Master of Science Thesis

MIOM05 Degree Project in Production Management

Optimizing Finished Goods Safety Stock Levels at a Swedish Food Company

Authors

Kristinn Jonasson & Emil Blomborg

Supervisor

Danja Sonntag

Preface

With the completion of this Master's thesis, our five-year journey studying Mechanical Engineering at LTH in Lund comes to a close. Throughout these years, we have acquired knowledge and experiences that have been instrumental in the development of this thesis.

We owe a great deal of gratitude to several individuals who have supported us throughout this journey. Foremost, we would like to express our sincere appreciation to our thesis supervisor, Danja Sonntag, Associate Professor at the Division of Production Management, for her invaluable support, guidance, and insights throughout the project.

We are also thankful to everyone at the case company who included us in their project, listened to our input, and allowed us to contribute. We recognize that such trust in our work is not commonly extended across the industry, and we believe that this project has significantly contributed to our growth and continued learning.

Lund, May 2024
Kristinn Jonasson & Emil Blomborg

Abstract

With rising interest rates increasing the cost of capital, businesses are focusing on reducing expenses. Tied-up capital, largely represented by inventory, is a critical cost driver. Therefore, optimizing inventory levels is crucial for controlling capital costs. Inventory levels consist of cycle stock and safety stock. Cycle stock covers average demand during replenishment, while safety stock mitigates the risk of stock-outs due to uncertainties. This thesis focuses on a Swedish food manufacturer that has implemented a new planning system, SAP IBP, featuring an Inventory Optimizer (IO) module designed to calculate safety stocks for its finished goods warehouses. Despite this, the company has been hesitant to rely on its outputs, leading to manual adjustments of safety stock levels based on staff experience. The primary purpose of this thesis is therefore to develop an inventory model that optimizes safety stock levels based on a target service level, enabling it to serve as a benchmark for assessing the IO's performance in the future.

To fulfill this purpose, the thesis utilizes the first four steps of the operations research framework: defining the problem and collecting relevant data, formulating a mathematical model, creating a computer-based solution based on that mathematical model, and validating the solution through practical testing. This resulted in an inventory model implemented in Microsoft Excel. The findings show that the new safety stocks could be reduced by 47% compared to the current levels at the case company while meeting the company's service level targets. Additionally, the model has been tested and the results indicate that it performs well enough to be utilized as a benchmark for the IO in the future.

Keywords: inventory control, forecast error, operations research

Contents

List of Figures

List of Tables

Abbreviations

3PL	Third-Party Logistics
CDF	Cumulative Distribution Function
CM	Contribution Margin
CV	Coefficient of Variation
IL	Inventory Level
IO	Inventory Optimizer
IP	Inventory Position
MAD	Mean Absolute Deviation
MAPE	Mean Absolute Percentage Error
OR	Operations Research
PDF	Probability Density Function
SKU	Stock-Keeping Unit
SL	Service Level
SS	Safety Stock
TSL	Target Service Level

1 Introduction

1.1 Background

With rising interest rates leading to an increased cost of capital, businesses are placing a high emphasis on reducing their expenses. In this context, tied-up capital becomes a critical cost driver to consider. Usually, a major part of this tied-up capital is represented by the inventory a company holds. Therefore, optimizing inventory levels is important to control a company's capital costs. Inventory levels are composed of cycle stock and safety stock. Cycle stock covers the average demand during the replenishment period. However, due to uncertainties, safety stock is needed to mitigate the risk of stock-outs.

Optimizing safety stock levels is a way to reduce tied-up capital while achieving a desired customer service level. Some companies rely on manual adjustments based on personnel experience, which often leads to safety stock levels being too high. The main reason for this is that there is a trade-off between low inventory levels and achieving high service levels. When service levels fall below the agreed-upon standards, companies typically incur penalty fees and a risk of losing customers to competitors. Without customers, there is no business. In an effort to minimize the risk of failing to meet required service levels, companies often maintain higher safety stock levels than necessary to meet the service level targets.

This thesis focuses on inventory management at a manufacturer in the Swedish food industry, referred to as the *"case company"*, that is facing the challenge of rising capital costs. The company currently relies on personnel experience when determining safety stock levels and is consequently seeking an understanding of the current performance levels and if they can be improved.

1.2 Supply chain mapping and thesis delimitations

The case company is a large food manufacturer in Sweden, and part of a larger conglomerate. It specializes in production of a wide range of food products that are frequently found in Swedish households. These include, for example, frozen or refrigerated ready meals and pantry items.

Operating across Sweden, the company manages several production facilities and finished goods warehouses. This thesis narrows its focus to a specific finished goods warehouse, known as *"Warehouse Y"*. Notably, Warehouse Y operates under a third-party logistics (3PL) model, distinguishing it from the company's other warehouses due to the higher inventory holding costs associated with 3PL operations. Therefore, the case company decided to focus on this particular warehouse. This warehouse receives products from two production plants. The case company has recently initiated an attempt to lower the safety stock on some of the products that are produced at the plant, referred to as *"Plant X"*, and stored at warehouse Y. As a result, the case company requested that the focus should be to consider the safety stocks of products produced there. Therefore, this thesis focuses on products manufactured at Plant X and the other plant is out of scope. Although the primary emphasis is on inventory at Warehouse Y, some factors impacting safety stocks, such as lead times, are specific to the production plant, making Plant X relevant.

1.3 Problem formulation and purpose

This thesis follows a methodology known as *operations research* (OR), which is a six-step approach to solve an operational problem presented by Hillier and Lieberman (2010). The first step in this approach is to define the problem. In addition, the first step involves collecting relevant data, which is presented in chapter 4. The problem is that the case company implemented the planning system, SAP IBP, which includes an *Inventory Optimizer* (IO) module, designed to calculate safety stocks for all its finished goods warehouses. Despite this advanced tool, the company has been hesitant to rely on its output, primarily due to a lack of understanding about how the module functions and the input parameters it requires. As a result, safety stock levels have been manually adjusted based on staff experience. To address this issue, the company initiated a project aimed at optimizing the use of the IO by understanding how different input parameters influence its output, with the ultimate goal of using the IO's calculations. The primary purpose of this thesis is to develop an inventory model that optimizes safety stock levels based on a target service level, enabling it to serve as a benchmark for assessing the IO's performance in the future.

1.4 Report disposition

The thesis starts with a brief overview of the case company in chapter 2, touching on its policies and processes. In chapter 3, a theoretical framework on inventory control and forecasting is presented, providing the necessary foundation for building the model. This framework is then applied in chapter 4, where the theory is used to analyze the data and establish a basis for the modeling stages. The model and its results are then discussed in chapter 5, including an analysis of input parameters, as well as model validation and testing. A critical discussion follows in chapter 6, covering assumptions and validation results. Finally, in chapter 7, the report concludes with remarks and suggestions for future research.

2 Overview of case company

The following chapter provides an overview of the case company and how their inventory control process is set up. It begins with a description of how the safety stock values are set. Then, the concept of service level is introduced and related to the case company. Subsequently, review- and ordering policies at the case company are described. Lastly, since forecasting future demand is an important part of planning inventory, methods to forecast future demand at the case company and measurement of the performance of these are presented.

2.1 Service level and Safety stock

As mentioned in Section 1.3, the case company's ambition is to utilize an inventory optimizer to calculate product-specific safety stocks. However, the optimizer is not in use presently due to concerns about whether it is configured correctly. Instead, as previously noted, the case company employs a fixed value of the safety stock known as "Days of Safety Coverage". This value is based on an average of the expected daily demand quantities and is at Plant X set at 21 days. That is, the safety stock should be able to cover, on average, 21 days worth of demand. At Plant X this value is arbitrarily set rather than being based on data-driven analysis, which results in excessively high safety stock levels. To provide context, Figure 1 illustrates an example of safety stock when using days of coverage compared to the inventory optimizer calculation. Upon analyzing Figure 1, it becomes evident that the safety stock levels, currently employed by the case company, are inflated. Furthermore, in some cases, the inventory optimizer's recommended safety stock may seem very low without further analysis. This discrepancy could stem from the fact that the case company is unsure what input parameters the IO requires, as discussed in the problem formulation in Section 1.3. However, it could also be accurate, indicating that the fixed safety days of coverage is driven by a fear of stockouts when using the lower safety stock values provided by the IO. Ultimately, the use of 21 days as a safety stock threshold could be seen as an exaggerated measure to mitigate the risk of stockouts. This discrepancy highlights the importance of figuring out what the ”optimal” safety stocks are.

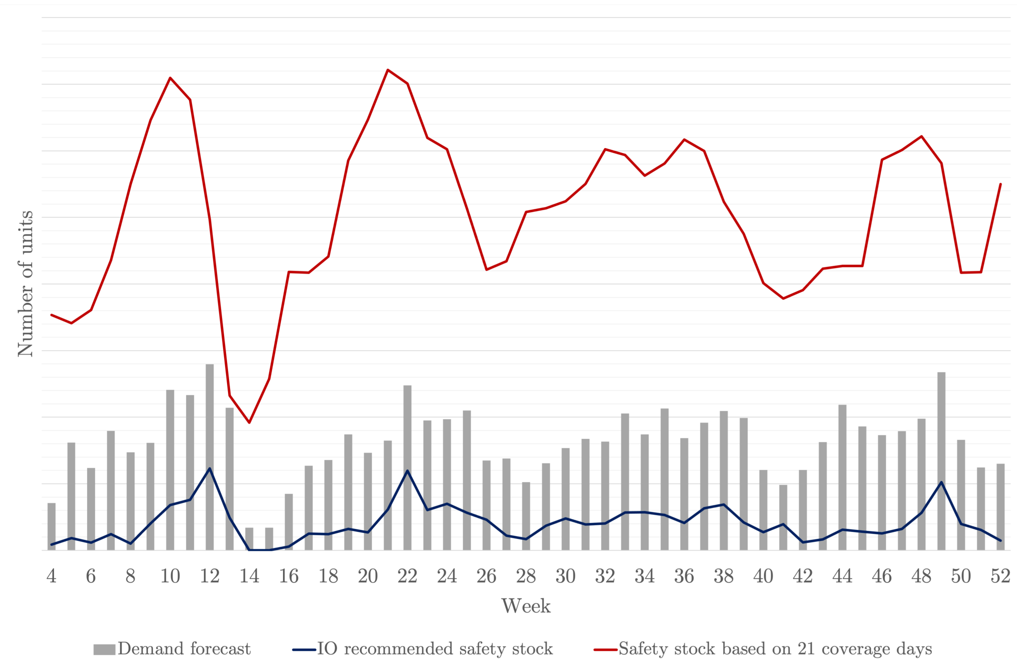

Figure 1: A comparison of the safety stock level provided by the IO and a fixed safety stock level

Safety stock levels are commonly optimized based on a service level constraint. The case company uses a service level known as fill rate, denoted as S_2, which is defined as the amount of demand that can be satisfied immediately from the stock on hand. Axsäter (2006) argues that from a practical point of view, it is important that the same type of service level definition is used throughout the company. Moreover, the overall fill rate goal at the case company is currently set to 98%. This number is agreed with customers and essentially means that the case company needs to achieve an average service level of 98% across all products. However, that doesn’t necessarily mean that every product needs to have the same fill rate requirements since the performance is measured as an average across all products. It could be unwise to have the same service level for all products since the impact that each product has on the profits may significantly vary due to sales volume, price etc. Instead, one could segment the products based on some common characteristic (e.g., with the help of an ABC-analysis). An ABC-analysis is a method to segment products into three categories; A, B, and C based on a predefined metric such as sales volume. In this case, the case company employs an ABC-analysis based on the contribution margin (CM), defined as $Quantity \cdot Margin$, where the quantity refers to the sales volume and the margin refers to the difference between a product's sales price and its variable costs. The segmentation follows Pareto's 80/20 rule, which is commonly applied to ABC-analyses, where 20% of the total volume corresponds to 80% of the total CM, 30% corresponds to 15% of the total CM, and 50% corresponds to 5% of the total CM. Using this rule, the products are categorized as shown in Table 1.

Table 1: An overview of the Pareto rule which the ABC-analysis is based on

Product category	Percentage of total CM
A	80%
B	15%
C	5%

Utilizing the results from the ABC-analysis, the case company has determined the target service level (TSL) that products within each category should meet. The specific target service levels for each product category are outlined in Table 2.

Table 2: An overview of the product categories and corresponding target service levels

Product category	Target service level
A	98.5%
B	97%
C	96%

It is important to clarify that while there is a product segmentation, it has not previously been used to calculate safety stock. This is because the safety stock has been set for a fixed number of days rather than determined by a service level constraint. However, there are plans to experiment with this segmentation in the future.

2.2 Inventory control system

The purpose of this section is to give a brief overview of how the inventory system works at the case company. The different concepts will be further discussed later on in section 3.1. As of now, the SAP IBP system handles a significant part of the ordering process, with supply planners making adjustments to its settings in response to changes in planning or unforeseen factors, such as promotions or production malfunctions. The inventory system at the case company can be represented by an (R, Q)-system, i.e. that an order of a fixed lot size (quantity), Q, is placed at a time when the inventory position is R units. Additionally, the system's review policy is effectively continuous, meaning that the inventory position is reviewed on a continuous basis. As of now, the lot size is seen as the amount of units a shift at Plant X is able to produce. This results in lot sizes that are not consistently equal, so the system data includes a minimum lot size and a maximum lot size.

In section 3.1, the significance of lead time in inventory control is discussed. Some years ago, there was a discussion on interpreting lead times and it was decided to represent the lead time for finished goods based on the cycle time of the entire product portfolio for each production line, i.e., how long it would take to produce the entire product portfolio. The case company is interested in exploring whether this is an effective approach for determining lead time values.

2.3 Forecasting

The case company is undertaking a forecast improvement project aimed at enhancing seasonality predictions using a seasonal model. Seasonality refers to some products exhibiting higher sales volumes during certain periods of the year. The project focuses on applying a statistical forecast method in conjunction with an index that accounts for seasonality variations, which is further described in section 3.3.1. Further details of the project are beyond the scope of this thesis since this thesis is not explicitly about forecasting, rather inventory control. However, the relationship between forecasting and inventory control is evident. An accurate forecast reduces uncertainty, providing an opportunity to lower safety stocks compared to a less accurate forecast.

