To Buy or not to Buy: A Thesis about Estimating and Accounting for Partial Substitution in Retail Inventory Management.

Master Graduation Thesis

Master Industrial Engineering and Management

D.M. Kösters BSc. (Daan)

To Buy or not to Buy: A Thesis about Estimating and Accounting for Partial Substitution in Retail Inventory Management.

A master graduation thesis submitted to the University of Twente as partial fulfilment for the degree of

Master of Science in Industrial Engineering and Management

Commissioned by
Slimstock B.V.

Internal supervisors
Dr. D.R.J. Prak (Dennis)
Dr. M.C. van der Heijden (Matthieu)

External supervisors
B. van Gessel (Bart)
M. Hozeman (Miranda)

By
D.M. Kösters BSc. (Daan)
d.m.kosters@student.utwente.nl

Enschede, August 19, 2023

Management Summary

Slimstock is a leading software and consulting firm specialising in inventory optimisation solutions. Their software platform SLIM4 helps businesses increase efficiency and reduce waste by providing real-time visibility into their inventory levels, accurately forecasting their future demand and optimising their inventory management. One of the things they see as an opportunity as an addition to their software is a way of dealing with the concept of substitution. Substitution is the phenomenon of customers, typically in the retail industry, opting for another similar product if their initially desired product is out of stock. With full substitution, all customers opt for a substitute upon a stock out, but product groups are far more likely to experience partial substitution. Here, a part of the demand is substituted once a stock out occurs, and the rest is considered lost sales as the customers leave the store to buy their desired product elsewhere. The following research question is formulated to research this phenomenon by designing a way of accounting for the substitution effect in the inventory management approach.

"What is the effect of accounting for partial substitution on the order-up-to level advice for inventory models?"

In literature, several theories have been found that approach this problem. All of them use a substitution parameter, which is the portion of demand that goes from product i to product j upon a stock out of product i. From these theories, two promising ways of estimating this parameter are found. These two methods, a weighted least-squares-based approach and a Maximum Likelihood Estimation (MLE), are combined, and one final parameter estimation method is devised. Next to the substitution parameter, the initial demand rate, being the stationary demand rate for product i without any substitution effects, should be estimated. Luckily, the same two papers propose methods for that.

The estimators for the substitution parameter and initial demand rate require knowledge of the inventory state, which describes whether a product is in stock or out of stock in a given period. Due to the fact that the provided data set does not have exact data regarding whether a product is in stock or not, transaction data is used to approximate this utilising the number of subsequent periods without sales and the corresponding probability if the product were in stock. The result of this statistical approximation, combined with the transaction data provided, is then used as input for the found estimation methods. These methods result in two sets of estimated parameters, which are evaluated based on three performance measurements to select the best-performing one given the provided data set. After this selection,

new order-up-to levels are heuristically determined based on a target fill rate and minimal inventory value. The inventory policy already in place in SLIM4 can then use these order-up-to levels.

Subsequently, multiple experiments are performed in order to validate this model's performance using theoretical data. This theoretical data is transaction data generated using a Poisson distribution with chosen values for the substitution parameters, initial demand rate, number of periods and stock out percentage. This data set, which, for the sake of demonstration, only contains transactions for two products substituting both ways, is used as input for all three separate parts of the created model. The results are then compared to the input parameters to validate the method's performance.

Regarding the inventory state approximation, the substitution parameter does not influence the method's performance. However, the initial demand rate does influence performance. For values lower than 12, this method cannot discern in stock from out of stock periods as well, meaning 12 should be the minimum value. The stock out percentage also influences the performance. For values lower than 55%, the performance is rather good, but after that, it worsens significantly. Luckily, having a stock out 55% of the time is not good anyhow, making this limit not as bad. More realistic values of this parameter perform well.

On top of that, the model's ability to correctly estimate the substitution parameters and the initial demand rate is analysed. From this, it can be concluded that the initial demand rate can be estimated rather accurately. However, the same does not hold for the substitution parameters with a low number of periods. To improve the estimation accuracy, the total number of periods is increased, which increases the accuracy. From this, the conclusion can be drawn that a minimum of 1500 periods must be analysed to estimate the substitution parameters with fair accuracy given a stock out percentage of 25%, which is a lot in practice. The same experiment was performed with a stock out percentage of 5%, resulting in a minimum number of periods of 2800. Although this stock out percentage is more realistic, the minimum required number of periods is over 7.5 years, which is a lot of data to hold. It is also shown that the substitution parameter is generally underestimated, meaning that the substitution effect is generally more significant than the model suggests.

Moreover, the comparison is made between accounting for substitution when computing the order-up-to levels and negating the substitution effect when it is actually present. This is done for both a real data set and a theoretical one. Here, accounting for the substitution effect causes the order-up-to levels of the product experiencing substitution to another product to decrease steadily. In contrast, the other product needs more inventory to keep up with the extra demand due to substitution. If the substitution is large enough, which happens from a substitution parameter of 0.4, the order-up-to level of the latter product decreases as well, and the total expected inventory value, when using the model, is lower than when negating substitution.

In concluding and answering the research question, the created model decreases the order-up-to level of a product experiencing substitution. At the same time, it increases that level of the product receiving the demand due to substitution. Nevertheless, upon a large enough substitution effect, accounting for it will result in a lower total expected inventory value, making it beneficial to use the model. Combining this with the notion that the substitution parameter is generally underestimated results in the conclusion that even more benefits could be attained than resulting from this model.

Therefore, it is recommended to first improve upon this model. The parameter estimation should be advanced, mainly on how this model deals with the input data. The model currently assumes stationary demand, meaning no trend, seasonality, or promotions are possible in the data set. In practice, this is not the case. Making sure the model is able to deal with these phenomena will further improve its performance. After that, Slimstock could look into adapting the model to fit SLIM4's approach and integrating it within the software.

Acknowledgements

For the completion of this graduation thesis, there are a few people that I would like to thank. First and foremost, these are my two primary supervisors, Dennis Prak and Bart van Gessel. Dennis guided me through this academic journey and supported me whenever necessary. I am incredibly thankful for our sparring sessions, discussing the ideal way of approaching the problem at hand, which genuinely felt like an equal conversation. On the other hand, Bart supported me more with practical eyes on my research. He helped me get out of the academic way of thinking and encouraged me to imagine the practical side of my results and conclusions. This other way of looking at things opened my eyes to how I will use it more in the future. This made it an incredible learning opportunity.

Second, I want to thank my secondary supervisors, Matthieu van der Heijden and Miranda Hozeman. They both provided me with feedback in another stage of this research. While Matthieu came into the picture towards the end of the process, Miranda mostly gave input at the beginning with the formulation of my core problem. Both contributions were invaluable to the end result.

Thirdly, I am genuinely thankful to my peers and colleagues for their constant support and occasional sparring sessions. They have helped me from a completely other point of view. Although they also have an academic background, their expertises lie from Mechanical Engineering to Business Administration and everything in between, so they could provide valuable insights from both within as well as outside my comfort zone.

I appreciate all the input and support I have received over the past half a year. The result would not be the same if I had not, so I thank you all for that.

Contents

List of Algorithms

List of Figures

List of Tables

Abbreviations and Parameters

Below, the abbreviations and parameters used in this thesis can be found. This overview is intended for easier insight into the meaning of these terms.

Abbreviations

EOQ	Economic Order Quantity.
IOQ	Incremental Order Quantity.
MAD	Mean Absolute Deviation.
MAPE	Mean Absolute Percentage Error.
MAX	Maximum Absolute Error.
MLE	Maximum Likelihood Estimation.
MOQ	Minimum Order Quantity.
POS	Point Of Sales.
SOP	Stock Out Percentage.

