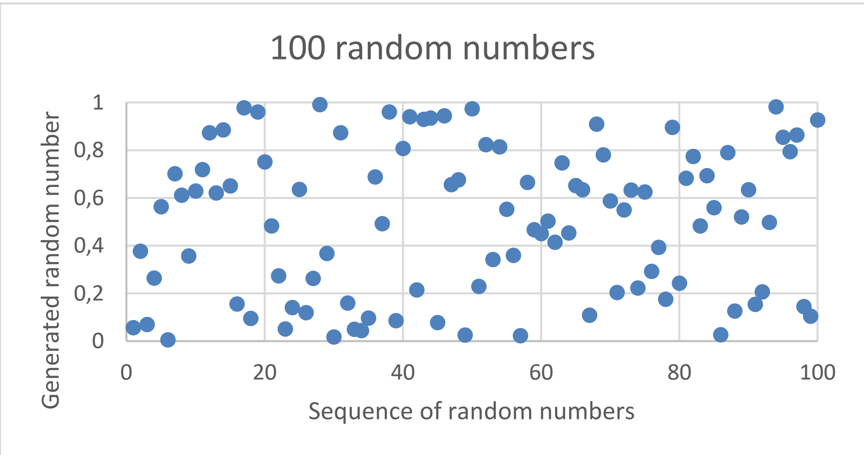

Figure C1: First 100 random numbers

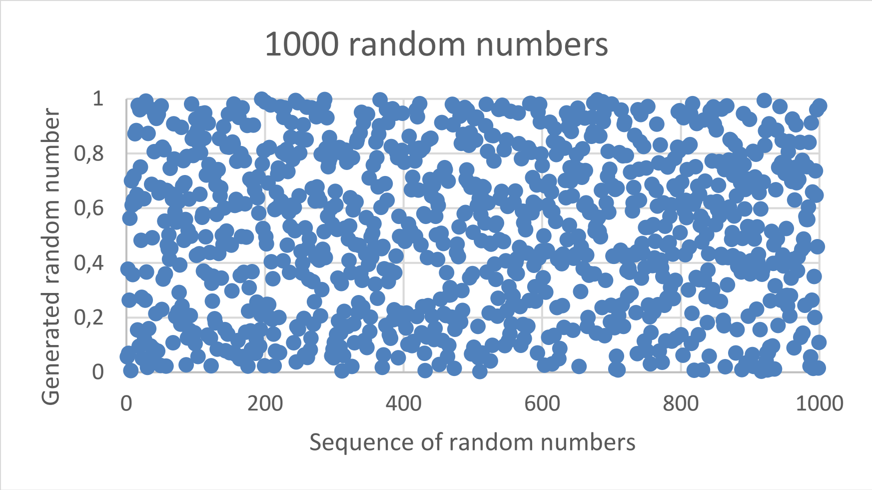

Figure C2: First 1000 random numbers

The sequence of random numbers is motivated to be uniformly distributed. The figure below shows the expected and counted number of random numbers in each bin, where each bin represent 0,01. Comparing the errors with the Chi-Square test, where the level of significance of the test equals α = 0,05 and 99 degrees of freedom, the distribution is motivated to be uniformly distributed on the first 100, 500, 5.000, 10.000, 25.000 and 50.000 random numbers.

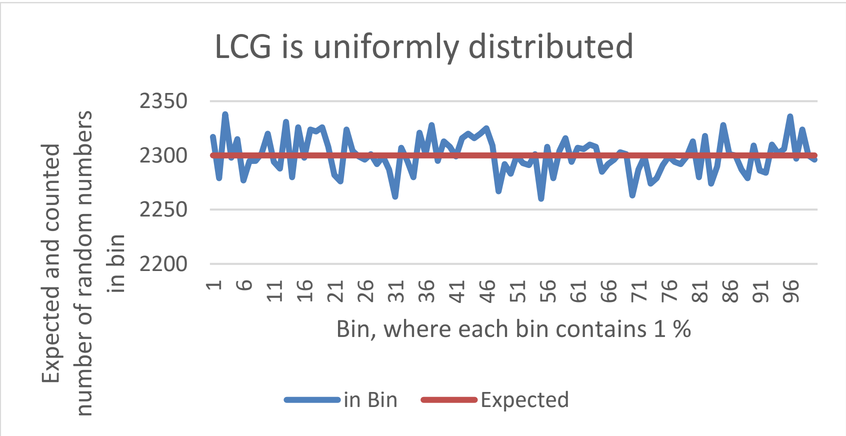

Figure 25: Analysis uniformly distribution of random numbers.