From an inventory control perspective, forecast errors like the *mean absolute deviation (MAD)*, which directly compares actual demand to the demand forecast, is a commonly used metric. The main reason for keeping safety stocks is to cover for demand uncertainty. Since the demand is forecasted, that demand uncertainty could be expressed as the errors in the forecast. Therefore, the forecast error is used to represent the demand uncertainty in this thesis. Although MAD is used in this thesis when calculating the safety stocks, the case company frequently uses the *mean absolute percentage error (MAPE)*, a metric that contextualizes the difference between sales and demand forecast by comparing it to the size of the demand forecast. The reason for introducing MAPE is that it is hard to determine whether a MAD value represents a high or low forecast accuracy, since it does not take the volume of the demand forecast into account. MAPE is therefore used in this thesis when discussing the performance of the forecast.

3 Theory

This chapter presents the theoretical framework essential for understanding the inventory control principles discussed in this thesis. It begins with an overview of inventory systems. Subsequently, the chapter explores how demand can be modeled with statistical distribution functions, recognizing the critical influence of demand on inventory control decisions. Due to the demand's importance, different demand forecasting methods are presented but more importantly, the concept of forecast accuracy is introduced, which is relevant when determining safety stock levels. Lastly, the chapter addresses the optimization of inventory systems, integrating the aforementioned concepts to create a final approach.

3.1 Inventory control systems

The objective of inventory control is to maintain a balance between the capital invested in inventory and meeting customer demand (Axsäter 1991). To manage this trade-off effectively, an inventory control system is essential. Such a system is guided mainly by two key decisions: when to order and how much to order, i.e., the lot size (Q). It was previously mentioned that there is a minimum and maximum lot size in the system. However, The IO currently utilize the *minimum lot size* as the value of Q. This is currently not adjustable and thus, the focus shifts primarily to the timing of orders, i.e., when to place an order. This decision is based on an *ordering policy* but to fully understand this policy, the concept of inventory position is first introduced. The inventory position can be described as a thorough perspective on inventory levels, since it does not only account for the physical stock on hand but also pending transactions such as outstanding orders and backorders. Axsäter (2006) introduces the following relationship: p

$$\textit{Inventory position} = \textit{stock on hand} + \textit{outstanding orders} - \textit{backorders}$$

where *stock on hand* is the actual number of 6)sunits in stock, *outstanding* *orders* are the orders that have not arrived, and *backorders* is the demand that has been ordered but not yet delivered.

An important policy in inventory control is the review policy. The review policy determines how often the inventory position is reviewed. There are two main types of review policies discussed in the literature, *continuous* and *periodic* (Axsäter 2006). When utilizing a periodic review, the inventory position is reviewed on a regular basis, with a predetermined time between each review. That could for instance mean that one reviews the inventory position once a week and at every review decides whether to place an order or not. As previously discussed in section 2.2, the concept of continuous review is used at the case company, where the inventory position is monitored continuously.

Whilst the review policy is important to determine how often the inventory position is set to be reviewed, the ordering policy is necessary to fully understand when to place the order. The two most common ordering policies discussed in the literature are often denoted *(R,Q) policy* and *(R,S) policy* (Axsäter 2006). In both cases, R is referred to as the reorder point, i.e., a point where once the inventory position falls below, an order is placed. Instead, the difference is how much to order, where S refers to an order up to level. Essentially, that means that once the inventory position falls below the reorder point R, an order is placed so that the inventory position reaches a level S. As previously mentioned, the case company utilizes an (R,Q) policy which will therefore be the basis of the model created in this thesis. As opposed to the order up to level S used in the (R,S) policy, an (R,Q) policy places an order with the fixed lot size Q whenever the inventory position reaches the reorder point. To fully understand how such an ordering system works, figure 2 illustrates how an (R,Q) policy works when applying a continuous review policy.

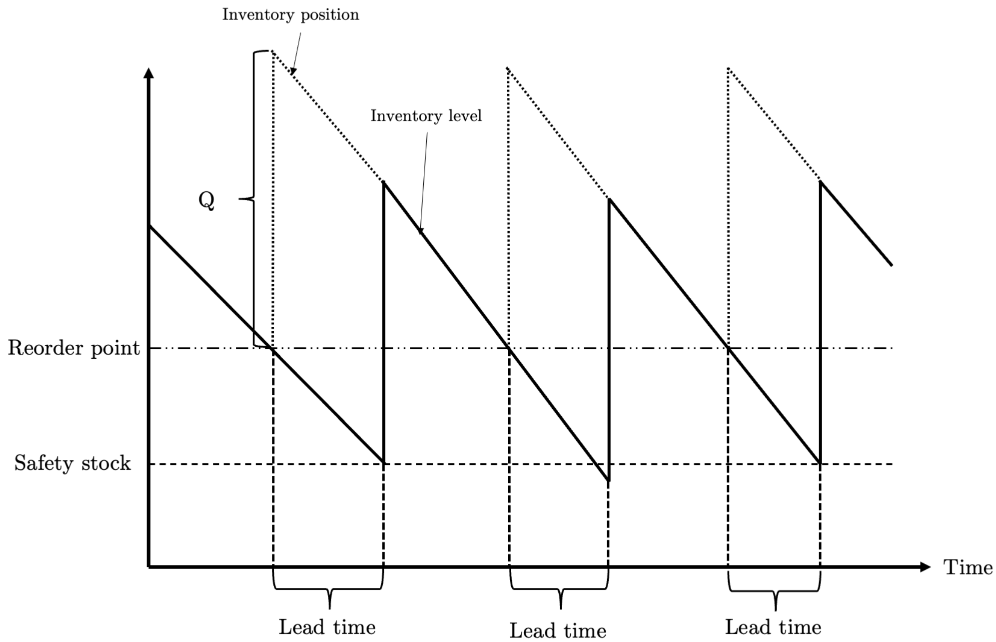

Figure 2: A continuous (R,Q) policy

As illustrated in Figure 2, lead time is the period between placing an order (when the inventory position reaches the reorder point) and its arrival at the warehouse. Once an order is placed, the inventory position increases with the lot size, Q, even though the order has not arrived at the warehouse. Then, the safety stock can be defined as the expected amount of stock on hand when that order arrives. Aside from the target service level, the primary consideration is demand during the lead time when determining safety stocks. Thus, understanding demand is essential for effective inventory management and will be elaborated on in the following section.

3.2 Stochastic demand

As explained in section 3.1, the inventory position does not only consist of stock on hand, but also considers the outstanding orders and backorders. In this context, the outstanding orders are important since their arrival is dependent on a lead time. Since the safety stock is the expected amount of stock on hand just before an order arrives, it is therefore important to consider the demand during this lead time period to prevent running out of stock. This is done by fitting a suitable demand distribution for the quantity demanded during the lead time. Given the stochastic nature of demand, the conventional method involves assuming demand during the lead time adheres to a certain suitable probability distribution (Mattsson 2007a). Although this is not the case for all demand distributions, there are two parameters that describe the demand in this thesis. Those are the mean (μ), which is the average demanded volume, and the standard deviation (σ), which describes the variability of the demand. These are defined with equations (1) and (2), where x_t is the demand quantity during time t and N is the number of data points.

$$\mu = \frac{1}{N}\sum_{t=1}^{N} x_t \tag{1}$$

$$\sigma = \sqrt{\frac{1}{N}\sum_{t=1}^{N}(x_t - \mu)} \tag{2}$$

For instance, if the data is weekly, the mean and standard deviation represent the average quantity demanded per week and the variability of that quantity each week. Since demand during the lead time is what is important in inventory control, these values are adjusted for the lead time and are marked with an apostrophe, as shown in equations (3) and (4).

$$\mu' = \mu(L) = \mu L \tag{3}$$

$$\sigma' = \sigma(L) = \sigma\sqrt{L} \tag{4}$$

Further, sales data is often used to measure the demand in practice since the actual demand is hard to measure. However, it is crucial to acknowledge, as noted by Axsäter (2006), that sales figures and actual demand are not synonymous. The distinction lies in the fact that sales data fails to account for lost sales, which can lead to discrepancies between reported sales and true market demand. Although using sales data as a stand-in for demand often provides a reasonable estimate, potential irregularities should not be overlooked. To mitigate this fact, the service level can be considered before the analysis. A low service level implies a significant volume of lost sales, suggesting that the sales data may not be a fully reliable indicator of demand. However, products from Plant X consistently achieve high service levels, with an average of 99.8% for the past year, and the sales data is therefore used as a proxy for the demand.

3.2.1 Demand distributions & distribution fitting

The importance of fitting a suitable demand distribution has previously been stated. The distribution fitting process is described later in this section. Demand in real-world scenarios is discrete, meaning it almost always is represented by a positive integer. Consequently, employing a discrete demand model is often suitable (Axsäter 2006). However, although demand is discrete in practice, a continuous demand model serves as a practical approximation and is commonly used when demand levels are high which is the case when analyzing the demand of the case company's products. The minor discrepancies resulting from the continuous nature of the model can be considered negligible in this case (Axsäter 2006). Two continuous demand models that are often used in the context of inventory control are the *normal distribution* and the *gamma distribution*. These are briefly introduced in the following paragraph.

When managing products with high demand, a discrete process can be approximated by a continuous normal distribution over an extended time period. With a known mean μ' and standard deviation σ', a normal distribution can be defined. For normally distributes demand, the probability density function (PDF) is shown in equation (5). The cumulative distribution function (CDF) represents the probability of the demand size being equal to or less than x, and is shown in (6). The notation $x(L)$ is used to describe that the functions represent the demand, x, during the lead time, L.

$$f(x)_{x(L)} = \frac{1}{\sigma'\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{x-\mu'}{\sigma'}\right)^2\right] \tag{5}$$

$$F(x)_{x(L)} = \int_{-\infty}^{x} f(x)_{x(L)} dx \tag{6}$$

However, it's important to note that one characteristic of the normal distribution is its allowance for negative demand values, which can result in inaccurate approximations for small demand levels and for products with a high coefficient of variation *(CV)*, which is the ratio between the standard deviation and the mean, σ'/μ' (Axsäter 2006).

When the CV is high, there is a high probability for negative demand when using the normal distribution. In these cases, it might be a safer choice to use the gamma distribution, since it does not allow negative demand sizes. The gamma distribution has the PDF shown in (7), and the CDF shown in (8).

$$g(x)_{x(L)} = \frac{\omega^r x^{r-1} e^{-\omega x}}{\Gamma(r)}, \ x \leq 0 \tag{7}$$

$$G(x)_{x(L)} = \int_{-\infty}^{x} g(x)_{x(L)} dx \tag{8}$$

where $\Gamma(r)$ represents the gamma function (see (9)) and r and ω are defined as in (10) and (11).

$$\Gamma(r) = \int_{0}^{\infty} x^{r-1} e^{-x} dx \tag{9}$$

$$r = \left(\frac{\mu'}{\sigma'}\right)^2 \tag{10}$$

$$\omega = \frac{\mu'}{(\sigma')^2} \tag{11}$$

To offer a graphical comparison of the two continuous functions, distribution curves for the normal and gamma distributions have been plotted with identical means and standard deviations. They are presented in figure 3, where the CV is 0.6, and in Figure 4, where the CV is 0.2. These figures illustrate how the normal probability density function (PDF) permits negative demand values when the CV is high (0.6), whereas the gamma PDF does not. In the case of a CV corresponding to 0.2, the graphs look more similar.

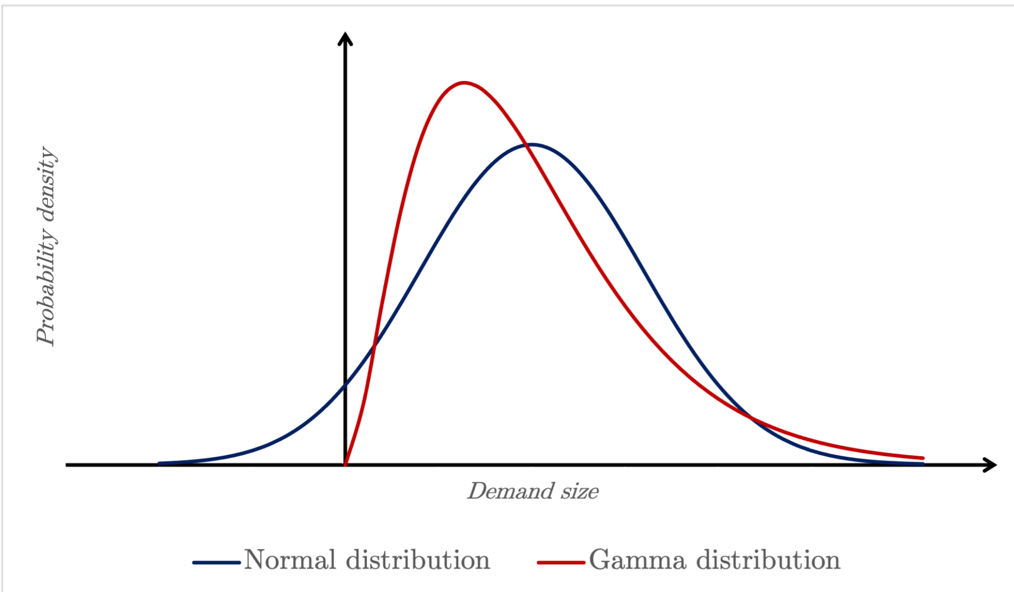

Figure 3: Normal- and gamma distributions with the same μ and σ with a CV of 0.6.