Parameters

$\boldsymbol{\alpha_{ij}}$	The substitution parameter from product i to product j.
$\boldsymbol{A_j(q_j)}$	The expected probability that product j will be out of stock given a certain inventory level q_j at the beginning of the replenishment cycle
$\boldsymbol{D_{i,t}}$	The sales recorded for moment t.
$\boldsymbol{\delta_{i,I_k}}$	The binary inventory substate for product i of inventory state I_k indicating whether the product is in or out of stock.
$\boldsymbol{e_{i,I'_0}}$	The error per product per inventory state.
$\boldsymbol{I_k}$	The set of inventory substates of inventory state k.
$\boldsymbol{\lambda_i}$	The initial demand for product i.
$\boldsymbol{\lambda_{i,I_k}}$	The initial demand for product i per inventory state I_k.
$\boldsymbol{\Lambda_i}$	The demand including substitution for product i.
$\boldsymbol{\Lambda_{i,I_k}}$	The demand including substitution for product i per inventory state I_k.
$\boldsymbol{\Lambda_{i,t}}$	The demand including substitution for product i per period t.

$\boldsymbol{L}$	The lead time.
$\boldsymbol{\mathcal{L}}$	The likelihood function used in MLE.
$\bar{\boldsymbol{s}}_{\boldsymbol{i},\boldsymbol{I_k}}$	The average sales per product i per inventory state k.
$\boldsymbol{\psi(d_i \mid \lambda_i)}$	The probability of the demand being d_i given a certain initial demand λ_i and a Poisson distribution.
$\boldsymbol{\psi(d_i \mid \lambda_i, \ldots, \lambda_N,\ q_i, \ldots, q_N)}$	The probability of the demand including substitution being d_i given a certain initial demand λ_i, the inventory level at the beginning of the cycle being q and a Poisson distribution.
$\boldsymbol{\Psi(d_i \mid \lambda_i, \ldots, \lambda_N,\ q_i, \ldots, q_N)}$	The cumulative distribution based on $\psi(d_i \mid \lambda_i, \ldots, \lambda_N,\ q_i, \ldots, q_N)$.
$\boldsymbol{R}$	The review period.
$\boldsymbol{\rho_i(q_i, \ldots, q_N)}$	The fill rate of product i given a set of inventory levels at the beginning of the replenishment cycle q.
$\boldsymbol{S_i}$	The order-up-to level of product i.
$\boldsymbol{W_{I_k}}$	The number of POS intervals inventory state I_k is assigned to.

1 | Introduction

This chapter is divided into two subjects. On the one hand, the company Slimstock and its customer are introduced in Section 1.1 to provide context. Furthermore, the reason for and situation surrounding this research will be elaborated on in Section 1.2, which leads to the other hand: the identification and explanation of the core problem. This is denoted in Section 1.3.

1.1 | Company Introduction

Slimstock is a leading software and consulting firm specialising in inventory optimisation solutions. The company was founded in the Netherlands in 1993 and has since expanded to have a global presence, serving customers in various industries such as retail, manufacturing, and distribution. Slimstock's software platform SLIM4 helps businesses increase efficiency and reduce waste by providing real-time visibility into their inventory levels and presenting accurate forecasts of future demand. In addition to their digital products, Slimstock also provides consulting services, including training and support, to help businesses implement and maximise the benefits of their inventory optimisation strategies. Through its innovative solutions, Slimstock has established itself as a trusted partner for businesses seeking to improve their inventory management processes and drive sustainable growth.

One of their customers is Verfsterk Groep. This family-owned corporate group comprises five companies, all with a significant presence in the paint industry. Each of these companies has its unique role in the group, ranging from wholesalers to production companies and retailers.

1.2 | Context Explanation

As indicated before, this thesis will surround the retail sector. This vast and diverse industry encompasses selling goods and services to end consumers. This industry generates trillions of dollars in revenue annually, with total worldwide retail sales reaching approximately $25.8 trillion in 2022. (Statista, 2022) Retailing encompasses a diverse range of settings, from physical stores to online shops and everything in between. One popular retail category is home improvement, which focuses on products and services related to home renovation, construction, and repair. Among other subcategories, paint stores are a standard fixture, offering customers a wide range of paint products.

Physical paint stores have been a staple of the retail industry for many years, and they continue to be popular among consumers who value the ability to see and touch products before making a purchase. These stores often carry various types of paint, including interior and exterior paint, as well as specialised paint products such as primers, sealers, and stains.

Like any other industry, the retail paint industry has its challenges. While paint stores are a crucial part of the retail sector, they struggle with the constant risk of stock outs. Such instances are detrimental to the business, causing a significant reduction in sales. However, not just the loss of revenue poses a problem, but also the dilemma it creates for customers.

Imagine a customer standing in front of the shelves filled with wall paints. They have come to buy a specific product, but unfortunately, it is out of stock. Now, the customer faces a difficult decision: to either leave the store or opt for a substitute product. This decision-making process can be stressful, especially if the customer has a specific paint colour in mind or is looking for a certain quality or brand.

In such situations, the customer might choose a similar product from a different brand or a higher or lower price point. However, it is not just the price that matters to the customer; the paint's quality and consistency are also essential.

Moreover, the decision to substitute a product or leave the store could also depend on the urgency of the need for the paint. If the customer had a pressing need for the paint, they might opt for a substitute product. However, if they had the luxury of time, they might prefer to wait for the desired product to be back in stock.

In conclusion, stock outs can create a dilemma for customers and challenge retailers. While substitute products might provide a temporary solution, ensuring they are of similar quality and consistency to the desired product is crucial. Retailers must keep their inventory up to date to avoid stock outs and provide their customers with a satisfactory shopping experience.

1.3 | Problem Identification

This section entails the identification of the core and subproblems. This is done by analysing the phenomenon stated before, after which the subject is built up from several parts that contribute to the problem's solution.

1.3.1 | Core Problem

The dilemma mentioned in Section 1.2 could occur whenever there is a stock out. At Slimstock, the phenomenon explained is currently not taken into account when deciding on what inventory to hold. However, it could be important when optimising inventory as stated in Section 1.1 that Slimstock provides. What if it is more beneficial to hold less of a particular product knowing that customers will opt for a, perhaps more expensive, substitute when the product is inevitably out of stock?

The phenomenon of customers choosing substitutes over a particular product is called substitution. However, in this case, not all customers opt for the substitute when their desired product is out of stock. Therefore, the more specific kind of substitution here is called partial substitution. This concept, explained more in detail in Section 3.2, includes a substitutability parameter that indicates what portion of customers opt for a substitute and what portion decides to leave the store.

From conversations with Slimstock, it has become apparent that they recognise the existence of partial substitution for their customers and see the potential value of its introduction into their existing inventory models. However, they do not yet know how to achieve this. Consequently, this research aims to investigate the effect of incorporating this concept into inventory models, which is then validated with a case study of one of Slimstock's customers. This purpose results in the following research question.

"What is the effect of accounting for partial substitution on the order-up-to level advice for inventory models?

1.3.2 | Sub-problems

The current situation must first be investigated to answer the main question appropriately. This exploration should not only go into detail about Slimstock's current situation regarding inventory policies but also look at the situation of Slimstock's customer. Therefore, the first sub-question is as follows, which will be answered in Section 2.

1. **What does the current situation look like for both Slimstock and their customer?**

    a) What is the inventory policy in use by Slimstock?

    b) What does the data look like for Slimstock's customer?

    c) How big is the phenomenon mentioned earlier at Slimstock's customer?

These questions provide insight into the current situation in practice. However, also the current academic situation should be catalogued. To do this, the current academic landscape should be studied. The existing papers pertaining to (partial) substitution should be found and understood, and an overall overview of the existing literature should be created. The latter leads to the second sub-question, elaborated on in Section 3.

2. **What does the current academic landscape look like regarding (partial) substitution?**

a) What literature has already been written about substitution?
b) What literature has already been written about partial substitution?
c) What literature has already been written about the quantification of partial substitution?

Once both the practical and academic situations are known, the answer to the research question can begin to be formed. As can be seen from that question, the concept of partial substitution should be studied and quantified. This is done in Section 4.3.

3. **How can partial substitution be quantified?**

a) How should the quantification of partial substitution be estimated?
b) How should the initial demand be estimated without partial substitution?