Appendix D explains how Synchronicity is built into the model.

Appendix D – synchronicity

In table D.1, the inventory policies of two C&S's are depicted. In the green cells, the orders are shown. The fifth order is placed in week 11 and it arrives in week 21, the red cell.

	2020 1	2020 2	2020 3	2020 4	2020 5	2020 6	2020 7	2020 8	2020 9	2020 10	2020 11	2020 12	2020 13	2020 14	2020 15	2020 16	2020 17	2020 18	2020 19	2020 20	2020 21	2020 22
Beginning On Hand inventory	16628,67	16628,67	16628,67	13558,48	13558,48	13558,48	10980,7	10980,7	10980,7	4013,894	4013,894	4013,894	15038,21	15038,21	15038,21	14266,08	14266,08	14266,08	9925,985	9925,985	9925,985	22065,18
Received replenishment orders												13608									13608	
Initial Demand	0	0	3070,197	0	0	2577,773	0	0	6966,81	0	0	2583,684	0	0	772,1294	0	0	4340,095	0	0	1468,807	0
Forecasted demand Ingredients			4398,501			4680,357			8669,029			5113,313			6134,297			5278,515			4170,821	
Initial Demand fulfilled from stock	0	0	3070,197	0	0	2577,773	0	0	6966,81	0	0	2583,684	0	0	772,1294	0	0	4340,095	0	0	1468,807	0
Adjusted demand	0	0	3070,197	0	0	2577,773	0	0	6966,81	0	0	2583,684	0	0	772,1294	0	0	4340,095	0	0	1468,807	0
Ingredient Demand backordered	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Ending On Hand Inventory	16628,67	16628,67	13558,48	13558,48	13558,48	10980,7	10980,7	10980,7	4013,894	4013,894	4013,894	15038,21	15038,21	15038,21	14266,08	14266,08	14266,08	9925,985	9925,985	9925,985	22065,18	22065,18
Total Pipeline			13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608		
Inventory position	16628,67	16628,67	27166,48	27166,48	27166,48	24588,7	24588,7	24588,7	17621,89	17621,89	17621,89	28646,21	28646,21	28646,21	27874,08	27874,08	27874,08	23533,99	23533,99	23533,99	22065,18	22065,18
Forecast during R+L+SL	16628,67	[illegible]	17258,06	17513,57	23099,16	18850,21	19102,64	23919,01	15497,68	15769,75	19607,68	14729,06	14968,89	21294,61	15335,09	15622,86	21768,65	16712,22	16953,75	22876,83	19030,05	19325,21
Replenishment order		13608									13608											
		1st order									5th order											