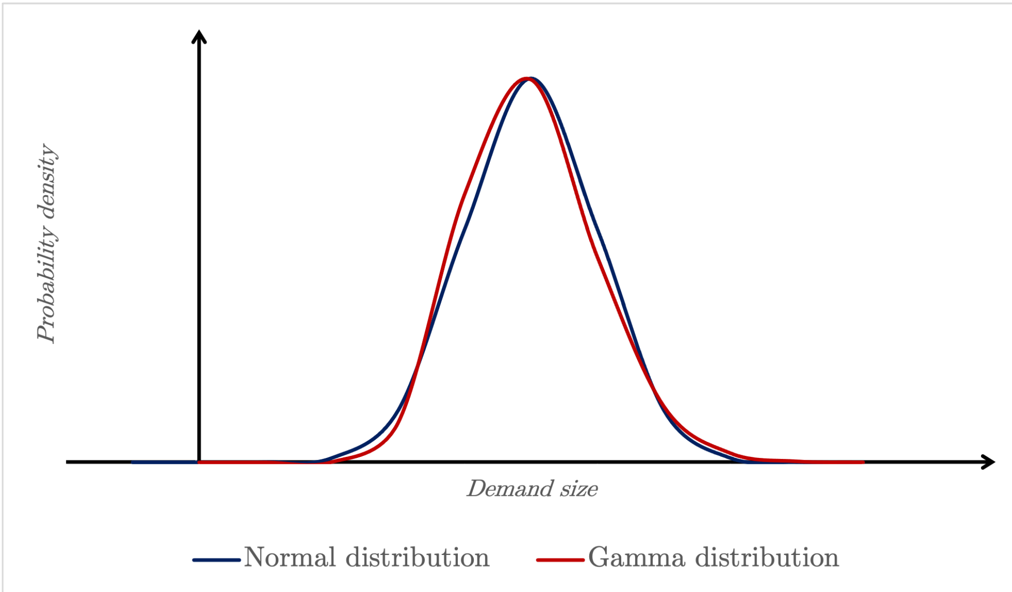

Figure 4: Normal- and gamma distributions with the same μ and σ with a CV of 0.2.

In estimating demand over a specific period, normal and gamma distribution serve as common choices. A frequently used technique in this process is conducting a statistical goodness of fit test (e.g. Kolmogorov-Smirnov test), which is a statistical method that compares historical data to the expected outcomes of a chosen model, e.g. the normal distribution. It is important to highlight that this test should ideally be applied to each Stock Keeping Unit (SKU) individually, though, in reality, applying it to a vast number of SKUs can be impractical. This is due to the fact that a goodness of fit test usually require a software or a more advanced computation and in this case, the case company might want to replicate the analysis without relying on a computer based program. Alternatively, a simpler method can be employed to ascertain the most fitting distribution for a dataset.

In the context of continuous demand, distribution fitting is widely discussed in the literature. To investigate this, a literature review was conducted using databases like Scopus, EBSCOhost, and Google Scholar, all accessible through Lund University. Keywords such as *Demand distributions*, *Distribu-*

tion fitting, and *Continuous demand* were used to identify relevant articles.

As a result of the literature review, it was concluded that the CV during the lead time, serves as a benchmark for assessing the feasibility of approximating demand with a normal distribution. H. S. Lau and H. L. Lau (2003) and Tyworth and O'Neill (1997) suggest that a CV below 0.5 during the lead time indicates suitability for normal distribution approximation. However, it is crucial to note that higher CV values increase the risk of encountering negative demand figures. For instance, a CV of 0.5 corresponds to a 2.28% risk of negative demand. H. S. Lau and H. L. Lau (2003) further discusses that even relatively small CV values (e.g., 0.3) may lead to noticeable errors, with a corresponding 0.04% risk of negative demand. Snyder (1984) argue that to minimize the risk of negative demand to a negligible level, the CV value should not exceed 0.2 for the normal distribution to remain appropriate in demand approximation. When the CV is 0.2, the probability for negative demand is $2.87 \cdot 10^{-5}$%. In this thesis, the strict CV threshold for normal distribution was applied to the model, due to it leading to a negligible probability for negative demand.

3.3 Forecasting

Although this thesis does not directly focus on forecasting, the significance of the concept should not be underestimated. Safety stocks aim to account for demand uncertainties, which can be quantified using forecast errors. An accurate forecast closely matching actual sales volumes enables lower safety stock levels. Forecast errors thus represent demand variability within the model. Therefore, forecasting methods are briefly introduced, followed by a description of how forecast errors are measured.

3.3.1 Forecasting methods and demand models

The predominant forecasting method involves analyzing historical data under the assumption that future demand will mirror past patterns to some extent (Axsäter 2006; Olhager 2013). Statistical methods are then employed to model this future demand, accommodating factors like trends and seasonality. However, many inventory models default to the simpler *constant model*, which posits that future demand will deviate randomly around a stable mean, x_t (Axsäter 2006). In the constant model, $\hat{x}_{t,l}$ represents the demand forecast made for a given period t, made in period l and ε_t represents the independent stochastic deviation. Consequently, the forecast for period t can be represented by (12). The model introduces the concept of lag, denoted by l, which is the time difference between making a forecast and the actual event. That is, a lag of l means that the forecast is made l time units prior to the sale that is forecasted.

$$\hat{x}_{t,l,constant} = x_t + \varepsilon_t \tag{12}$$

The forecast improvement project that the case company has initiated, presented in section 2.3, applies a model that is also found in the literature, where Olhager (2013) expands the constant model by adding a seasonal index for period t, F_t.

The seasonal index for an average period within a year, such as a month or quarter, is denoted as $F_t = 1.0$. If the index is, for instance, $F_t = 1.2$, it indicates an expected demand increase of 20% compared to the average for period t due to seasonal factors. Therefore, across T periods in a year, the cumulative sum of all seasonal indices equates to T. The demand for period t can then be represented by (13), where the lag remains denoted as l.

$$\hat{x}_{t,l,seasonal} = x_t + F_t t + \varepsilon_t \tag{13}$$

The concept of seasonal index is used to asses the seasonality of the products prior to choosing a suitable time period for which the safety stocks are set (see section 4.1.2), which makes the concept worth presenting.

3.3.2 Forecast errors

In the context of inventory control, estimating the mean of the future demand is not sufficient. As discussed, uncertainties are a big factor when dealing with inventory control at the case company and to be able to determine a suitable safety stock, it becomes necessary to know the extent of these uncertainties, i.e., how large the forecast errors tend to be (Axsäter 2006). The various methods for calculating forecast errors can be divided into two categories; *absolute forecast error* and *relative forecast error*. An absolute forecast error involves calculations based on the absolute error, e_t, which is defined as the difference between the demand for that period and the forecast for that period (Olhager 2013; Shcherbakov et al. 2013), see (14):

$$e_t = x_t - \hat{x}_{t,l} \tag{14}$$

Further, the most common method in inventory control is to quantify variation around a mean, by using standard deviation, σ (see equation (2)). A common assumption is that the forecast errors are normally distributed, which leads to the standard deviation being estimated using (15) (Axsäter 2006).

$$\sigma = \sqrt{\frac{\pi}{2}} \cdot MAD \approx 1.25 \cdot MAD \tag{15}$$

Prior to the advancements in information technology and computational capabilities, the *mean absolute deviation (MAD)* was commonly used to represent variability and is calculated with (16), where n represents the forecast horizon. Despite the capability to estimate σ with modern technology, many forecasting systems continue to estimate the standard deviation as per (15) (Axsäter 2006; Olhager 2013; Shcherbakov et al. 2013), including the IO at the case company.

$$MAD = \frac{1}{n} \sum_{t=1}^{n} |e_t| \tag{16}$$

Meanwhile, the relative forecast error, p_t, is shown in equation (17), which is essentially the ratio of the absolute error $|e_t|$ to the actual demand x_t (Shcherbakov et al. 2013). This measure contextualizes the error in relation to the forecasted quantity, providing a useful metric for evaluating forecast performance. However, many inventory models rely on the standard deviation to measure variation, which makes relative error metrics less practical for direct incorporation into these models.

$$p_t = \frac{|e_t|}{x_t} \tag{17}$$

Despite its name, the *Mean Absolute Percentage Error (MAPE)* is classified as a relative forecast error in this context as it is a percentage value that is compared to the demand forecast size. It stands out as the most widely used measure (Fildes and Goodwin 2007; McCarthy et al. 2006), including at the case company. MAPE computes the average of the absolute percentages of p_t capturing the error's proportion relative to actual demand, as indicated by (18).

$$MAPE = \frac{t = 1}{n} \sum_{t}^{n} 100 \cdot |p_t| \tag{18}$$

3.4 Optimization of continuous review (R,Q)-policy

As mentioned, the core aim of inventory control is to balance the service level achieved against the costs of maintaining inventory, where of safety stocks serve to maintain a competitive delivery capability (Mattsson 2007b). Hence, the optimization of these safety stock levels is typically based on some predefined service level constraint set by the company (Axsäter 2006).

3.4.1 Inventory level distributions and service level constraints

Considering an (R, Q) ordering system, an order is triggered as soon as the inventory position (IP) is less than or equal to the reorder point R with the lot size, Q. Knowing that, the relationship between the reorder point, R, and the safety stock SS is shown in (19) (Axsäter 2006). It is noteworthy that this is an approximation that relies on that the reorder point is always exactly hit in practice.

$$SS = R - \mu' \tag{19}$$

Normally distributed demand

When considering continuous and normally distributed demand, the inventory level's CDF, for all values of x, can be found using (20), where IL represents the inventory level, and $P(IL \leq x)$ is the probability that the inventory level falls below x (Axsäter 2006).

$$P(IL \leq x) = \frac{\sigma'}{Q}\left[G\left(\frac{R - x - \mu'}{\sigma'}\right) - G\left(\frac{R + Q - x - \mu'}{\sigma'}\right)\right] \tag{20}$$

where G is the loss function, obtained by the density and cumulative distribution functions, φ and Φ, of the standard normal distribution (see (21).

$$G(x) = \varphi(x) - x(1 - \Phi(x)) \tag{21}$$

The probability of not being able to satisfy the demand with the inventory on hand (S_2) can then be expressed as $1 - P(IL \leq 0)$, as seen in (22) (Axsäter 2006). Equation (22) can then be used to obtain the reorder point, and the safety stock is then computed with equation (19).

$$S_2 = 1 - P(IL \leq 0) = 1 - \frac{\sigma'}{Q}\left[G\left(\frac{R - \mu'}{\sigma'}\right) - G\left(\frac{R + Q - \mu'}{\sigma'}\right)\right] \tag{22}$$

Axsäter (2006) also notes that equation (22) is often approximated using equation (23), as this approximation generally performs well for large values of Q. Therefore, equation (23) is useful for manual calculations.

$$S_2 \approx 1 - \frac{\sigma'}{Q}\left[G\left(\frac{R - \mu'}{\sigma'}\right)\right] \tag{23}$$

Gamma distributed demand

For products with gamma distributed demands, calculating safety stocks involves similar principles but with specific methods for assessing the probability of demand being higher than the reorder point (Burgin and Norman

1976; Keaton 1995). In the literature, safety stocks are commonly calculated with a service level defined as cycle service level, S_1. However, there are fewer examples of methods that calculate the fill rate, S_2. Still, a method that utilizes the fill rate is discussed by Mattsson (2007a) and is presented as an iterative method, employing the gamma density function, $g(x)$, to estimate expected backorders, $E[IL^-]$ for a given demand size, x, with the reorder point, R, shown in equation (24).

$$E[IL^-] = \sum_{x/x \leq R} (x - R)(1 - g(x)) \tag{24}$$

Mattsson (2007a) describes the *allowed backorder quantity*, $E[IL^-]_{S_2}$, as the number of backorders that would lead to a specific fill rate, where the index S_2 represents that the backorders are calculated based on the target fill rate. This is seen in equation (25), where Q represents the lot size.

$$E[IL^-]_{S_2} = Q(1 - S_2) \tag{25}$$

To determine the optimal reorder point, iterations are conducted over various possible reorder points using equation (26). This process identifies the most suitable reorder point. Then, equation (19) is applied to compute the safety stock. The notation $E[IL]^-$ denotes the function that calculates backorders for each reorder point R. The calculated reorder point represents the reorder point where the expected backorders come closest to the allowed backorders.

$$min \left| E[IL^-](R) - E[IL^-]_{S_2} \right| \tag{26}$$

4 Analysis

This chapter outlines the analysis process, detailing the groundwork for customizing the model based on the distinct characteristics of the available data.

4.1 Data cleaning

One part of the first step in operations research involves gathering relevant data (Hillier and Lieberman 2010). This step was performed during the data analysis. The initial phase of the data analysis involves filtering and cleaning the company data to align with the project's objectives and boundaries. This step starts with gathering data from the case company's planning system and the data that was provided for this thesis dates back to early 2021. The initial dataset comprised 68 products. Post-cleaning, 36 products remained, with the reduction primarily due to the exclusion of discontinued items, i.e., items that are no longer included in the product portfolio. Additionally, three recently introduced products, having been on the market for about 16 weeks, were identified. Given their short history, these items presented unique challenges for the analysis, particularly in calculating historical forecast errors and means. Incorporating them would result in a skewed analysis due to the lack of data points (refer to section 4.3). Due to this, these products were excluded from the analysis.

After identifying the relevant product set, an examination of the data unveiled certain inconsistencies, notably weeks missing forecast data points – meaning that some weeks where represented by a zero for all products. If these weeks would have been included, the value of the forecast error would have been misrepresentative. Approximately four weeks per year were affected by this issue. This could be seen as a systematic error since all products had the value of zero in the forecast for the same weeks. To maintain data integrity and reliability, these periods were excluded from both the demand and forecast datasets.

4.1.1 Demand distributions

The planning system inputs suggest an assumption that all product demands follow a normal distribution. However, it is not certain that all the products can be assumed to be normally distributed. Moreover, it is also not certain that all products can be approximated by a continuous distribution. To determine if that is the case, a distribution fitting analysis was performed.

Working with continuous distributions is computationally efficient and building an inventory model that assume discrete demand distributions is not in the scope of this thesis. Therefore, the analysis began by understanding whether the products' demand distributions could be closely approximated by a continuous distribution. This assessment involved examining the demanded quantities for each product and their CV. The findings indicated that a continuous distribution model could represent all products' demands, barring two exceptions. This result is shown in figure 5, which shows the CV for each of the 36 products. This decision is based on the criteria that Axsäter (2006) presents, where a demand can be represented by a continuous distribution if the demand during lead time consistently exceeds 10 units and the CV stays below 0.5. In this case, the demand of all products consistently exceeds 10 units during lead time. However, the two products above the continuous CV threshold were removed from the analysis.