Once these parameters are estimated, the phenomenon's scale is known, but it still needs to be accounted for in the inventory model. To properly do this, the following sub-question is formulated. This sub-question is further explained in Section 4.4.

4. **How can the estimations be used to account for partial substitution in the inventory model?**

a) How will the demand distribution change?
b) How can a new order-up-to level be determined per product?

This model can then be evaluated, and insights into the effect of accounting for partial substitution in the demand estimation can be gained. This is expanded upon in Section 5.

5. **What are the results of this created model?**

a) What result does applying the constructed model at Slimstock's customer yield?
b) What is the accuracy of the constructed model?
c) What other insights can be acquired from using this created model?

Finally, it is important to know what these results will bring. The takeaways from and opportunities to build upon this research should be identified. This is elaborated on in Section 6.

6. **What are the takeaways from this research?**

a) What can Slimstock do with the result of this research and by extent the created model?

b) What can the academic environment do with the result of this research?

c) What possibilities exist for further research to deepen the understanding of this phenomenon?

1.4 | Theoretical Relevance

The theoretical relevance of this thesis mainly consists of the combination and application of several previously studied topics. Although the concepts of (partial) substitution and its quantification are not innately new in theory, the combination has yet to be found in literature, as shown in Section 3. On top of that, applying partial substitution to the determination of the order-up-to levels has not been done before. More elaboration on the existing literature regarding this subject can be found in Section 3.

1.5 | Practical Relevance

Based on that being said in Section 4.4, the created model and the results of this research can be used by Slimstock as a building block to integrate this into their own software. At the end of this thesis, the model is evaluated on one case, but doing this on multiple cases will ensure the significance and confidently portray the effect of implementing the concept of partial substitution.

2 | Current Situation

In this section, the current situation will be given. This section will start with an explanation of the process of SLIM4 in Section 2.1, in which a short introduction to SLIM4 will be recapitulated. Then, the data received from Verfsterk Groep will be explored, and an initial descriptive analysis will be provided in Section 2.2. This exploration is done to get an overview of their situation, especially regarding stock outs during which substitution could occur.

2.1 | SLIM4's Process

As stated before, the SLIM4 software, as well as the process of SLIM4, will be elaborated on here. First, the former will be explained. As can be read in Section 1.1, SLIM4 is software that optimises inventory management. It does this based on four main steps, as shown in Figure 1.

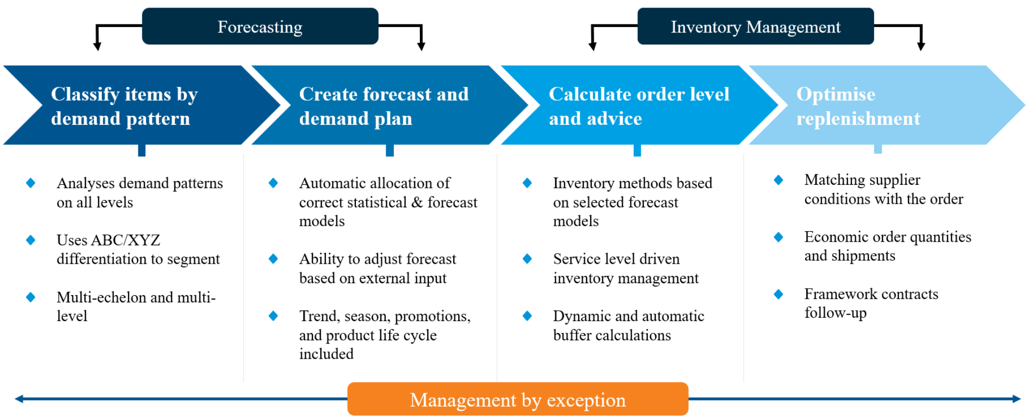

Figure 1: The four main steps of SLIM4. (Slimstock B.V., 2023)

As shown, this process starts with a demand classification. This categorisation is mainly based on demand size and frequency. Every item is classified into one of five demand classes to determine the forecast accurately. Next to the demand class, the possible trend and seasonality are determined to improve the expected future demand further. These input parameters, along with the historical demand, are used to forecast the upcoming months and calculate an order-up-to level. In this order-up-to level, a safety stock is included to account for the stochasticity of demand. This is then used to advise the user on what to order and when, which is then monitored to finalise the process.

2.2 | Data Exploration

As part of this thesis' analysis, the data set received is a backup of the SLIM4 database of one of Slimstock's customers. The data set was extensive, including data deemed irrelevant for this research, such as SLIM4 analysis results and SLIM4-specific data. The data had to be explored to find the most exciting parts. The goal of the exploration is not necessarily to differentiate relevant data from irrelevant data but rather to find a subset of promising data within the relevant data to use for the model to be created. Also, a general idea of the customer's current situation will be formed.

Upon analysing the data, it has been found that on average a product experienced a stock out 2.09% of the time for the entire data set of two years. However, it is also discovered that a trend of improvement is present. This negative trend of 0.43 percentage points shows a positive sign for the company as it indicates that their inventory management has improved. This stock out percentage is calculated by first determining the number of days a product has been out of stock, which is then divided by the total number of days in the data set.

The separate product groups are examined individually to provide insight into the differences between groups. This is done similarly to the approach explained before and shown in Figure 2. In this figure, only the top and bottom 25% are shown for clarity, which is the most interesting for this analysis. The entire figure can be seen in the appendix in Figure 36. As can be seen in this figure, some product

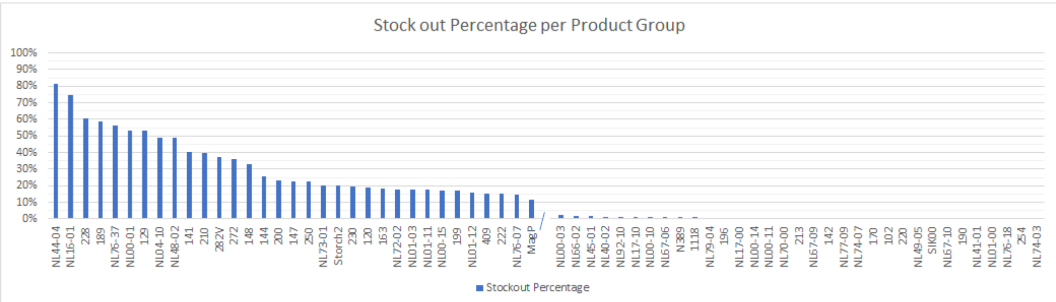

Figure 2: Stock out Percentage per Product Group.

groups have a large stock out percentage. However, upon closer examination, it was revealed that these product groups generally have between one and five products. This revelation might indicate that they are harder to manage or that those groups are at the beginning of their introduction, making demand forecasting difficult as there is little to no historical sales data.

Nevertheless, most of the product groups have a stock out percentage of less than 0.05, meaning that these perform relatively well based on the previous two years. Regarding the choice for the most relevant subset of products for this research, selecting one or more product groups is intuitively sound as these will have a higher chance of having substitutes within the group, which increases the likelihood of partial substitution.

Next, the expected lost sales due to stock outs were calculated by multiplying the number of days with stock outs per product by the average sales per day and aggregating that for all product groups. This time, the graph is split into two, again for clarity reasons. The product groups with a total number of average lost sales per year larger than 5000 are shown in Figure 3, and a total number of average lost sales per year smaller than 2500 shown in Figure 4. The whole figure can be seen in Figure 37 in the appendix. Putting these two figures together shows similarities

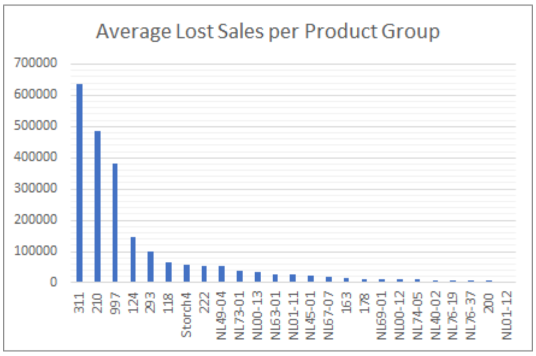

Figure 3: Average Lost Sales per Product Group Larger than €5000.