	2020 1	2020 2	2020 3	2020 4	2020 5	2020 6	2020 7	2020 8	2020 9	2020 10	2020 11	2020 12	2020 13	2020 14	2020 15	2020 16	2020 17	2020 18	2020 19	2020 20	2020 21	2020 22
Beginning On Hand inventory	8508,886	8508,886	8508,886	5556,546	5556,546	5556,546	3077,726	3077,726	3077,726	4378,352	4378,352	4378,352	5893,849	5893,849	5893,849	9151,359	9151,359	9151,359	4977,868	4977,868	8177,868	6765,445
Received replenishment orders									8000			4000			4000					3200		1600
Initial Demand	0	0	2952,341	0	0	2478,82	0	0	6699,374	0	0	2484,504	0	0	742,4895	0	0	4173,491	0	0	1412,424	0
Forecasted demand Ingredients			4229,655			4500,691			8336,249			4917,027			5898,819			5075,888			4010,715	
Initial Demand fulfilled from stock	0	0	2952,341	0	0	2478,82	0	0	6699,374	0	0	2484,504	0	0	742,4895	0	0	4173,491	0	0	1412,424	0
Adjusted demand	0	0	2952,341	0	0	2478,82	0	0	6699,374	0	0	2484,504	0	0	742,4895	0	0	4173,491	0	0	1412,424	0
Ingredient Demand backordered	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Ending On Hand Inventory	8508,886	8508,886	5556,546	5556,546	5556,546	3077,726	3077,726	3077,726	4378,352	4378,352	4378,352	5893,849	5893,849	5893,849	9151,359	9151,359	9151,359	4977,868	4977,868	8177,868	6765,445	8365,445
Total Pipeline			8000	8000	8000	12000	12000	12000	8000	8000	8000	7200	7200	7200	4800	4800	4800	6400	6400	3200	8800	7200
Inventory position	8508,886	8508,886	13556,55	13556,55	13556,55	15077,73	15077,73	15077,73	12378,35	12378,35	12378,35	13093,85	13093,85	13093,85	13951,36	13951,36	13951,36	11377,87	11377,87	11377,87	15565,44	15565,44
Forecast during R+L+SL	8508,887	[illegible]	12239,73	[illegible]	17112,17	12901,02	13018,44	[illegible]	10393,25	10534,12	15387,23	10611,07	10746,82	14605,95	8799,272	8894,608	15305,25	10314,71	10400,06	16646,55	12833,07	12977,18
Replenishment order		8000			4000			4000			3200			1600			1600			5600		
		2d order			3d order			4th order														

Table D.1

In the following situation, see table D.2 the second chunk (lowest table) has an extra week safety lead time. Therefore the orders of the second chunk happen one week before the former situation. In addition to this, the system has placed an extra order.

	2020 1	2020 2	2020 3	2020 4	2020 5	2020 6	2020 7	2020 8	2020 9	2020 10	2020 11	2020 12	2020 13	2020 14	2020 15	2020 16	2020 17	2020 18	2020 19	2020 20	2020 21	2020 22
Beginning On Hand inventory	16629	16629	16629	13558	13558	13558	10981	10981	10981	4013,9	4013,9	4013,9	15038	15038	15038	14266	14266	14266	9926	9926	9926	8457,2
Received replenishment orders												13608										13608
Initial Demand	0	0	3070,2	0	0	2577,8	0	0	6966,8	0	0	2583,7	0	0	772,13	0	0	4340,1	0	0	1468,8	0
Forecasted demand Ingredients			4398,5			4680,4			8669			5113,3			6134,3			5278,5			4170,8	
Initial Demand fulfilled from stock	0	0	3070,2	0	0	2577,8	0	0	6966,8	0	0	2583,7	0	0	772,13	0	0	4340,1	0	0	1468,8	0
Adjusted demand	0	0	3070,2	0	0	2577,8	0	0	6966,8	0	0	2583,7	0	0	772,13	0	0	4340,1	0	0	1468,8	0
Ingredient Demand backordered	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Ending On Hand Inventory	16629	16629	13558	13558	13558	10981	10981	10981	4013,9	4013,9	4013,9	15038	15038	15038	14266	14266	14266	9926	9926	9926	8457,2	22065
Total Pipeline			13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	13608	
Inventory position	16629	16629	27166	27166	27166	24589	24589	24589	17622	17622	17622	28646	28646	28646	27874	27874	27874	23534	23534	23534	22065	22065
Forecast during R+L+SL	16629	[illegible]	17258	17514	23099	18850	19103	23919	15498	15770	19608	14729	14969	21295	15335	15623	21769	16712	16954	22877	19030	19325
Replenishment order		13608									13608											
		2nd order									6e orde											