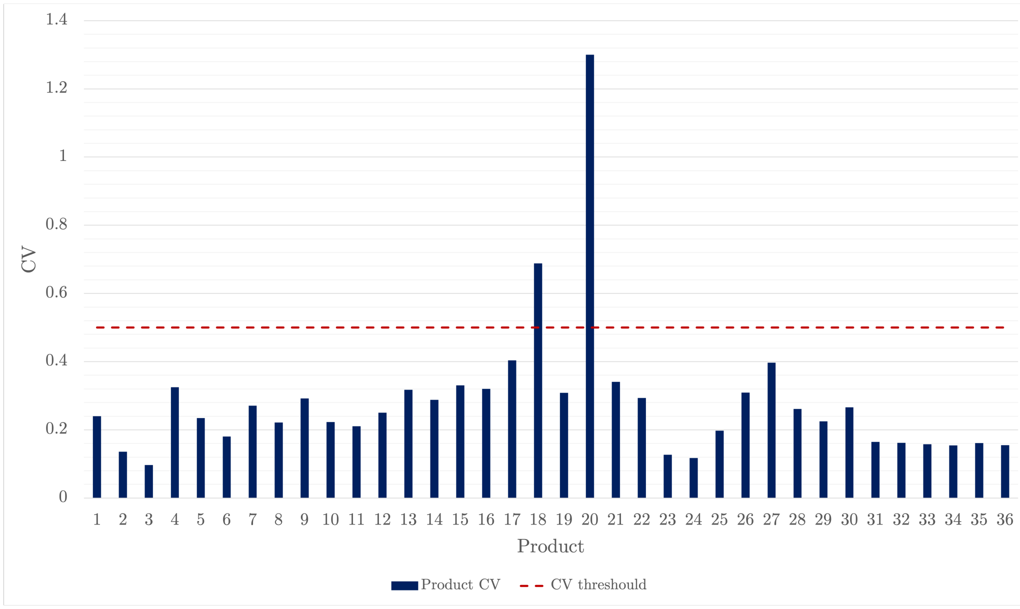

Figure 5: Each product's CV plotted with the threshold for being able to assume a continuous demand distribution.

Subsequently, the decision between using a gamma or normal distribution for modeling demand needed to be addressed. To simplify the process and enhance its replicability across different warehouses and product lines in the future, the analysis opted to use the CV as a guiding metric rather than performing a Kolmogorov-Smirnov test for each SKU. With the strict guideline of the maximum CV of 0.2 provided by Snyder (1984), which carries a risk of $2.87 \cdot 10^{-5}\%$ chance of negative demand, the analysis showed that 12 products could be approximated using a normal distribution while the remaining 22 should be represented by a gamma distribution. The results of this analysis, after the exclusion of the two non-continuous products, are shown in figure 6, where each products' CV is shown in relation to the 0.2-threshold.

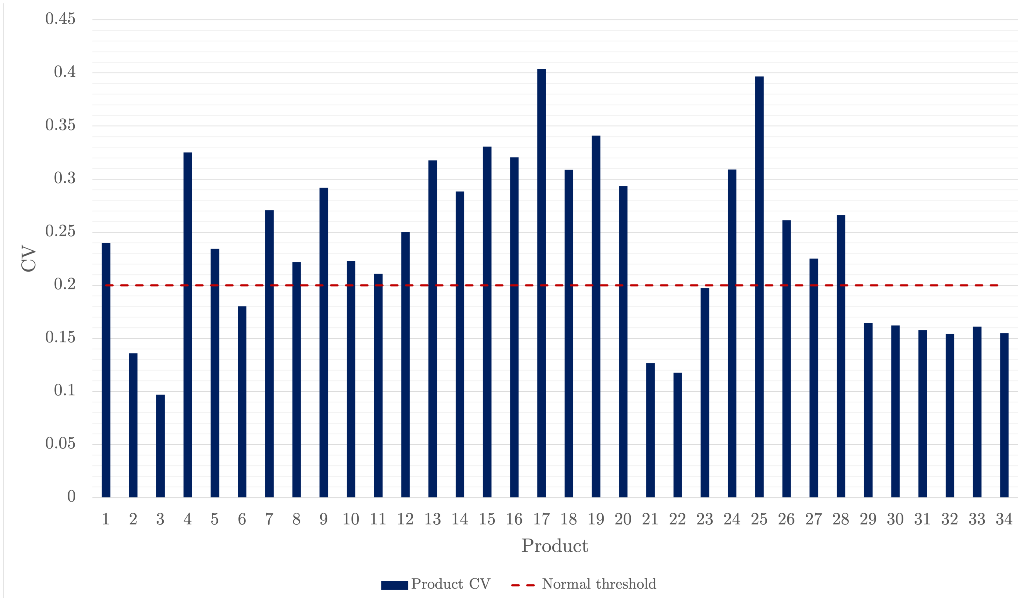

Figure 6: Each remaining product's CV plotted with the threshold for being able to utilize the normal distribution.

4.1.2 Considering seasonality in demand

When assessing seasonality quantitatively, data availability often presents challenges. In this case, only two complete years of sales data are available, making it difficult to conduct a robust seasonality analysis with statistical significance. Given that each month across two years provides only two data points per product, the analysis is limited in its ability to draw definitive conclusions about seasonal patterns.

However, discussions with supply planners and others familiar with the products at Plant X have identified that only product 28 is considered seasonal. This product is closely associated with a Swedish holiday in late summer, during which it experiences significant increases in sales volumes. To illustrate this seasonal effect, the seasonal model described in section 3.3.1 was used to provide a graphical overview of how sales volumes for this product peak during the holiday month. Figure 7 shows the seasonal index, F_t, for this product alongside several non-seasonal products, highlighting the distinct seasonal pattern compared to the more uniform indices of the non-seasonal items. For example, as seen in figure 7, the seasonal index for product 23 shows that its sales are approximately 20% higher than its average in September indicating possible seasonality. However, due to the data availability, this cannot be confirmed. Furthermore, discussions with the people that have extensive experience with these products, confirmed that this product is not considered seasonal.

Given the assumption of stationary demand patterns in this thesis, this specifically seasonal product was excluded from the broader analysis. This exclusion ensures that the results for the remaining products, presumed non-seasonal, are not skewed by the marked seasonality of one outlier. This approach suggests that the limited dataset does not significantly impact the analysis outcomes for the non-seasonal products.

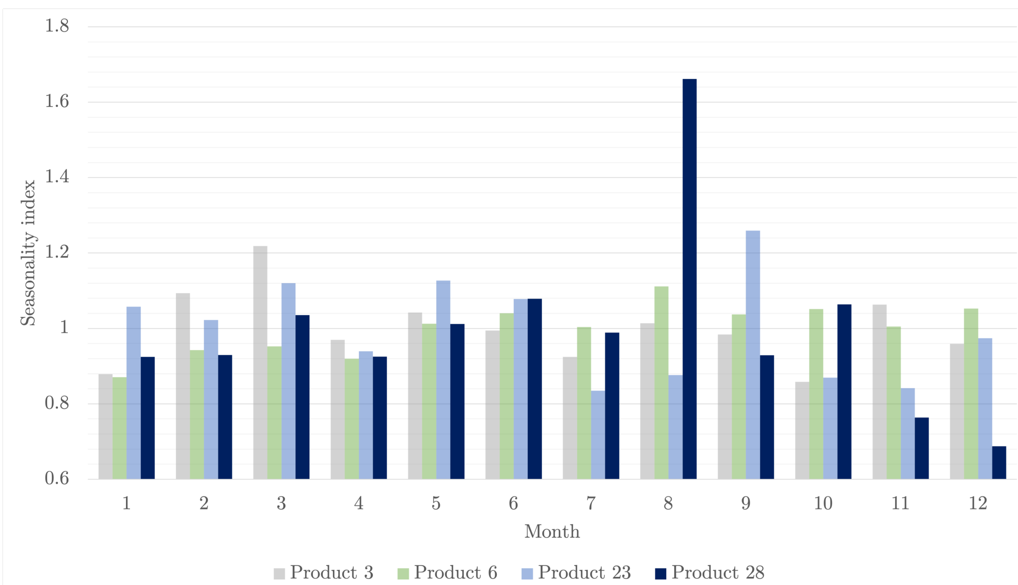

Figure 7: A comparison of the seasonality index, F_t, of four different products.

4.2 Lead times

As mentioned, the lead times were previously interpreted as the cycle time for the entire product portfolio, i.e., how long it takes to produce all products. However, this classification relied predominantly on general experience from the production plant. Moreover, the soundness of interpreting the lead times as a cycle time were questioned at the case company. Therefore, regarding lead time, the analysis is divided into two steps. First, it determines how lead times should be interpreted to most accurately reflect reality. Second, it involves extracting and analyzing data to establish values based on this interpretation.

The clarification of lead times and their conceptual interpretation entailed informal discussions with representatives from the case company. While the conventional term for gathering qualitative data is *interview*, in this study, the term *informal discussions* is used. This distinction arises from the absence of structured interviews with predetermined question lists for data collection (Höst et al. 2006). In this case, the discussions were particularly held with individuals overseeing operations at the targeted production plant connected with the investigated warehouse. Initially, the focus was on defining lead time, particularly in the context of calculating safety stocks. However, the presence of multiple products on the same production line posed a challenge for the case company in establishing predefined lead times for each product, as simultaneous production of all products was impractical. Furthermore, considering the short production and transportation duration, wherein specific products were produced within a few hours and transported to the finished goods warehouse the following day, the discussions shifted towards determining what other factors that have an impact on the overall lead time.

As discussed, lead times had previously been conceptualized with a cyclic mindset, executed at a higher level rather than being product-specific. In this scenario, the determination of lead time for each individual product emerged as a primary focus, given that safety stock calculations were intended for every product. Through discussions with the case company, it was concluded that interpreting lead time as the cycle time for each product was a logical approach. That is, instead of using the cycle time for the whole product portfolio, each product's individual cycle time would be used. Consequently, it was mutually agreed that the lead time used in safety stock calculations should be based on the average time between a product's production occasions. The main reason for this decision is the fact that if the lead time was interpreted only as the actual time it takes to produce a batch, store it at the plant, and transport that batch to the warehouse, the *queue time* will not be accounted for. The queue time represents the time it takes for a product to wait to be produced and in this case, the queue time represents the majority of the total lead time, which is showed in figure 8.

In this particular case, there are several factors that influence the queue time. Firstly, the production adheres to a production plan which is fixed for the following week. This production plan shows what product that is produced every day. Furthermore, due to strict hygiene requirements and therefore long changeover times (the time it takes to switch production from one product to another), only one product per production line is produced each day. That means that the entire cycle, including the queue time, becomes the time that needs to be accounted for when making decisions about how much cycle and safety stock are needed. In practice, lead times may vary depending on the queue time but in this thesis, constant lead times are assumed.

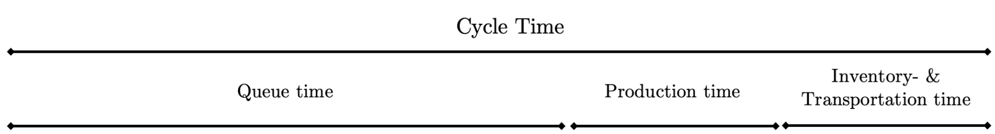

Figure 8: A representation of how the cycle time can be divided into Queue-, production, and inventory & transportation time.

Based on the aforementioned discussion, an approach to calculate new lead time values of the respective products was formed. In this case, data from the planning system was collected for the analysis. The collected data represented the date for every order that was produced for each product. However, the data that was collected needed cleaning. This was mainly due to to fact that there was a ramp up in production frequency from 2021 until the summer of 2023 because a reconstruction of Plant X was scheduled and the plant would therefore be closed for a longer time period during the summer 2023. To cover for this break in production, excessive stock was needed and hence, the ramp up in production. Therefore, it was decided that the production data from this period was excluded from the analysis to ensure that the data set provided the best possible representation of reality. The data that was used is therefore from week 32 in 2023. Also, the calculation took into account the fact that the production plant was closed for the last three weeks of that particular year.

After data cleaning procedures, the mean duration between production cycles of individual products was computed. The outcomes, along with a comparison against prevailing values, are presented in table 3. When analyzing the values, it becomes apparent that certain products display notable disparities between the lead time values calculated using the new approach and those that are currently set in the system. One reason for this could be the fact that there is a considerable variability in the number of data points across different products in this analysis. This is due to the fact that some products are produced less frequently than others and thus had a limited number of data points within the period since the reconstruction. As a result, their average cycle time might yield less reliable estimates. This source of error is further discussed in chapter 6.

Notably, it was previously discussed in section 2.2 that the lead times were set based on the time it would take to produce the whole product portfolio. However, as seen in table 3, the values of lead times currently set in the system are not uniform and some values for L_2 are larger than for L_1, which does not represent the prior approach. This is due to an unknown system-related technical issue but since these are the values that the IO uses, these become the benchmark.

Table 3: An overview of the current values of the lead times L_1 (weeks) that are in the system compared to the values calculated with the new approach L_2 (weeks)

Product	L_1	L_2	Product	L_1	L_2
1	3	4.9	**18**	3	4.3
2	1.5	6.7	**19**	10	5.1
3	8	4.4	**20**	10	4.5
4	5	2.2	**21**	8	2.6
5	8	6.3	**22**	5	6.5
6	8	1.2	**23**	5	6.2
7	5	3.3	**24**	5	4.0
8	5	4.5	**25**	3	6.2
9	5	4.1	**26**	5	3.3
10	5	2.2	**27**	5	5.3
11	5	2.0	**28**	8	9.5
12	5	9.7	**29**	8	5.7
13	5	2.1	**30**	10	1.8
14	8	2.6	**31**	8	5.3
15	5	1.9	**32**	10	2.1
16	5	1.4	**33**	8	3.4
17	5	0.9			

4.3 Demand variability

Demand variability and uncertainty is the primary reason for maintaining safety stocks, making it a crucial parameter in the presented model. To determine the variability of the demand, there are two common approaches:

1. Using forecast errors to represent variability, with calculations based on equations (15) and (16).

2. Calculating demand variability directly from the historical demand data.

The case company has decided that the first approach will be utilized when implementing the SAP Inventory Optimizer. The reason for this is that the inventory plan ahead is based on a demand forecast, and hence, the forecast error gives an indication to the variability that the safety stock need to account for. A possible downside to this is that a low forecast accuracy naturally leads to higher forecast error and ultimately, higher safety stock levels. In these cases, the second approach can become more useful, due to it not being reliant on a high forecast accuracy. However, a downside with using historical standard deviation to represent demand variability is that it includes demand peaks that may have been accounted for by the forecast, potentially leading to a higher variability than necessary.