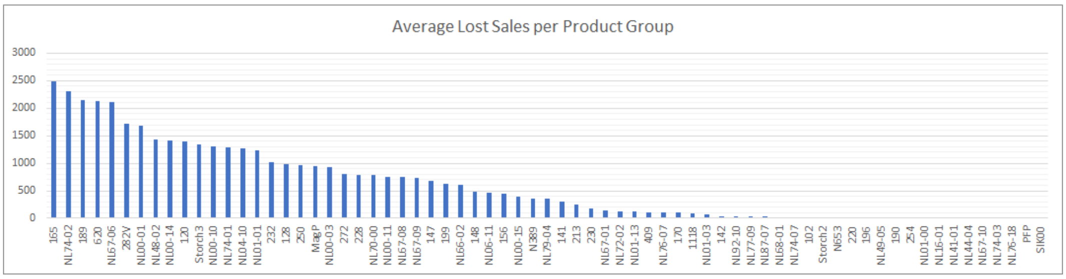

Figure 4: Average Lost Sales per Product Group Smaller than €2500.

to the negative exponential distribution. Contrary to the previous analysis, the top three product groups have a high number of products. Nevertheless, these product groups' exceptionally high values beckon the question that there is something at play other than insufficient stock.

However, to truly understand the financial impact of stock outs, the average lost sales are multiplied by the selling price of each product and aggregated for all product groups to find the average lost revenue per product group. This calculation provides a more insightful view of the potential revenue loss due to stock outs. Here, the total figure is shown in the appendix in Figure 38 and a split-up version can be seen in

Figures 5 and 6. The cut-off point in the previous analysis was set at 5000 pieces, but here, this threshold is €175,000.00. This point is based on the rounded multiplication of 5000 pieces and the average sales price of €35.58. For the sake of a clear overview, in the figure depicting the lower end of the graph, the threshold is set at €40,000.00. Similar patterns can be seen in this analysis. Especially looking at Figure 5, two

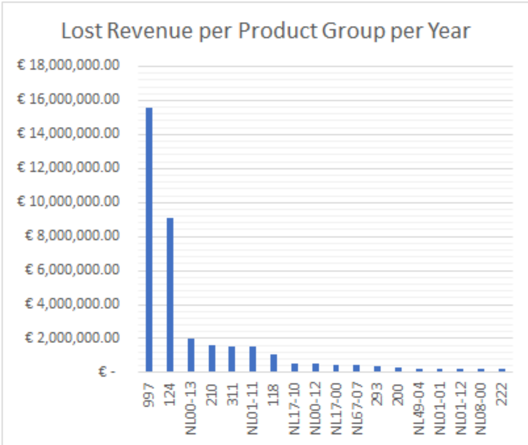

Figure 5: Average Lost Revenue per Product Group Larger than €175,000.00.

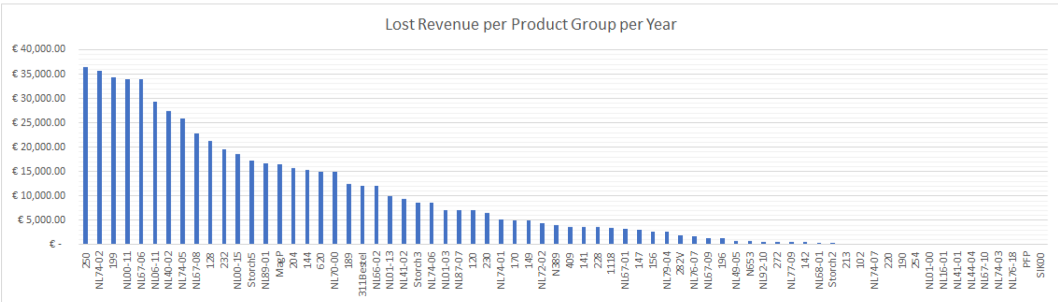

Figure 6: Average Lost Revenue per Product Group Smaller than €40,000.00.

product groups stand out, one of which also ranked high in the previous analysis. This enforces the idea that something is at play other than insufficient stock.

Finally, as it is the main subject of this thesis, a rough estimate of the level of substitution will be computed for the data set to gain insights into the severity of the problem. Due to the fact that the data set only includes data about how many days per specific week a particular product has been out of stock and not about what specific days the product is out of stock, an estimation has to be made. This estimation is based on the work of Karabati et al. (2009), in which they state that, in this case, the determination of stock out days can be approximated by computing the chance that a specific product has no sales for l consecutive days. If this probability is less than a threshold ϵ, which they set at 10^{-4} after numerical experiments, the l

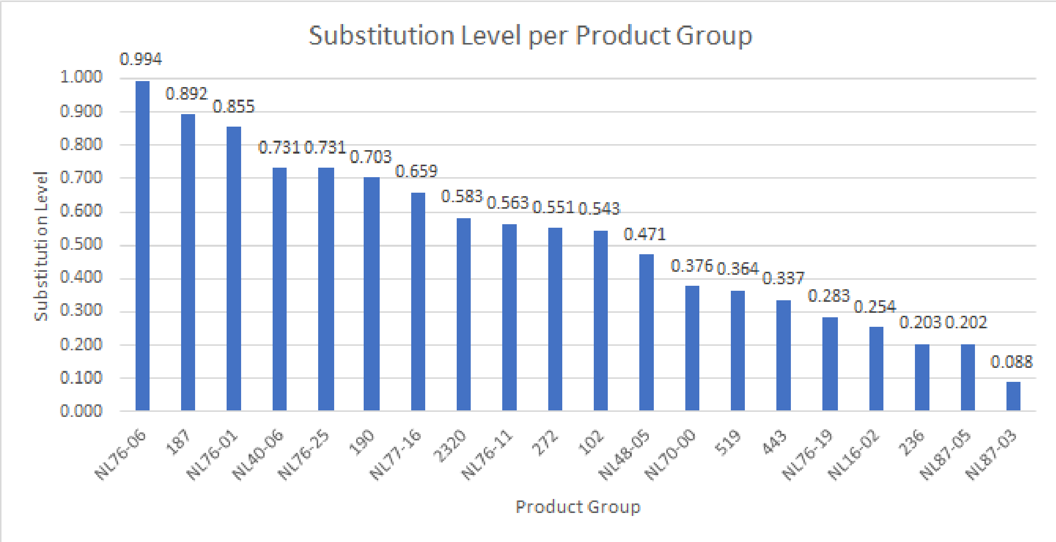

Figure 7: Top 15 Substitution Level per Product Group.

consecutive days will be considered out of stock days. This estimation is done for the entirety of the data set, and at the end, all the days in which a product is not out of stock are automatically in stock. It will be done using the Poisson distribution, just like in the case of Karabati et al. This method will be elaborated on more in Section 4.2.

After this determination, two cases can be determined for every product. On the one hand, the daily average demand of the product group when the product is in stock can be calculated, and on the other hand, this can be done for the situations in which the product is out of stock. Once this has been completed for all products, the level of substitution can be computed by dividing the difference between the demand when product i is in stock and when product i is out of stock by the average daily demand of product i. This division is shown in the equation below.

$$Substitution\ Level_i = \frac{Demand_{In\ stock} - Demand_{Out\ of\ stock}}{\frac{1}{T}\sum_t^T Demand_{i,t}} \tag{1}$$

This value, which is always between 0 and 1, denotes how much of the demand for product i will flow to other products in case of a stock out of product i. This level of substitution is averaged over the product group to give an overview of the product groups that have the most to gain from this research. The top fifteen product groups are shown in Figure 7, and the full figure is shown in the appendix in Figure 39. This exploratory analysis will be used when selecting a suitable product group in Section 5.