	2020 1	2020 2	2020 3	2020 4	2020 5	2020 6	2020 7	2020 8	2020 9	2020 10	2020 11	2020 12	2020 13	2020 14	2020 15	2020 16	2020 17	2020 18	2020 19	2020 20	2020 21	2020 22
Beginning On Hand inventory	8508,9	8508,9	8508,9	5556,5	5556,5	5556,5	3077,7	3077,7	11078	4378,4	4378,4	8378,4	5893,8	5893,8	9893,8	9151,4	9151,4	12351	8177,9	8177,9	8177,9	8365,4
Received replenishment orders								8000			4000			4000			3200				1600	
Initial Demand	0	0	2952,3	0	0	2478,8	0	0	6699,4	0	0	2484,5	0	0	742,49	0	0	4173,5	0	0	1412,4	0
Forecasted demand Ingredients			4229,7			4500,7			8336,2			4917			5898,8			5075,9			4010,7	
Initial Demand fulfilled from stock	0	0	2952,3	0	0	2478,8	0	0	6699,4	0	0	2484,5	0	0	742,49	0	0	4173,5	0	0	1412,4	0
Adjusted demand	0	0	2952,3	0	0	2478,8	0	0	6699,4	0	0	2484,5	0	0	742,49	0	0	4173,5	0	0	1412,4	0
Ingredient Demand backordered	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Ending On Hand Inventory	8508,9	8508,9	5556,5	5556,5	5556,5	3077,7	3077,7	11078	4378,4	4378,4	8378,4	5893,8	5893,8	9893,8	9151,4	9151,4	12351	8177,9	8177,9	8177,9	8365,4	8365,4
Total Pipeline		8000	8000	8000	12000	12000	12000	8000	8000	8000	7200	7200	7200	4800	4800	4800	3200	3200	3200	8800	7200	7200
Inventory position	8508,9	16509	13557	13557	17557	15078	15078	19078	12378	12378	15578	13094	13094	14694	13951	13951	15551	11378	11378	16978	15565	15565
Forecast during R+L+SL	[illegible]	16359	12240	[illegible]	17112	12901	18369	18612	10393	15164	15387	10611	14394	14606	8799,3	15023	15305	10315	16303	16647	12833	18583
Replenishment order	8000			4000			4000			3200			1600			1600			5600			3200
	1st order			3rd order			4th order			5th order												

Table D.2

In the first table, table D.1, we saw that the 5th order of the simulation arrived in week 22. But due to the change in the settings of the second chunk, the 5th order is now the 6th order of the simulation. And therefore, the same order has gained a different random number. The result is that the settings of the first chunk are not changed, but its fill rate is decreased due to different random numbers.

The solution is either to make the run length longer, to erase such inconsistencies, or to prevent these inconsistencies to occur. It shows that we should not only have a fixed sequence of random numbers, but we have to make sure that every order has always the same random number. Our solution is to draw a random number for each possible order. In this way, we prevent those inconsistencies to occur.

This synchronicity does not only prevent the described inconsistencies, but it also takes care that synchronicity is maintained for all replications, and not only for the first replication.

Appendix E – Develop logarithmic functions

A logarithmic regression line is defined as following:

$$y = A + B * ln(x)$$

In this research, y represents the fill rate of a chunk or swirl and x takes the value for either the SS or the inventory value.

The parameters A and B are retrieved with the following steps:

i) First compute the mean:

$$\overline{ln(x)} = \frac{\sum Ln(x)_i}{n}$$

$$\overline{y} = \frac{\sum y_i}{n}$$

ii) Secondly, calculate S_{xx},, S_{yy} and S_{xy}:

$$S_{xx} = \sum ln(x_i)^2 - n * \overline{ln(x)}^2$$

$$S_{yy} = \sum {y_i}^2 - n * \overline{y}^2$$

$$S_{xy} = \sum ln(x_i)y_i - n * \overline{ln(x)} * \overline{y}$$

iii) Compute A and B as following:

$$B = \frac{Sxy}{Sxx}$$

$$A = \overline{y} - B\,\overline{ln(x)}$$