In relation to the analyzed products, the forecast error method is found to generate, on average, a 19% higher variability than the standard deviation obtained by only considering past demand data. The product specific comparison is shown in table 4. This concludes that, in order to accurately use the forecast error, as intended at the case company, the forecast needs to be improved. To provide an overview of how this decision affects the final safety stock level, a comparison is presented in section 5.6.1.

Table 4: A comparison between using the forecast error, σ_{MAD}, and using the standard deviation of historical demand, σ_{demand}.

Product	σ_{MAD}	σ_{demand}	**Product**	σ_{MAD}	σ_{demand}
1	115	130	**18**	94	78
2	71	57	**19**	40	46
3	91	78	**20**	69	55
4	752	740	**21**	181	147
5	41	30	**22**	127	90
6	497	381	**23**	102	58
7	723	607	**24**	706	646
8	459	385	**25**	475	337
9	355	289	**26**	609	454
10	1038	826	**27**	285	291
11	856	748	**28**	324	299
12	229	179	**29**	480	330
13	450	397	**30**	387	275
14	959	875	**31**	152	137
15	1400	1216	**32**	318	264
16	1684	1424	**33**	172	154
17	1915	1570			

4.4 Mean demand

Although the purpose of safety stocks is to cover for demand variability, the demand volume during the lead time is an integral part when calculating reorder points and should therefore be discussed. In this context, two approaches can be utilized. The first approach is to base the mean demand during the lead time on historical data, essentially operating without considering forecast data. This approach may be adopted for reasons such as low forecast accuracy or the absence of forecasts. For this method to be relevant, historical demand must be stationary and not expected to change in the future. However, demand volumes have dropped during the past two years. This is due to the food industry being sensitive to macro-economic factors such as rising inflation rates. This indicates that using historical data alone can lead to inaccuracies. Specifically, if demand decreases in the future, reorder points calculated on past data may be too high. For example when considering the drop in demand volumes during the last year, a reorder point set for the year of 2023 will result in reorder points that are unnecessarily high if this sales volume drop is not anticipated by the forecast. Conversely, if demand increases they might be too low.

Given that the average MAPE has been relatively low at 13.1% over the past year at Plant X and considering the recent decrease in demand volumes. Calculating the mean of the forecasted demand instead of using historical data was therefore considered a reasonable approach.

In this thesis, the reorder point will remain fixed for a certain time period ahead. Considering the fact that the model assumes stationary demand, the demand during the fixed time period must be stationary. This period will be referred to as the *span* for the reorder point calculation. Several choices can be made regarding this span. From a demand perspective, the reorder point relies on two primary demand parameters: variability and volume. While the safety stock is influenced solely by demand variability, the reorder point also depends on the mean demand during the lead time (μ'), as outlined in equation (19). The selection of span affects these reorder points because the mean demand might change when the span changes and thus, the span becomes an important consideration. To decide a suitable span, it is important to determine if there is any nonstationarity during this span. Since it has been established that there are no seasonal products in the analysis, the reorder points set are set for the remainder of the year of 2024 in this thesis. However, the choice of span remains important even without seasonal products, since it becomes harder to be accurate in forecasting as the span increases.

5 Implementation and results

This chapter firstly presents the final model. It begins with a description of how the model is created based on the previous chapters, which is followed by a description of how the model works. Thirdly, the results from the model calculations are presented along with results from testing the performance of the model. Lastly, results on how the different values of the input parameters affect the final safety stock is presented.

5.1 The model

According to Hillier and Lieberman (2010), the second step in operations research involves formulating a mathematical model that represents the problem at hand. In this case, the mathematical models that represents the problem are (23) for normally distributed products, and (24)–(26) for gamma distributed products.

The third step in operations research involves developing a computer-based procedure to derive solutions from the model (Hillier and Lieberman 2010) The final model is presented as an Microsoft Excel-file consisting of multiple spreadsheets. In this case, the model differentiates between products following a gamma distribution and those following a normal distribution. The primary interface sheet displays a list of products with their respective names and numbers, allowing manual inputs such as product ID and target service level adjustments. Appendix A offers an overview of the interface sheet. Another sheet contains relevant data for each product necessary for safety stock calculations, including:

1. Lead time
2. Lot size
3. Mean forecasted demand

4. Forecast error
5. Statistical distribution

Upon entering a product number, all relevant parameter values are fetched from the data sheet, presenting them in a user-friendly table format (see figure 9).

Product number	[redacted]
Target service level	98.0%
Product name	
Production lead time (weeks)	**0.91**
Lot size	**3216**
Mean demand during lead time	**3718.36**
Standard deviation during lead time	**1827.75**
Statistical distribution	**Gamma**

Figure 9: Relevant product data shown to the user

As seen in the figure, the lead time and the lot size for the product is presented. The mean demand and the forecast error during lead time is automatically calculated based on the product data. Lastly, the statistical distribution is presented. This is based on the CV mentioned above and has the threshold of 0.2 for normal distribution. Based on these values and whether the product is normally or gamma distributed, the safety stock is calculated in number of units and then translated to weeks and days using the average weekly demand volumes. The results are presented to the user as seen in figure 10.

Safety stock (pcs)	**2266.566 pcs**
Safety stock (weeks)	**0.556 weeks**
Safety stock (days)	**3.889 days**

Reorder point (pcs)	**5984.924 pcs**
Reorder point (weeks)	**1.467 weeks**
Reorder point (days)	**10.269 days**

Figure 10: An example of how the results are presented to the user

5.2 New safety stocks with a TSL of 98%

The model implemented recalculated fixed safety stock levels for all 33 products for the remainder of 2024 based on a TSL of 98%. There are a few reasons to why the span is set as the remainder of 2024 and why the safety stocks are fixed during that time period. Firstly, a longer forecast period tends to decrease accuracy, contributing to potential errors. Additionally, the forecasts for 2025 were not sufficiently updated, prompting a decision to confine the analysis within 2024. Setting fixed safety stock levels assumes a stationary system, meaning that the demand's mean and standard deviation are fixed for the entire period. Additionally, it has been concluded that the influence of seasonality is minimal, apart from one product which was removed from the analysis. Therefore, it was decided that setting fixed safety stocks for the remainder of 2024 was reasonable without significant risk of error. The results show that the average days of safety stock can be reduced by 47% compared to the current safety stock level of 21 days, as illustrated in Figure 11. In terms of units, this resembles a reduction from approximately 101 000 units to 37 000 units, a decrease of 64 000 units.

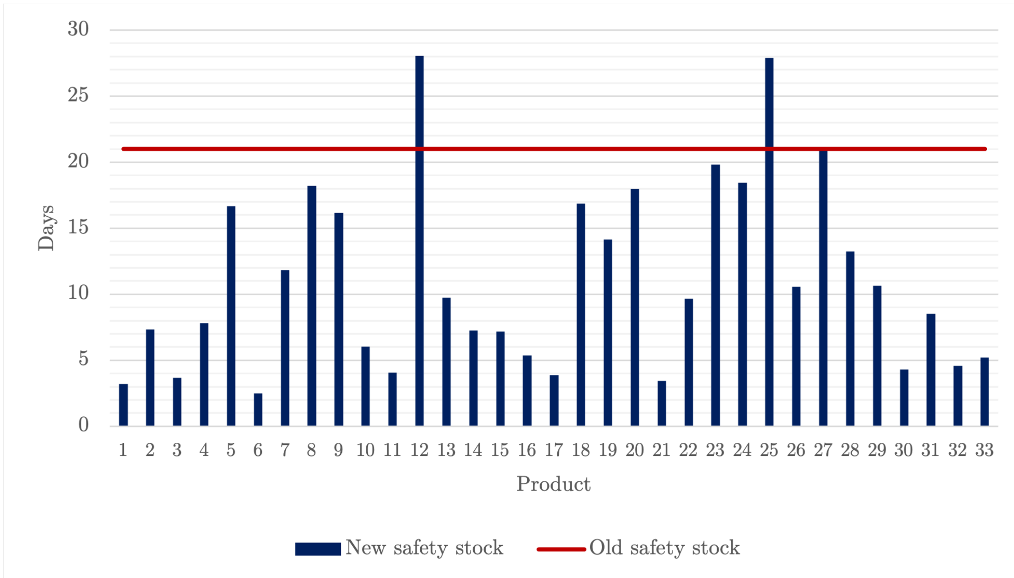

Figure 11: New safety stock levels compared to the old safety stock levels.

The results show significant difference in safety stock levels across different products, primarily due to differences in lead times and forecast accuracy. Products with longer lead times typically require more safety stocks, while those with shorter lead times need less. Additionally, products with higher forecast accuracy, which are easier to predict, necessitate lower safety stocks. Conversely, those with lower forecast accuracy require higher safety stocks to compensate for the increased uncertainty in their demand forecasting.

Examining Figure 11, two outliers, product 12 and product 25, are highlighted. For these products, a reduction in safety stock is not feasible; instead, their safety stocks need to be increased from the current level of 21 days. Product 12, as shown in table 3, has the highest lead time among the analyzed products at 9.7 weeks. While this contributes to its high safety stock levels, a comparison with product 28, which has a nearly equivalent lead time of 9.5 weeks yet significantly lower safety stock, suggests that lead time alone does not drive the need for higher safety stocks. The primary distinction lies in their forecast accuracy. Product 12 has a MAPE of 42%, significantly higher than product 28's MAPE of 20%.

Additionally, consider product 25, which, although not having the longest lead time at 6.2 weeks, has a safety stock level comparable to that of product 12. This similarity is attributed to its low forecast accuracy, with a MAPE of 46% over the past year. In contrast, products 6 and 21, which have much lower safety stocks, have impressive MAPE values of 17% and 15%, respectively. This pattern highlights the important role of forecast accuracy in determining safety stock levels, as safety stocks primarily serve to buffer against demand uncertainties stemming from forecast errors.

5.2.1 Validating the model

The operations research framework outlined by Hillier and Lieberman (2010) is integral to this thesis, specifically the fourth step which entails testing the model using historical data to ensure it adequately addresses the intended problem. This validation step is crucial for the case company because the results suggest a significant reduction in safety stocks compared to prior levels, potentially leading to skepticism or hesitation if not properly validated.

To validate the model, three six-month periods were examined: the second half of 2022, the first half of 2023, and the second half of 2023. The validation required specific historical data, particularly the forecast data for the subsequent period, which was extracted from SAP. Since the safety stock calculation is forecast-dependent, a new safety stock calculation was conducted for each period, with the same TSL of 98% for all products, but with a different forecast for each testing period. This means that the forecast data is different for each period, leading to a different mean demand during lead time, and therefore a new reorder point. Additionally, since the forecast error calculation is based on one year prior to the beginning of the period, these will also differ across the periods, impacting the safety stocks. The performance was then assessed using actual sales data to determine how the recalculated safety stocks would have influenced service levels over these six-month spans. The tests yielded average achieved fill rates of 98.95%, 98.17%, and 98.67% for the three respective six-month periods, with product-specific results detailed in figures 12, 13, and 14.

Figure 12: The results of testing the model on the second half of 2022.

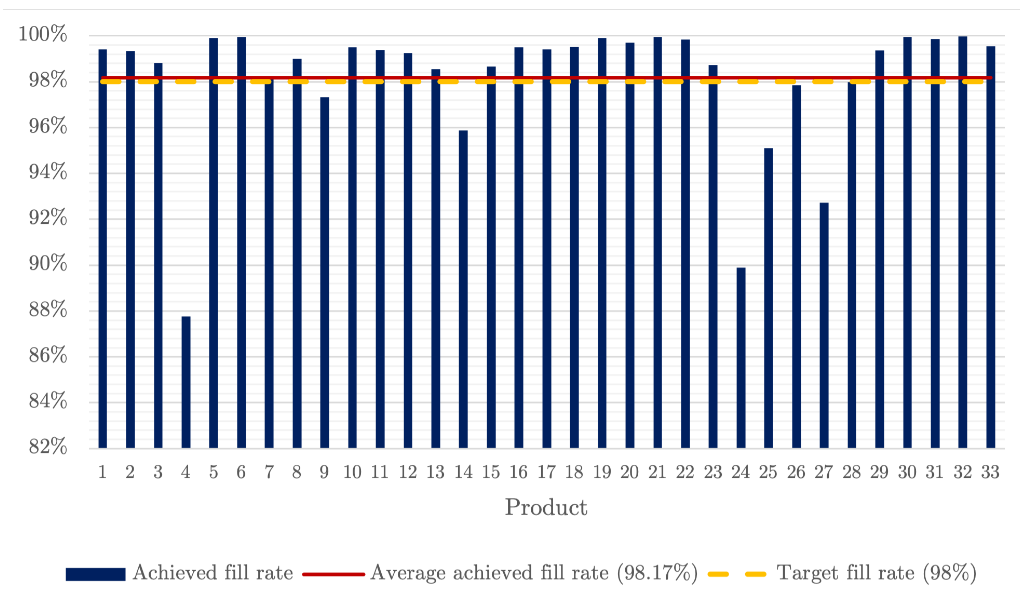

Figure 13: The results of testing the model on the first half of 2023.