2.3 | Conclusions

Concluding, the process of SLIM4 is consists of four primary steps every product goes through: demand pattern classification, forecasting, order level and advice calculation, and replenishment optimisation. Next, the data received from Slimstock's customer is explored. It has been noted that there is a rather significant difference in stock out percentage between the product groups, with an average over the entire assortment of 2.09%. However, it has also been determined that over the past two years, Slimstock's customer has been improving their stock out percentage by 0.43 percentage points each year. On top of that, the financial impact of the stock outs is put into perspective, showing that some product groups are out of stock quite a lot, meaning that there is a significant sum of money the customer misses out on currently. The data exploration is finalised by approximating the substitution levels. The top 15 product groups are shown, and it can be concluded that substitution does exist within the data set provided. However, the effect is not in every product group as large, since there are quite some product groups that are not in the figure presented, meaning they have a substitution level of less than 0.088.

3 | Literature Review

Silver et al. (2016) state that since the mid-1980s, Inventory Management has become increasingly popular. Due to the fact that uncertainty is more often than not a key player in this process, effective decision-making in such a situation requires accurate forecasts of the demand in the future. One of the causes of uncertainty is the phenomenon of substitutability. Substitutability is the phenomenon during which a customer opts for a certain product over another product for any reason. Shin et al. (2015) indicate three different scenarios of substitution mechanisms: assortment-based substitution, inventory-based substitution and price-based substitution. In this framework, the focus is on inventory-based substitution. This concept, alternatively called customer-driven substitution (Yu et al., 2017), describes situations where a stock out of a certain product causes a customer to buy another substitutable product. This will be elaborated on in Section 3.1. Such substitutions can happen fully or partially. Partial substitution happens when during a stock out of product A, some customers choose substitute B, and some leave the store, which can be considered lost sales. This concept will be elaborated on in Section 3.2.

3.1 | Inventory-based Substitutability

Inventory-based substitution occurs when a particular product A is out of stock, forcing the customer to opt for product B (Shin et al., 2015). One of the first studies incorporating this concept into an inventory model is done by McGillivray and Silver (1978). They created a model for dealing with substitutable demand in a case with a periodic review policy, including an order-up-to-level. Later on, this concept was also applied to an EOQ model in a paper by Drezner et al. (1995). Here, contrary to the paper by McGillivray and Silver, three scenarios are evaluated: no, partial, and full substitution. Section 3.2 will elaborate on partial substitution. Both studies conclude that incorporating this may be beneficial as costs will reduce. This notion is also most prominent in the paper by Bassok et al. (1999), later confirmed by the findings of Hsu et al. (2005). Nevertheless, especially Drezner et al. state that this reduction is not always as significant.

However, the aforementioned papers assume one-way substitution, meaning a customer can opt for product B instead of product A but cannot opt for product A instead of product B. This phenomenon is also demonstrated in papers by Rajaram and Tang (2001), Rao et al. (2004), Shah and Avittathur (2007), Dutta and Chakraborty (2010) or Deflem and Van Nieuwenhuyse (2013). Khouja et al. (1996), building on the work by Parlar (1988) and Pasternack and Drezner (1991), continue to change this notion of one-way substitution in their two-item newsvendor problem with substitutability in which they assess two products that can substitute both ways, which is later further developed by Mahajan and Van Ryzin (2001b). The problem of Khouja et al. is then jointly optimised using a Monte Carlo simulation to find the order quantity that maximises profit. Here, they conclude that incorporating substitution into an inventory model positively affects the joint profit. Nevertheless, this is still based on two products, while most retailers have more than that in a particular product group or as a whole. Therefore, later studies expanded these aforementioned concepts to multiple substitutable products. Huang et al. (2011) consider a multi-product competitive newsvendor problem with partial substitution in which they consider n substitutable products held by different decision-makers. This situation is also researched by Mahajan and Van Ryzin (2001a) and Ye (2014). Again, Huang et al. evaluate this problem on the basis of joint profit, which is further explained in Section 3.2.

3.2 | Partial Inventory-based Substitutability

As noted before, the literature not only elaborates on situations with full substitution but also partial substitution is assessed (see, for example, Drezner et al. (1995), Khouja et al. (1996) or Huang et al. (2011)). In all these situations, partial substitution happens when stock-out of the desired product occurs. With a specific substitution parameter α_{ij}, the customer opts for the substitutable product j over product i and with probability $1 - \alpha_{ij}$, the customer leaves the store, resulting in lost sales (Huang et al., 2011). It can be considered that for full substitution α_{ij} equals 1 as, in that case, all customers opt for another product upon a stock-out instead of leaving the store.

Smith and Agrawal (2000) state a model in which they assume unique choice sets of products which contain potential substitutes. Customers only opt for a substitution once a stock-out for the initially desired product occurs. They compute α_{ij} by combining information concerning a products market share with the fraction of time a product is out of stock. Huang et al. (2011) add onto Smith and Agrawal by incorporating α_{ij} into the demand distribution to find the effective demand, which is what they call the demand including the substitution effect based on a newsvendor model. Huang et al. then optimise the order quantity by introducing an iterative algorithm. They start this by initialising the order quantity according to the demand distribution function of product i. This initialisation is used to find a new effective demand and the new parameter for the demand distribution function of product i. With that, the new order quantity is determined and compared to the previous order quantity. The algorithm is stopped if the difference between the two is smaller than or equal to 0.1. Otherwise, the algorithm is started again from the step where the effective demand is computed. For practical reasons, the near-optimal order quantity is then rounded to the nearest integer.

Another approach for accounting for the substitution effect is proposed by Smith and Agrawal (2000). They create a new demand distribution based on the assumption that the demand is negative binomial distributed. They define $g_i(x,m)$ as the probability that the m^{th} arriving customer demands product i, either initially or as a substitute, based on their decision variable x. This variable is a set of binary values that state whether every product should be stocked (1) or not (0). If it should be stocked, they assume a known order-up-to level is set for that product. Contrary to the optimisation goal of Huang et al., Smith and Agrawal strive to find the optimal assortment choices x by using $g_i(x,m)$. This is mathematically defined as the fraction of customers that initially desire product i plus the fraction of customers that chooses product i as a substitute as stated below in Equation 2. Here, the assumption is made that all parameters are known and only x must be decided on.

$$g_i(x,m) = f_i + \sum_{j \neq i} f_j[1 - x_j A_j(x,m)]\alpha_{ji} \tag{2}$$

Here, $A_j(x,m)$ is defined as the probability that product j is in stock for the m^{th} arriving customer given assortment choices x. Although Smith and Agrawal state this can exactly be determined using a discrete, transient Markov process formulation, it is somewhat complex and better approximated using bounds. These bounds are managerially set as Smith and Agrawal point out that $A_j(x,m)$ may not be smaller than the target fill rate r_j and intuitively cannot be larger than 1, resulting in the following bounds.

$$r_i \leq A_j(x,m) \leq 1, \ \forall m. \tag{3}$$

Using this, they devise a method of computing the inventory fill rate including substitution. They continue that by formulating an optimisation problem that optimises the assortment and, with that, the decision whether or not to stock a product based on expected profit maximisation. Smith and Agrawal do not include the stock levels as a separate variable to be optimised. By contrast, they state that once a product is stocked, the order-up-to level is set such that the probability of the demand including substitution being less than or equal to the order level, must be larger than or equal to the target fill rate.