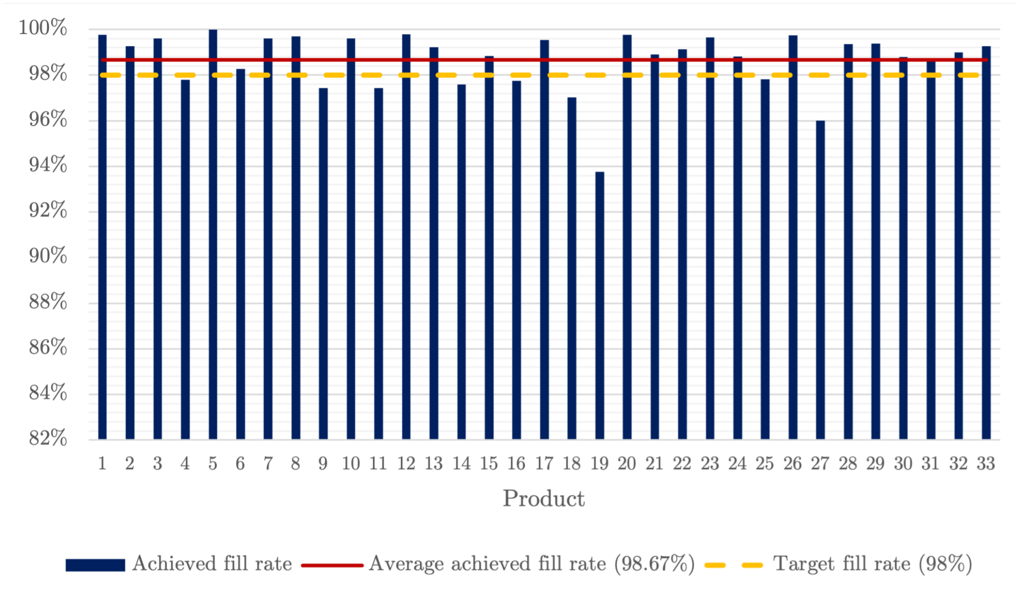

Figure 14: The results of testing the model on the second half of 2023.

The tests demonstrate that the model calculations generally achieve the target fill rates for all periods. However, products 4, 19, 24, 25, and 27 all recorded fill rates below 95% for at least one of the tested periods. These deviations are attributed to forecast errors being higher than expected. The model utilizes past forecast errors as a proxy for anticipated errors in upcoming periods; however, this assumption can prove inaccurate, affecting the fill rates. This challenge highlights the risks associated with managing demand uncertainties in inventory control. When dealing with multiple products, it is likely that stockouts will occur for some items, reflecting the unpredictable nature of demand and the limitations of forecasting models.

5.2.2 The safety stocks' sensitivity to the TSL

Given that safety stock levels are optimized with a target service level constraint, and this constraint is a manual input into the model, analyzing how different service level constraints affect the output is relevant. Therefore, this analysis provides insight into how sensitive the safety stock is to the choice of TSL. Also, as outlined in 1.1, there exists a trade-off between the target service level and inventory levels, which frequently presents strategic dilemmas for companies. Therefore, seeing how the SL constraints affect the results can become a valuable insight for the company. Firstly, a comparison between safety stock levels with a TSL of 97% and 98% was made and is shown in figure 15. Further, an illustration of how a TSL of 99% compares to the base of 98% is provided in figure 16.

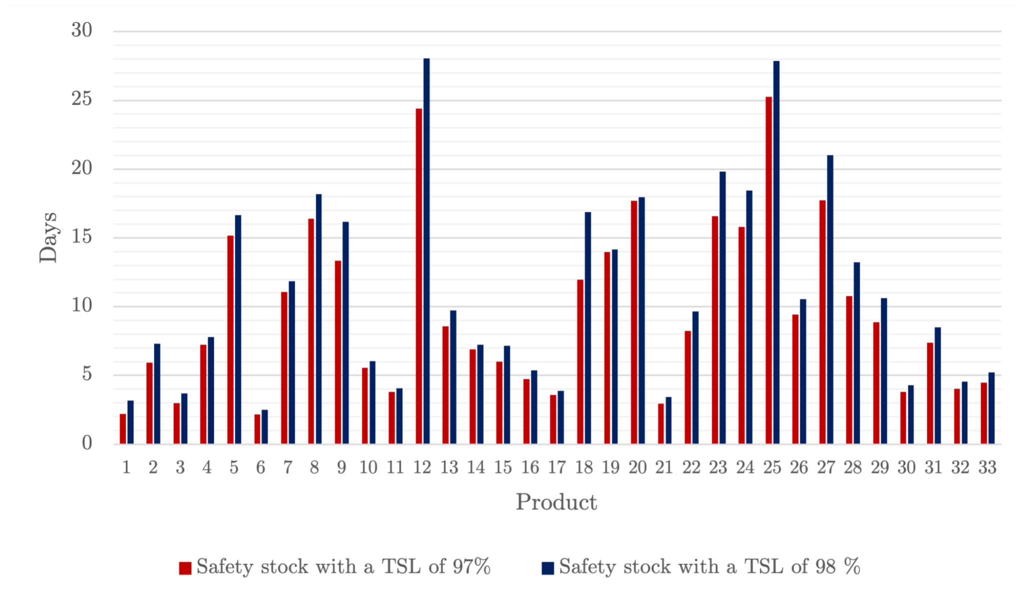

Figure 15: An overview of the safety stock levels calculated with service levels 97% and 98%.

Figure 16: An overview of the safety stock levels calculated with service levels 98% and 99%.

When it comes to implementing a uniform TSL of 97% instead of 98% results in a 13% reduction in safety stock on average. Using a TSL of 99% leads to an 18% increase compared to a TSL of 98%.

However, as illustrated in Figure 15, the impact of lowering the TSL varies significantly across different products. For instance, the differences between the outcomes for products 18 and 19 under the two scenarios are notably distinct. This variation can be attributed to the model's approach of each product having a unique probability of no stockout $(1 - P(IL \leq 0))$. The fill rate, S_2, is the same as the the probability for positive stock when it comes to continuous demand distributions and can therefore be defined as $1 - P(IL \leq 0)$ (Axsäter 2006). Since each product's $P(IL \leq 0)$ is distinct and not linear, the effects of changing the TSL manifest differently across products, leading to varying impacts on safety stock levels. Table 5 showcases the results of the comparison between the different target service levels.

Table 5: The average safety stock level calculated with three different target service levels

Target service level	**Average safety stock (days)**
97%	9.61
98%	11.07
99%	13.11

5.3 New safety stocks with segmented TSL

As discussed in section 2.1, products are segmented based on the contribution margin and as a result, each product group, A, B, and C, is assigned a target service level of 96%, 97%, and 98.5%, respectively (see table 2). Due to the fact that safety stock is currently not determined using a service level constraint, product segmentation has not been utilized in setting these levels. The case company is now interested in exploring whether this segmentation should be incorporated into the IO rather than utilizing a uniform target service level of 98% for each product and what the final safety stock level would be compared to using a uniform TSL. Therefore, new safety stock levels were calculated based on this segmentation as a request from the case company. That is, each product's target service level is based on their respective segmentation group in the ABC segmentation. The results are presented in figure 17.

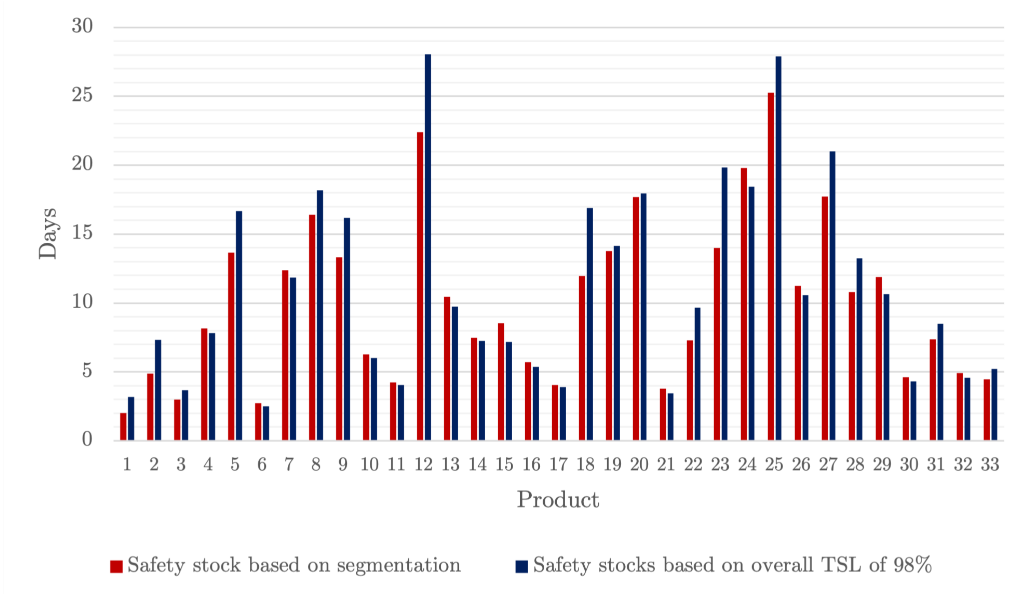

Figure 17: Safety stock values calculated with a TSL based on a ABC-segmentation compared with safety stocks calculated with an overall TSL of 98%

As the result shows, every product showcases different values compared to the initial calculation which is expected since the TSL constraint is different in the two cases. Also, it is not surprising to see that some products display a larger difference compared to the previous result due to the fact that the B and C products have a TSL of 97% and 96% respectively. On average, the safety stock was 9.5% lower when utilizing the segmentation as opposed to applying the same TSL on every product. Since some of the products have a TSL of 97% and 96%, it is interesting to see whether the safety stock levels that are based on the segmentation are high enough to achieve the average service level goal of 98% across all products. This will be evaluated in the following section.

5.3.1 Validating the model

The model was tested using the same approach as described before, but using product-specific TSL based on the company ABC-segmentation. Although it is evident that these safety stocks result in a slightly lower achieved service level compared to using an overall TSL of 98% for all products, all periods resulted in an overall achieved service level above the target of 98%. Figure 18 shows how the achieved service levels for each segments differ across the three periods.

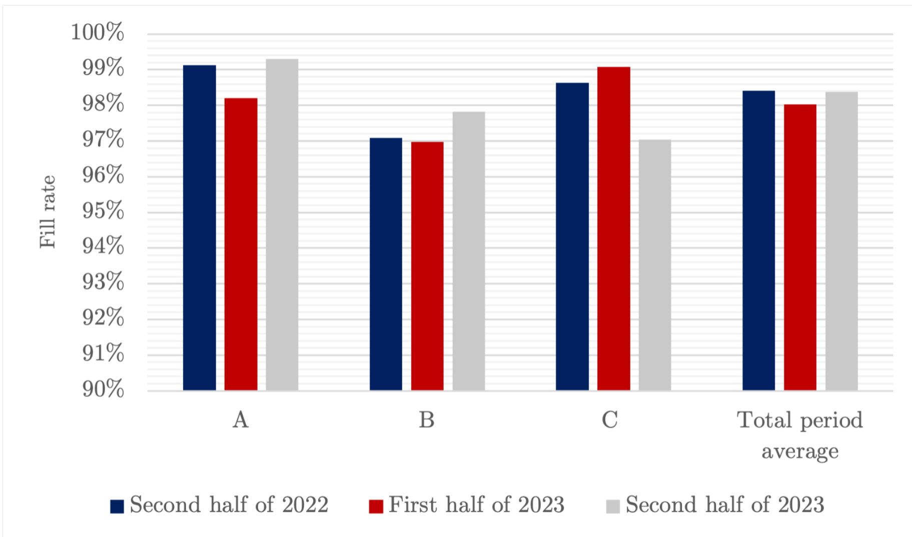

Figure 18: The resulting achieved service levels for each product segment for all three six-month periods.

The results show that all six-month periods achieved an average service level above the target of 98%. As seen in figure 18, the achieved service levels for mainly products in the C-category are significantly higher than their target of 96%, particularly in the first two time periods. The overachievement in the C-segment has to do with the products' forecast errors being lower than anticipated on average, leading to higher service levels than planned for. Based on this analysis, the segmentation could be seen as a reasonable option to set the target service levels.

5.4 New safety stocks with safety stock thresholds

The case company requested the evaluation of two additional scenarios for safety stock calculations: one scenario with a minimum safety stock of one week for all products, and another with a two-week minimum. This request was made for several reasons. Firstly, the company operates on weekly planning cycles, and safety stocks that deviate significantly from 7 or 14 days could complicate supply planning. Secondly, the model is intended primarily as a benchmark tool for the company, useful for assessing the IO's outputs and providing a general overview of which products might have high potential for safety stock reductions. Thirdly, the model does not incorporate all potential variables, such as longer maintenance stops or other factory closings, which will be discussed later in section 6. Finally, implementing safety stock reductions is a gradual process that requires careful planning and time, often extending over several months, making the incremental weekly step changes an alternative for implementation (Zanakis et al. 1980). The results for the scenario with a one-week minimum enables a safety stock reduction of 43%. This is shown in figure 19 and resembles a reduction approximately of 57 000 units in safety stocks on average.

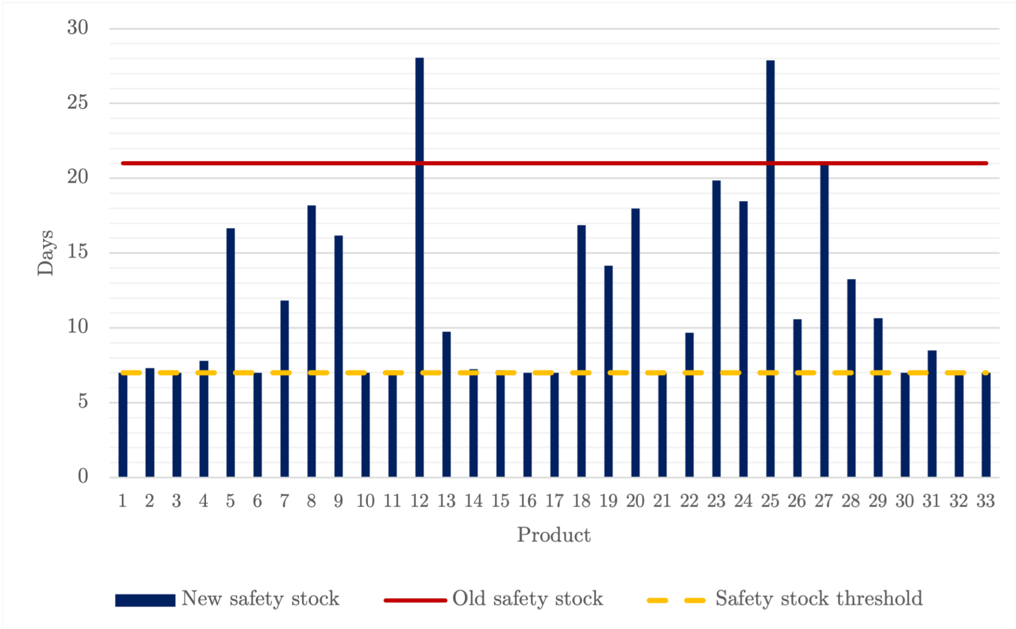

Figure 19: New safety stock levels compared to old safety stock levels when the minimum safety stock is set to 7 days.

The outcomes of the minimum two-week scenario are depicted in Figure 20, resulting in a reduction of 25% in days and approximately 33 000 units on average. These results, under conditions of minimum safety stock, highlight the excessive levels of current stock holdings, thereby illustrating that there is significant potential for stock reductions even when adopting a cautious approach to implementing changes.

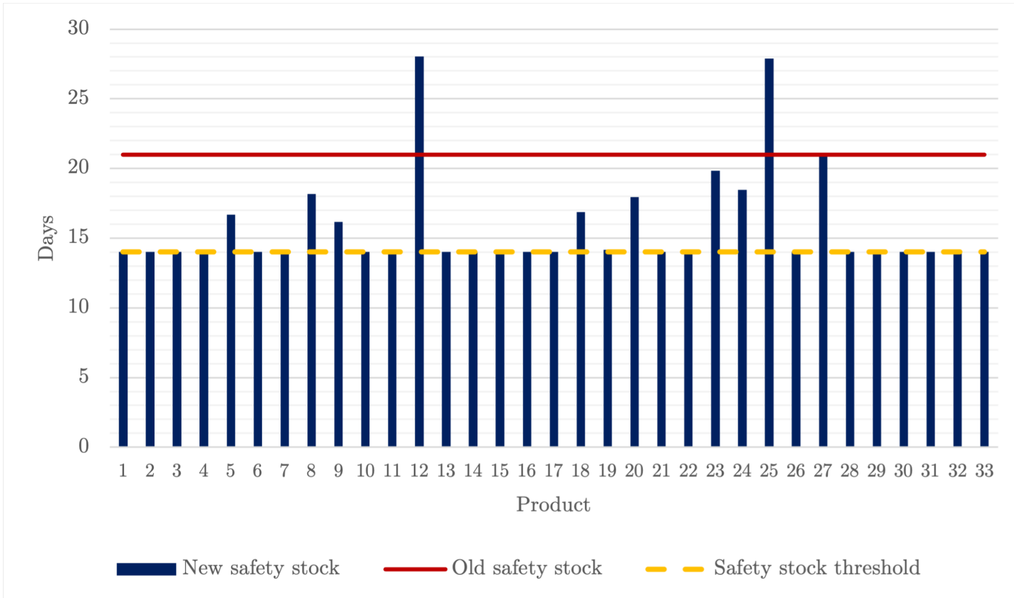

Figure 20: New safety stock levels compared to old safety stock levels when the minimum safety stock is set to 14 days.

5.5 Summary of new safety stocks and validation results

To summarize the results, it is evident that there is potential to reduce the safety stock levels at the case company. Different approaches have been tested, where every approach results in a lower average safety stock compared to the current levels at the case company. An overview of the new average safety stock levels is presented in table 6.

Table 6: A summary of the resulting safety stocks based on chosen approach

Approach	**Average days of safety stock (days)**
Current company solution	21.00
Uniform TSL of 97%	9.61
Uniform TSL of 98%	11.07
Uniform TSL of 99%	13.11
TSL based on ABC segmentation	10.06
Min. 7 days of safety stock with a uniform TSL	12.00
Min. 14 days of safety stock with a uniform TSL	15.86

An important part of the result analysis was to determine the performance of the model. The fact that reducing safety stock is a sensitive process has previously been discussed. Therefore, it is important to demonstrate that the model works and that the new safety stock values are reasonable. The results of testing the model showcase that the model provides reasonable values of the safety stock that in turn are enough to achieve the overall average service level goal of 98% that the case company has agreed with its customers, even though some products showcase lower achieved service levels.

Table 7: A summary of the achieved fill rates for the tested time periods based on applied approach

Period	**Average achieved fill rate using a uniform TSL of 98%**	**Average achieved fill rate using a TSL based on ABC segmentation**
Second half of 2022	98.95%	98.40%
First half of 2023	98.17%	98.02%
Second half of 2023	98.67%	98.37%

5.6 The impact of the input parameters

The case company is interested in recalculating safety stocks and comparing these new calculations with the outputs from the IO. While the model facilitates this comparison, it also serves as a tool for gaining insights into how various factors influence safety stock levels. Currently, there is a limited understanding of the specific contributions of different factors to safety stocks. Therefore, further analyses and discussions exploring various scenarios and their impacts on the model's output are presented below.

5.6.1 Using MAD vs historical demand standard deviation

Recalling the discussion in 4.3, two different approaches could be used in regards to demand variability. As mentioned above, the forecast error was ultimately used as the variability parameter. However, it was noted that there was a significant difference in the average variability when comparing the standard deviation based on the forecast error and the historical demand standard deviation. Since the variability significantly impacts the safety stock calculations, the two different scenarios where tested to provide an overview to what extent it actually impacts the final results. In figure 21, an illustration of the result comparison between using forecast error or standard deviation is presented.

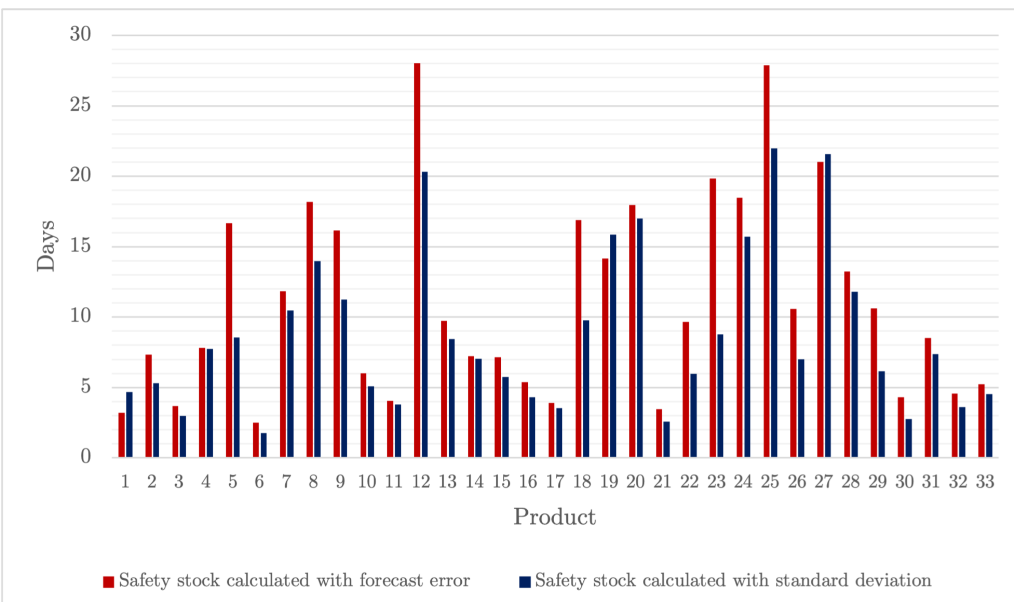

Figure 21: An overview of the safety stock levels when using forecast error and standard deviation as a metric for variablitiy

As demonstrated in the figure, the variability is generally higher when using forecast errors compared to historical standard deviation. On average, the variability using forecast errors is 19% higher, resulting in an average increase in safety stock of 21%. This result indicates the need for improvements in the forecast, since it accounts for a higher variability than necessary. Notable examples of products that exhibit significant differences are products 5 and 23. For product 5, the variability parameter is 38% higher when using forecast errors, leading to a safety stock that is 94% higher than what would be calculated using standard deviation. Similarly, for product 23, the variability parameter is 74% higher with forecast errors, resulting in safety stocks more than doubling, with an increase of 126%. These examples highlight the important role of the demand variability parameter in the model, as it is elastic and significantly influences the output. Moreover, the differences in safety stock levels for these products raise concerns regarding the accuracy and quality of the forecasting process. The forecasts for these products appear to account for an excessively high level of variability, consequently leading to inflated safety stocks.

5.6.2 Impact of lead times

The discourse regarding lead times and how to interpret them resulted in a new approach to calculate the values for them. This approach is used in the model as standard values for the lead time. Since the mean and the standard deviation of demand are directly impacted by the lead time in the model, the safety stock is ultimately affected as well. As the IO takes on the current system values of lead times, L_1, a comparison between how the resulting safety stocks differ between using L_1 and L_2 and shown in figure 22.

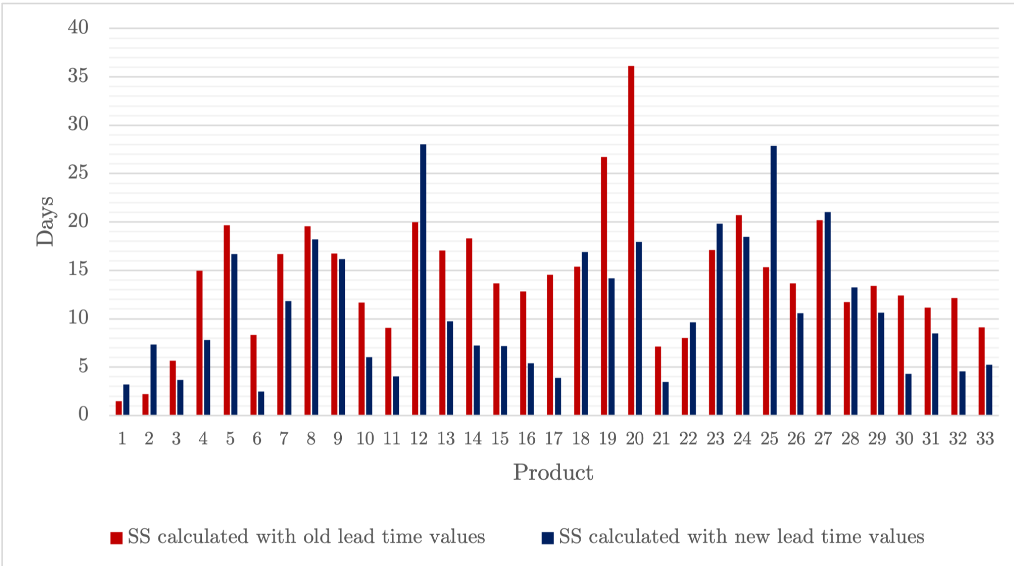

Figure 22: A comparison of safety stock levels the value of lead time is calculated with the new approach compared to old values of lead time.

When the values in table 3 are compared with figure 22, a clear pattern emerges: reductions in lead times lead to decreases in safety stock levels. This fact is not surprising, since a longer lead time will translate to a longer time where stockouts can occur and thereby a higher probability of stockouts during the lead time. For instance, product 20 experiences a 55% reduction in lead time, which correlates with a 49% reduction in safety stock. Furthermore, the new lead times are, on average, 30% shorter than the old lead times across all products. However, the corresponding safety stocks have only decreased by 23% when the new lead times are applied. This discrepancy indicates that while lead time reductions significantly impact safety stock requirements, their effect is less elastic compared to the forecast error parameter discussed earlier.

5.6.3 Impact of lot sizes

As previously noted, the value of the lot size, Q, is currently non-adjustable in the Inventory Optimizer as the IO defaults to using the minimum lot size specified in the planning system. Therefore, the minimum lot size, Q_{min}, has been used as the standard value for the lot size in the model. However, there has been some discussion about whether this limitation significantly affects safety stock levels, making the case company wondering whether SAP can offer a solution to modify this setting. An alternative approach under consideration at the case company is to use the average of achieved historical lot sizes for each product as the lot size input for the IO. This method might offer a more realistic representation since it reflects actual ordering behavior rather than a potentially arbitrary minimum set in the system. To calculate the average lot size, historical data of the lot sizes for each product were collected and the average of these were calculated. To evaluate the impact that these two different approaches have on the safety stocks, an analysis was conducted. Table 8 presents the lot sizes for all products involved in this analysis, highlighting substantial differences between the two approaches. The results of the comparison, which assess how using a historical average versus the minimum lot size affects safety stock requirements, are presented in Figure 23.

Table 8: A comparison of each product's minimum lot size and average lot size.

Product	Q_{min}	Q_{avg}	Product	Q_{min}	Q_{avg}
1	2010	2339	**18**	120	536
2	618	1838	**19**	120	365
3	618	2128	**20**	120	536
4	3216	4764	**21**	618	2613
5	166	620	**22**	618	2472
6	892	2776	**23**	672	1902
7	2920	4709	**24**	3216	6131
8	2190	4269	**25**	3832	4623
9	2190	3249	**26**	3216	5452
10	2920	5183	**27**	3795	4213
11	2920	4527	**28**	3186	6187
12	2190	3088	**29**	3186	6104
13	2280	4032	**30**	606	2736
14	3216	5948	**31**	606	2509
15	3216	5238	**32**	606	2818
16	3216	6160	**33**	606	2509
17	3216	4670			

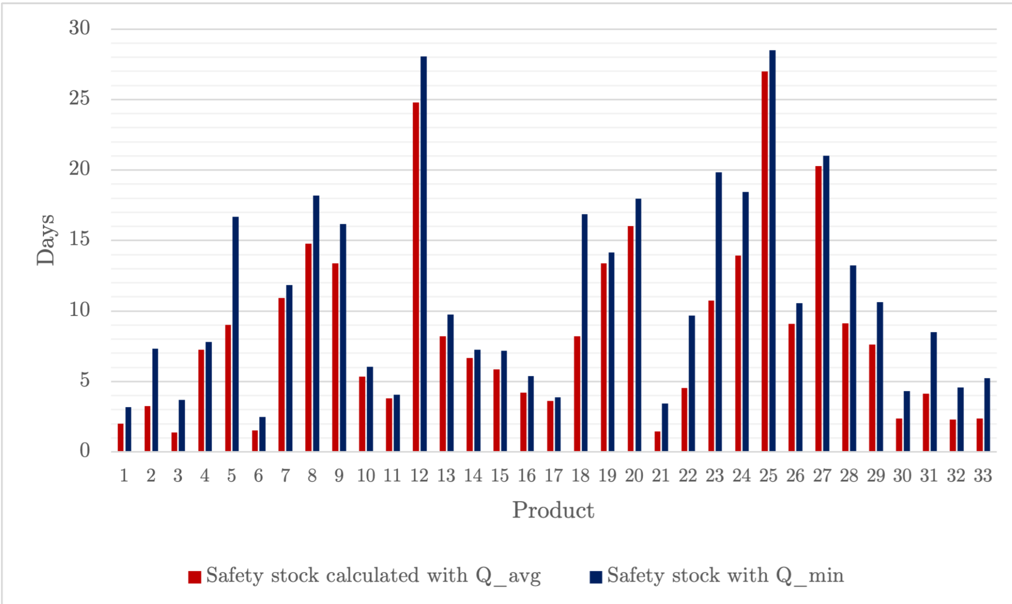

Figure 23: A comparison of safety stock levels the value of lead time is calculated with the new approach compared to old values of lead time.