3.2.1 | Substitution parameter

Combining the two most elaborated papers leaves one parameter to be estimated but no approach to do so. Upon further literature research given the topic of estimating

α_{ij}, two promising methods for this estimation are found. First, Karabati et al. (2009) used point-of-sale (POS) data to estimate the substitution parameter. Karabati et al. do this by assigning different inventory states to the POS intervals. The example they give is the situation of a retailer with three products working ten hours a day and seven days a week with POS data every hour. At the beginning of the week, the retailer gets his replenishment, meaning there are 70 POS intervals in the replenishment cycle. For every interval, the inventory substate of the products can be either in stock (1) or out of stock (0). The inventory substates per POS interval together form an inventory state. They move on by defining a formula that declares that the average sales per POS interval of product i should equal the sum of the initial demand of product i, the substitution effect going to product i and an error term. This formula should hold is elaborated on more in Section 4.3.1. This should be the case for every individual inventory state. The error term is essential in the research of Karabati et al., because they continue their algorithm by introducing an optimisation problem that tries to minimise the weighted sum of squared errors coming to estimations for both the initial demand rate λ_i of every product i as well as its substitution probabilities α_{ij} from product i to product j. Later, this concept is built upon further by Honhon et al. (2010), which is then used in a paper by Honhon and Seshadri (2013). However, these papers focus more on assortment optimisation instead of inventory optimisation, like the paper by Smith and Agrawal (2000).

Second, while papers like Talluri and van Ryzin (2004), Kök and Fisher (2007), Vulcano et al. (2009) and Conlon and Mortimer (2013) propose similar approaches, Anupindi et al. (1998) use the Maximum Likelihood Estimator (MLE) for the estimation of the demand and the substitution probabilities. The sample likelihood function is computed by assuming a two-product system with Poisson demand for both product i and j. After computing the log-likelihood function and maximising that, the estimations for both the demand of both products and their substitution parameters can be estimated. A thorough search of the relevant literature yielded no case studies from any of the papers above.

Upon comparing the two methods, Karabati et al. state that both methods perform similarly when the substitution parameter is medium. However, when the parameter gets smaller, the method by Karabati et al. yields better results, and when the parameter increases, the MLE method by Anupindi et al. computes a better estimate with lower error. Therefore, it depends on the situation as to which method is better to be used. Karabati et al. perform this analysis by comparing the Mean Absolute Difference (MAD), Mean Absolute Percentage Error (MAPE) and the maximum absolute error (MAX) over 100 iterations for both methods of computing the values for $\hat{\alpha}_{ij}$ and $\hat{\lambda}_i$.

3.3 | Summary

Summarising, the found literature describes the concept of inventory-based substitution and methods presented to incorporate the concept into inventory management. Mainly three papers stand out and are especially useful for further sections. Firstly, the papers by Karabati et al. and by Anupindi et al. describe how the parameters for α_{ij} and λ_i can be estimated. On top of that, Karabati et al. show ways to evaluate and compare both methods. Secondly, the paper by Smith and Agrawal denotes an approach to use the estimations and incorporate them into inventory management. These methods, although perhaps slightly adjusted to fit this thesis' situation, will be used in Section 4 to form a complete model to be used in Section 5.

4 | Methodology

This section presents the methodology used for answering the earlier stated research question. As discussed earlier, the provided data set does not include accurate daily information about whether a product is in stock or not, which is needed for the parameter estimation discussed in Section 3. Therefore, this section describes an approach of how to approximate these states in Section 4.2. Subsequently, two parameters α_{ij} and λ_i, elaborated on in Section 3, will be estimated using these inventory states. The approach proposed for that can be seen in Section 4.3. After that, the parameters can be used to account for substitution in inventory management. This usage will be presented in Section 4.4. However, before all that is discussed, it is imperative that the mathematical context is discussed. Therefore, this section is started with that in Section 4.1.

4.1 | Mathematical Context

To ensure clarity throughout the explanation of the model, the context of it will be elaborated on. Before going into depth about the mathematics of the solution, the process this model tries to encapsulate will be explained. The model will estimate the substitution within a particular product group. This substitution occurs when a specific product i is out of stock, which causes the customer to buy product j, temporarily increasing demand for product j that initially belonged to product i. This model will estimate the initial demand rate for product i and the proportion of demand of product i that goes to product j upon such a stock out. These estimations will be done for all products, which will be used to determine the demand per product including substitution and discover the near-optimal order-up-to levels using this demand, which will be the model output.

Beginning with the mathematical context, several sets are used in this model. First, there is the set for the products. There are a total of N products, with a minimum of two due to the substitution's core concept, for which a single product is denoted by either an i or a j and $i \neq j$. Second, as input for this model, transaction data is used spanning over T periods. Throughout the model, a single period is a day, denoted by t. Finally, the two methods described in Section 4.3 work with inventory states, elaborated more in detail in that section. The number of possible states found in a data set, denoted by K, depends on the number of products and is equal to 2^N. From the earlier notion that at least two products should be present in the data set, it can be derived that the minimum value of K equals four. The sets used are summarised below.

Products	i = 1, 2, ..., N
	j = 1, 2, ..., N, where i≠j
Periods	t = 1, 2, ..., T
States	k = 1, 2, ..., K

To elaborate on the variables and parameters, the two most significant ones will first be explained. As already noted in Section 3, the substitution parameter is denoted as α_{ij}. This parameter, which indicates the fraction of demand being substituted from product i by product j once product i is out of stock, is bounded between 0 and 1. Furthermore, λ_i indicates the initial demand rate of product i. This initial demand rate is the demand without substitution. In literature, it is also called the base demand, as it is the base upon which the demand including substitution is built.

At last, the assumptions are denoted. The first assumption is that the initial demand rate for the products is stationary and Poisson distributed. The transaction data used as input for the model is that of products sold in physical retail stores. Therefore, it can be assumed that the products are relatively slow movers, which provides cause to assume a Poisson distribution (Boylan et al., 2008). Next to that, the existence of partial substitution within the product groups is assumed. This partial substitution is not only one-way but asymmetrically two-way, meaning that product i can substitute for product j and vice versa. However, these two substitution parameters do not need to be the same. It is more logical that these are not the same due to stochasticity in demand. Finally, an (R, S)-model is assumed with positive lead times and review periods.

4.2 | Inventory State Approximation

Remember from Section 4.1 that the data used does not contain accurate inventory levels or daily information about whether or not a product is in stock. However, that is a vital part of the parameter estimation as explained in Section 4.3. To solve this, the inventory substates, stating in a binary fashion whether or not a product is in stock, can statistically be approximated. This approximation is made by determining the average sales of all nonzero periods. Using these, the probability of l consecutive periods with zero sales is determined using a Poisson distribution assuming the product is in stock. If this probability falls below a threshold of 10^{-4}, it is rather unlikely that the product is indeed in stock, meaning that it can be assumed that the product has been out of stock for the past l consecutive periods. Take the example sales in Table 1. The average sales of all nonzero periods is 4 in this case. In case period 5 and 6 did not exist, there would only be one period with zero sales. The

Period	Sales
1	4
2	3
3	5
4	0
5	0
6	0

Table 1: Example Sales per Period

probability of that happening is Poisson distributed with a rate of 4, which equals 0.018. This is not lower than the threshold, so the assumption that the product is in stock holds. Now, period 5 is added. The probability of two periods with zero sales is Poisson distributed with a rate of 8, because in two periods you would expect $2 * 4 = 8$ items sold. This probability equals 0.00035, meaning it is less likely to happen assuming the product is still in stock, but still not lower than the threshold. However, adding period 6 changes that. The probability of this period also having zero sales is Poisson distributed with a rate of 12, which equals 0.000006. This is lower than the threshold of 10^{-4}, meaning that three consecutive periods having zero sales assuming the product is still in stock is too unlikely, causing the three periods to be considered as stock outs. This constitutes Table 2.

Period	Sales	Inventory substate
1	4	1
2	3	1
3	5	1
4	0	0
5	0	0
6	0	0

Table 2: Example Sales per Period with Inventory Substates

4.3 | Parameter Estimation

As noted in Section 4.1, two parameters need to be estimated. The first parameter is the substitution parameter, also called the substitution rate or probability in literature, α_{ij}. The second is the initial demand rate, indicated by λ_i. The estimation to be done will be based on the literature found in Section 3 of Karabati et al. and Anupindi et al., which is more elaborately explained in Sections 4.3.1 and 4.3.2

respectively. These estimations can then be used to account for the substitution in Section 4.4.