A notable observation in this analysis is that safety stock levels are generally lower when calculated using the average lot size, showing a 21% decrease on average across products compared to calculations based on minimum lot sizes. This suggests that lot size significantly influences safety stock levels, a factor that the case company should consider. However, it is important to note the considerable differences in lot sizes between Q_{min} and Q_{avg} shown in table 8. Utilizing average lot sizes results in lot sizes that are, on average, 82% larger than the minimum lot sizes. This indicates that while lot sizes impact safety stock levels, the magnitude of the effect varies with the differences in Q values.

Intuitively, the influence of lot size on safety stocks may seem minimal when considering the inventory level fluctuations over a brief period, as illustrated in figure 2. Initially, the size and timing of the order appear to be independent factors. However, over an extended period, Q demonstrably affects safety stock levels. Increasing Q reduces the frequency of orders, leading to fewer ordering cycles. An ordering cycle is the period during which cycle stock is held to satisfy demand, and safety stock is added to the cycle stock as a buffer against demand uncertainties. With fewer stock cycles over a longer duration, the likelihood of stockouts decreases, allowing for reduced safety stock during such periods (Mattsson 2007a; Natarajan and Goyal 1994; Korponai et al. 2017).

6 Critical discussion

It's now crucial to reflect on the approach and results with a critical perspective. Several assumptions were made throughout the project, and practical issues raised by the case company were not fully considered. The following chapter addresses these factors. The discussion will be framed from the perspective of research quality, emphasizing reliability and validity, presented by Höst et al. (2006). Reliability concerns the consistency of a measure, or the ability to reproduce results under the same conditions. Validity, on the other hand, is the accuracy of a measure, indicating whether the results truly represent what they are intended to measure.

6.1 The accuracy of the lead time approach

The first aspect to consider is the approach used to interpret lead times and establish lead time values. When it comes to reliability, which is the consistency of measure, data from the planning system was extracted manually to ensure consistent collection. However, as noted in section 3, a production ramp-up occurred prior to the summer of 2023. Including this data would have compromised reliability because the production frequencies would not accurately represent current conditions. Thus, data points from before the summer of 2023 were excluded, reducing the dataset and impacting validity, which is the accuracy of measure. This trade-off between reliability and validity was managed by ensuring that all products had undergone enough production cycles to provide representative results. Additionally, as new data points emerge each week, the validity will improve over time. If the study were conducted in a year, the validity would likely be higher due to the accumulation of new data.

Another factor to discuss from a validity perspective is the precision of the lead time definition. Lead time refers to the interval between when an order is placed and when it is stocked. Quantifying this is challenging due to limited knowledge of the exact ordering time, as orders are often managed by the supply planning system according to a production plan, and capacity constraints often lead to extended replenishment times, primarily due to queuing. The exact queuing time is difficult to determine without a cyclical approach because the order time is not fully known.

The chosen approach prioritized time efficiency and data availability. An alternative would be a detailed case study, which could have delved into product-specific lead times by accounting for technical system functions and planning policies. However, after informal discussions with production planners, it was agreed that assuming cyclical planning would yield a "good enough" lead time estimate. This approach enabled the use of quantitative system data rather than relying on qualitative interviews and process analysis, simplifying the analysis. Additionally, a qualitative approach would have required higher research quality standards, as it could lead to results influenced by interpretation rather than concrete system data, further complicating validity.

6.2 Lot sizes and real-world production factors

As mentioned, the IO currently uses the minimum lot size in its safety stock calculations, and since the main goal of the model presented in this thesis is to serve as a benchmark to the IO, the minimum lot size was utilized to optimize safety stocks. However, as demonstrated, the impact on the output is influenced by the lot size, which makes its use noteworthy. The adoption of the minimum lot size in safety stock calculation has its drawbacks since it may not accurately reflect real-world conditions. The results from calculating a more representative lot size, i.e., using average lot sizes, revealed significant differences. This discrepancy could suggest that using the minimum lot size to calculate safety stock in this thesis might be viewed as a questionable decision. Nonetheless, this decision was influenced by the current limitations of the IO.

Furthermore, the method for determining average lot sizes involved data collection from the planning system, with data prior to the summer of 2023 excluded to ensure reliability. Like the approach for lead times, incorporating real-world production factors, such as the annual closure of Plant X in July, requires adjustments. This closure requires an increase in safety stock levels to offset the absence of production during that period. While this situation could affect the reliability of data collection due to potential drastic production ramp-ups, discussions with supply planners at the case company indicated that the ramp-up is gradual, with production volumes increasing steadily over a long period before the closing. Consequently, the risk of compromising data reliability was considered low, and this factor was not included in the data collection. However, it is important for users to be aware that the model does not accommodate for this period.

6.3 Limitations of the demand- and forecast data

An essential aspect to address is the demand data. The primary issue is that the model does not account for promotions, which may affect its validity. During data analysis, weekly sales data was extracted from SAP Analytics. This data did not specify whether it was from a promotional period or linked to a specific campaign, potentially increasing perceived variability. While promotions are a part of reality and should generally be considered in modeling, the primary issue here is that the case company often knows about promotions in advance and adjusts supply plans accordingly. The model, however, assumes promotions do not exist, limiting its ability to anticipate these fluctuations.

Another implication of excluding promotions is the discrepancy between actual and forecasted sales data. During the new project managed by the demand planning department, a *base forecast* is established, where historical data is cleaned to remove the effects of past promotions. The forecast quantity is then adjusted based on a seasonal index (see equation (13)). Although earlier forecasts were not created using this exact method, similar data cleaning practices were used, potentially creating a gap between actual and forecast data, leading to an overestimation of forecast errors. Despite this potential validity problem, consistent data use in both setting safety stocks and validating the model allows the results to remain useful for the intended purpose. The model aims to act as a benchmark and highlight areas for reducing stock levels, so this overestimation is not seen as a critical issue.

6.4 Discussing the data in general

Another noteworthy aspect is the limited data availability, as the data provided by the case company only dates back to early 2021. This issue is partly discussed in section 4, where e.g. nonstationary demand volumes during the data period are noted. For instance, in the case of the one seasonal product discussed in section 4.1.2, it is not statistically proven due to the analysis being based on only two years. It might merely be coincidental that the volume rises during the same month two years in a row. However, as a result of discussions with experienced supply planners at the case company, it was qualitatively confirmed that this is indeed a seasonal product because it is strongly linked to a Swedish holiday.

The validity problem that might arise with the limited data availability cannot be completely resolved. However, section 4 outlines how the data analysis process considered this limitation when drawing conclusions, ensuring that the model's output is not overly reliant on data availability, though this remains a factor worth discussing.

In conclusion, achieving the highest possible research quality has consistently been a primary goal, aligned with the problem-solving approach. In operational research, focusing on the problem may sometimes compromise research quality due to prioritizing solutions. However, to ensure consistency and replicability when implementing this solution in additional production plants within the case company, research quality remained a key factor throughout the project.

7 Conclusions & future work

Reflecting on the thesis's purpose, a model has been developed to calculate safety stock levels for products manufactured at Plant X and stored at Warehouse Y. The model aims to replicate the inventory control processes of the case company as accurately as possible. The results demonstrate significant potential for reducing safety stock levels, even when considering a gradual reduction process as discussed in Section 5.2. The new safety stock levels represent the optimal level with respect to the TSL of 98%. However, a more cautious approach, which is setting minimum safety stock thresholds 14 days, could still lead to a 25% reduction in total average safety stock levels compared to the current safety stock levels at the case company. This conclusion offers the case company a broader perspective and could be used as a three-step approach for reaching the optimal level in the long term. Additionally, the results provide insight into products with high potential for significant safety stock level reductions and those that require more careful consideration.

Beyond optimizing safety stock levels, the goal was to use the model as a benchmark for the IO. As demonstrated when validating the model, the model yields reasonable results, further solidifying its value as a benchmark option. It's crucial to remember the critical aspects of the model as discussed in Section 6. Unlike the IO, which calculates safety stock levels weekly, resulting in potential fluctuations, this model calculates a fixed safety stock level. While this may occasionally lead to safety stock levels being slightly too high or too low, the case company could still use the model as a tool for benchmarking the IO's performance in the future.

The thesis also delves into various input parameters, with a particular focus on lead time. Extensive discussions throughout the thesis have led to an approach for setting lead time values, despite the limitations such as a scarcity of data points. Recalculating lead times with more data points in the future can yield more accurate values, as discussed in Section 6. Furthermore, the approach was created with the objective to be replicated on different production plants. When looking at lead time as per 8, one could replicate the approach on a different production plant even though conditions like production lead time might be different compared to this case.

Lastly, while parameters like lot size, target service level, and variability have not undergone as thorough analysis as lead time, the analysis presented in this thesis provides valuable insights for future discussions and decision-making within the case company. It can act as a guide for prioritizing parameters or conducting further careful analysis.

7.1 Future work

While this thesis concentrated on Warehouse Y and Plant X, the model and approach presented here can be expanded to include products from other plants and warehouses at the case company. Although products at other warehouses might be different compared to the products analyzed in this thesis, the approach has been designed to be applicable company-wide if required. Therefore, this thesis can serve as a guide for creating similar models tailored to specific plants or warehouses. As this thesis focuses only on products that fit continuous demand distributions, the model may not be suitable for discrete demand distributions. After successful implementation of the IO, further research should address modeling products that cannot be approximated with continuous demand distributions to systematically calculate safety stocks across all company products. Lastly, due to the primary goal of this study being to develop a model that serves as a benchmark for the IO during its implementation, future work should focus on using this model and the findings to compare and test against the IO's output.

References

Axsäter, S. (1991). *Lagerstyrning*. 1st ed. Lagerstyrning. Studentlitteratur. ISBN: 9789144334912.

– (2006). *Inventory Control*. 3rd ed. International Series in Operations Research & Management Science. Springer International Publishing. ISBN: 9783319157290.

Burgin, TA and Norman, JM (1976). “A table for determining the probability of a stock out and potential lost sales for a gamma distributed demand”. In: *Journal of the Operational Research Society* 27, pp. 621–631.

Fildes, Robert and Goodwin, Paul (2007). “Against your better judgment? How organizations can improve their use of management judgment in forecasting”. In: *Interfaces* 37.6, pp. 570–576.

Hillier, F.S. and Lieberman, G.J. (2010). *Introduction to Operations Research*. Introduction to Operations Research. McGraw-Hill Higher Education. ISBN: 9780071267670.

Höst, M., Regnell, B., and Runeson, P. (2006). *Att genomföra examensarbete*. svenska. Studentlitteratur AB. ISBN: 91-44-00521-0.

Keaton, Mark (1995). “Using the gamma distribution to model demand when lead time”. In: *Journal of Business Logistics* 16.1, p. 107.

Korponai, János, Tóth, Ágota Bányainé, and Illés, Béla (2017). “Effect of the Safety Stock on the Probability of Occurrence of the Stock Shortage”. In: *Procedia Engineering* 182. 7th International Conference on Engineering, Project, and Production Management, pp. 335–341. ISSN: 1877-7058.

Lau, Hon Shiang and Lau, Hing Ling (2003). “Nonrobustness of the normal approximation of lead-time demand in a (Q, R) system”. In: *Naval Research Logistics (NRL)* 50.2, pp. 149–166.

Mattsson, Stig A. (2007a). “Efterfrågefördelningar för bestämning av säkerhetslager”. In: *Lund University*.

– (2007b). *Standardavvikelser för säkerhetslagerberäkning*. Project report. Next Generation Innovative Logistics (NGIL).

McCarthy, Teresa M, Davis, Donna F, Golicic, Susan L, and Mentzer, John T (2006). “The evolution of sales forecasting management: A 20-year longitudinal study of forecasting practices”. In: *Journal of Forecasting* 25.5, pp. 303–324.

Natarajan, R and Goyal, SK (1994). “Safety stocks in JIT environments”. In: *International Journal of Operations & Production Management* 14.10, pp. 64–71.

Olhager, Jan (2013). *Produktionsekonomi: principer och metoder för utformning, styrning och utveckling av industriell produktion*. 2:3. Studentlitteratur. ISBN: 9789144067667.

Shcherbakov, Maxim Vladimirovich, Brebels, Adriaan, Shcherbakova, Nataliya Lvovna, Tyukov, Anton Pavlovich, Janovsky, Timur Alexandrovich, Kamaev, Valeriy Anatol’evich, et al. (2013). “A survey of forecast error measures”. In: *World applied sciences journal* 24.24, pp. 171–176.

Snyder, R.D. (1984). “Inventory control with the gamma probability distribution”. In: *European Journal of Operational Research* 17.3, pp. 373–381. ISSN: 0377-2217.

Tyworth, John E and O’Neill, Liam (1997). “Robustness of the normal approximation of lead-time demand in a distribution setting”. In: *Naval Research Logistics (NRL)* 44.2, pp. 165–186.

Zanakis, Stelios H, Austin, Larry M, Nowading, David C, and Silver, Edward A (1980). “From teaching to implementing inventory management: Problems of translation”. In: *Interfaces* 10.6, pp. 103–110.

A Model user interface