4.3.1 | Parameter Estimation by Minimising Weighted Error Terms

The first approach of estimating the parameters is based on one proposed by Karabati et al. (2009). They use POS data as input for an optimisation problem that minimises the weighted squared error. To do this, the input data is, first, grouped into predetermined POS intervals. The smallest amount of time for which sufficient available data exists is chosen for the size of the intervals. In the case of this model, that is a day. Next to that, in the case by Karabati et al., the inventory levels at the beginning of every POS interval are known. Suppose the data per interval look as in Table 3. This data is used as input to assign binary inventory substates, δ_i with

POS interval	**Inventory product i**	**Sales in Pieces of Product i**	**Inventory product j**	**Sales in Pieces of Product j**
1	15	4	16	1
2	11	3	15	2
3	8	2	13	3
4	6	3	10	2
5	3	3	8	1
6	0	0	7	2
7	0	0	5	3
8	0	0	2	2
9	0	0	0	0

Table 3: Example Data per POS Interval.

$\delta_i = 1$ if the inventory is positive and $\delta_i = 0$ otherwise, to the different products per POS interval, which will be utilised to assign an inventory state, I, to every POS interval. The inventory substates indicate whether the product is in stock or not, and all inventory substates in the product group together form the inventory state per POS interval. For Table 3, this would be as shown in Table 4. Note that not every possible combination is found in this example, which is also possible in actual data. Three out of four possible inventory states are prevalent, while the state (1, 0) is not. Nevertheless, this approach can still find an estimation for most of the parameters. The parameter which cannot be found in this case is $\alpha_{(1,0)}$ as there are no POS intervals during which product j is out of stock and product i is in stock. However, in order to continue the method, $\alpha_{(1,0)} = 0$ is assumed. This must later be updated when

POS interval	δ_i	δ_j	I_k
1	1	1	(1, 1)
2	1	1	(1, 1)
3	1	1	(1, 1)
4	1	1	(1, 1)
5	1	1	(1, 1)
6	0	1	(0, 1)
7	0	1	(0, 1)
8	0	1	(0, 1)
9	0	0	(0, 0)

Table 4: Example Inventory Substates and State Assignment.

information is available about this inventory state. This is continued by computing the average sales per inventory state, $\bar{s}_{i,I_k}$. Applying that to our example data creates Table 5 with total as well as average sales per inventory state. Karabati et al. use

I_k	$\bar{s}_{i,I_k}$	$\bar{s}_{j,I_k}$
(1, 1)	3	1.6
(0, 1)	0	2.667
(0, 0)	0	0

Table 5: Example Average Sales per Inventory State.

this to find the values of α_{ij} and λ_i according to the following formula. Here, e_{i,I_k} concerns an error term used to compensate for the stochasticity in demand.

$$\bar{s}_{i,I_k} = \hat{\lambda}_i + \sum_{j \neq i} \hat{\alpha}_{ji}\hat{\lambda}_i(1 - \delta_{j,I_k}) + e_{i,I_k} \tag{4}$$

This formula has a rather simple setup, meaning that the demand recorded should be equal to the sum of the initial demand rate, which excludes any substitution effects, and the demand that flows from products j to product i as a result of a stock out, being the substitution effect. The latter is built up as the substitution parameter of product j to i multiplied by the initial demand of product j and δ_{j,I_k}, the inventory substate of product j. This is summed for all products for which $j \neq i$. This formula ends with the error term, which denotes the deviation from the estimation and the recorded average sales per inventory state, $\bar{s}_{i,I_k}$. The goal is to minimise the weighted squared error terms, for which the weights, W_{I_k}, are based on the occurrence of the states in the data set. If this is applied to the earlier example, it can be seen that (1, 1) occurs five times, giving it the weights of 5. Applying that to the other states would also provide the weights of 3 and 1 for states (0, 1) and (0, 0) respectively. Combining this all will provide the following optimisation problem based on Karabati et al., which yields estimates for all parameters α_{ij} and λ_i.

$$\min \sum_{i \in N} \sum_{k \in K} W_{I_k} e_{i,I_k}^2, \tag{5}$$

subject to

$$\hat{\lambda}_i + \sum_{j \neq i} \hat{\alpha}_{ji} \hat{\lambda}_i (1 - \delta_{j,I_k}) + e_{i,I_k} = \bar{s}_{i,I_k}, \ \forall k \in K \text{ and } \forall i, \ j \in N \text{ if } i \neq j, \tag{6}$$

$$\sum_{j \in N} \hat{\alpha}_{ij} \leq 1, \ \forall i \in N, \tag{7}$$

$$\hat{\lambda}_i \geq 0, \ \forall i \in N, \tag{8}$$

$$\hat{\alpha}_{ij} \geq 0, \ \forall i, \ j \in N \tag{9}$$

4.3.2 | Parameter Estimation using MLE

Another estimation method is by using MLE, which is applied by Anupindi et al. (1998). They apply this statistical method of estimating parameters to the specific case of partial substitution. This estimation is started by determining the different states the product group can be in. This is done similarly to the approach of Karabati et al. with the difference that Anupindi et al. do not look at the average sales per state, but at the total sales per inventory state, $\bar{s}_{i,I_k} W_{I_k}$. When the previous section's example is revisited, these values can be seen in Table 6. Moreover, similar

I_k	$\bar{s}_{i,I_k} W_{I_k}$	$\bar{s}_{j,I_k} W_{I_k}$
(1, 1)	15	8
(0, 1)	0	8
(0, 0)	0	0

Table 6: Example Total Sales per Inventory State

to the approach of Karabati et al., it is essential to know something about the prevalence of the inventory states. This is done the same as in the method explained before, using W_{I_k}. Next, the likelihood functions are determined for every inventory state, which are multiplied to find the overall likelihood function for all states (see Anupindi et al. (1998) for the exact derivation). If the logarithm of this is taken, these multiplications change to additions and the following formula is determined for the example mentioned earlier.

$$\log \mathcal{L} = \sum_{k \in K} \sum_{i \in N} \Big((\bar{s}_{i,I_k} W_{I_k}) \log \hat{\Lambda}_{i,I_k} - W_{I_k} \hat{\Lambda}_{i,I_k} \Big) \tag{10}$$

Note that $\hat{\Lambda}_{i,I_k}$ is not the same as $\hat{\lambda}_{i,I_k}$ in Equation 6. In the model by Anupindi et al., $\hat{\Lambda}_{i,I_k}$ is the estimator of the demand including substitution for product i per inventory state, while Karabati et al. denote $\hat{\lambda}_{i,I_k}$ as the initial demand rate of product i per inventory state. This subtle detail leads to the further elaboration of the optimisation problem. Due to the core concept of substitution, the substituted demand of product j to product i can never be larger than the demand of product j alone. From that, the constraints are set up. If this is applied to the example situation, it constitutes the following.

$$\hat{\Lambda}_{j,(1,1)} \leq \hat{\Lambda}_{j,(0,1)} \leq \hat{\Lambda}_{i,(1,1)} + \hat{\Lambda}_{j,(1,1)} \tag{11}$$

Typically, there would also be such a constraint for product i, but since, in the example, there is no POS interval in which product j is out of stock while product i is, this cannot be set up. Therefore, as mentioned earlier, the substitution effect from product j to i cannot be determined in this example. Finally, completing this optimisation problem for this example will provide the following yielding estimates for all Λ_{i,I_k}.

$$\max_{i \in N,\ k \in K} \log \mathcal{L} \tag{12}$$

subject to

$$\hat{\Lambda}_{j,(1,1)} \leq \hat{\Lambda}_{j,(0,1)} \leq \hat{\Lambda}_{i,(1,1)} + \hat{\Lambda}_{j,(1,1)},\ \forall i,\ j \in N \text{ if } i \neq j \tag{13}$$

$$\hat{\Lambda}_{i,I_k} \geq 0,\ \forall i \in N \tag{14}$$

This problem will, by maximising the log-likelihood, provide estimates for $\hat{\Lambda}_{i,I_k}$ for all products and inventory states. However, these are still not estimations for α_{ij} and λ_i. For this, an additional calculation has to be done. The initial demand rate can be taken from the inventory state where all products are in stock, as no substitution is possible. The computation of α_{ij} requires more work. The demand flowing from product i to product j upon a stock out of product i can be computed from the estimations. If that is then divided by the initial demand rate of product i, it results in $\hat{\alpha}_{ij}$ as shown below.

$$\hat{\alpha}_{ij} = \frac{\hat{\Lambda}_{j,(0,1)} - \hat{\Lambda}_{j,(1,1)}}{\hat{\Lambda}_{i,(1,1)}} \tag{15}$$

4.3.3 | Performance and Method Selection

Due to the fact that not one of these two methods is more accurate in every case (Karabati et al., 2009), they are both implemented to estimate the substitution parameter. This results in a set of α_{ij} and λ_i per product per method. The only decision left is which method to use for the product group based on performance measures. Karabati et al. made the same evaluation based on MAD, MAPE and MAX. The formulas for these measures are given in Equations 16, 17 and 18 respectively.

$$MAD_i = \frac{1}{T}\sum_{t=1}^{T}|D_{i,t} - \hat{\Lambda}'_{i,t}| \tag{16}$$

$$MAPE_i = \frac{1}{T}\sum_{t=1}^{T}\frac{|D_{i,t} - \hat{\Lambda}'_{i,t}|}{|\hat{\Lambda}'_{i,t}|} * 100 \tag{17}$$

$$MAX_i = \max_{t \in T}|D_{i,t} - \hat{\Lambda}'_{i,t}| \tag{18}$$

Apart from parameters and sets earlier elaborated on in Section 4.1, two are yet to be explained. To start, $D_{i,t}$ is used. This is the demand of product i in period t as given by the input data. Furthermore, $\hat{\Lambda}'_{i,t}$ can be found in the formulas above. Note that this looks similar to the earlier mentioned $\hat{\Lambda}_{i,I_k}$ because it has a similar definition. $\hat{\Lambda}'_{i,t}$ denotes the estimated demand including substitution for product i in period t. Product i is either in stock or out of stock every period, denoted by the binary parameter $\delta_{i,t}$ in the formula below.

$$\hat{\Lambda}'_{i,t} = \hat{\lambda}_i + \sum_{i \neq j}\hat{\alpha}_{ji}\hat{\lambda}_j(1 - \delta_{j,t}) \tag{19}$$

After the performance measures are computed for all products and the two methods, the method selection phase comes into effect. Suppose the majority of the measures are lower for the method by Karabati et al. for a particular product. In that case, the parameters computed by that method are accepted, while the others are rejected, and vice versa. However, in the unlikely scenario that one or three measures have the same value, it does not matter which method is chosen as they both perform equally, and one of the two is selected randomly. This results in the final set of parameters per product, which can be used for the procedure described in Section 4.4.

4.4 | Substitution Incorporation in Inventory Model

First, the change in demand distribution when accounting for partial substitution will be addressed based on the work of Smith and Agrawal (2000) in Section 4.4.1. This will then be continued by determining the near-optimal order-up-to level per product based on the fill rate and the inventory value, denoted in Section 4.4.2.

4.4.1 | Demand Distribution with Partial Substitution

A new demand distribution should be established to estimate the stochasticity of the newly determined demand including substitution accurately. Smith and Agrawal (2000) do this by determining distributions for two occurring phenomena: the number of customers arriving at the store and the number of items these customers demand.

Combining the two constitutes the demand distribution without the substitution effect, being the initial demand rate. However, using Section 4.3, the initial demand rate is already known. The parameter λ_i estimated in that section is the initial demand rate based on a Poisson distribution. That would make the probability mass function of the initial demand rate the following.

$$\psi(d_i|\lambda_i) = \frac{\lambda_i^{d_i} e^{\lambda_i}}{d_i!} \tag{20}$$

In this equation, d_i equals the demand for product i and λ_i is the initial demand rate of product i estimated in Section 4.3. Smith and Agrawal continue by determining a probability distribution for $\bar{D}_i$, the demand of product i including substitution. This distribution is not necessarily the same as $\psi(d_i|\lambda_i)$ as the former incorporates the possible substitution effect due to stock outs of other products. To correctly do that, the following adjustment has to be made to the parameter of the Poisson distribution to incorporate this. This is denoted by Λ_i below.

$$\Lambda_i = \lambda_i + \sum_{j \neq i} \lambda_j [1 - A_j(q_j)] \alpha_{ji} \tag{21}$$

It can be noted that this equation is split into two parts. On the one hand, there is the initial demand rate for the product itself. On the other hand, the substitution effect from all other products in the product group to product i is denoted. For this, three parameters are used: λ_j, $A_j(q_j)$ and α_{ji}. λ_j and α_{ji} are already explained in Section 4.3, which leaves $A_j(q_j)$. This parameter denotes the probability that product j will be in stock given a certain expected inventory level q_j at the beginning of the cycle. To determine this value, the order-up-to level cannot simply be chosen as there is both a positive lead time and review time for all products, meaning the demand during these periods should be taken into account. Therefore, to approximate $A_j(q_j)$, the Poisson probability of the demand during $R+L$ periods being smaller than or equal to the order-up-to level must be computed and can be used for $A_j(q_j)$. Please note that this approximation is too pessimistic as the time between the reaching of the order-up-to level and the arrival of the order is most likely shorter than $R+L$, meaning more demand is assumed causing the approximation to be lower. This is because once the inventory position is reviewed and it is less than the order-up-to level, this happened somewhere between the previous review moment and the current one. The actual demand during that period of time is called the undershoot, but the probabilistic implementation of this concept falls outside of the scope of this thesis. Next to that, this approach uses the initial demand rate of product i when computing $A_j(q_j)$, but in reality, there is also a substitution effect to be taken into account, which can only be achieved using Equation 21, meaning it should be computed iteratively. To do this, $A_j(q_j)$ should be computed first with λ_i. This is then put in Equation 21 to calculate Λ_i, which can be used to compute $A_j(q_j)$. This will eventually converge to the real value of Λ_i and $A_j(q_j)$ given the estimates of λ_i and α_{ij}. However, this has not be implemented into the model due to it being outside the scope of this research. Because of this, a likely lower demand rate is considered, causing $A_j(q_j)$ to be higher than reality. Therefore, the too optimistic approximation of $A_j(q_j)$ is determined using $R+L$ periods and an initial demand rate.

Combining this all results in a demand distribution that includes the partial substitution effect. This formula, stated below, can later be used to compute the fill rate, which will be utilised in the order-up-to level determination.

$$\psi(d_i|\lambda_i,...,\lambda_N,\ q_1,...,q_N) = \frac{\Lambda_i^{d_i} e^{\Lambda_i}}{d_i!} \tag{22}$$
$$\text{where } \Lambda_i = \lambda_i + \sum_{j\neq i} \lambda_j A_j(q_j)\alpha_{ji}$$

4.4.2 | Order-up-to Level Determination

To determine the near-optimal order-up-to level when accounting for partial substitution, a heuristic, explained in Section 4.4.2.3, is used based on two performance measurements. First, the fill rate of a particular setup of order-up-to levels is determined. This is elaborated on in Section 4.4.2.1. Second, the inventory value is computed in Section 4.4.2.2.

4.4.2.1 | Fill Rate Calculation

The fill rate resulting from a particular set of order-up-to levels is mainly based on the demand distribution, including the substitution effect, denoted in Equation 22. Although Smith and Agrawal try to optimise the decision whether or not to stock a product based on a specific set of optimal order-up-to levels, this stocking decision has already been made in the situation of this research. However, the set of near-optimal order-up-to levels has not been determined yet. Therefore, the methodology and equations proposed by Smith and Agrawal are adapted slightly to fit the current optimisation goal as they focus on which products to stock given a set of order-up-to levels. However, this section aims to find order-up-to levels, given that all products are actually stocked. This constitutes the formula for the fill rate ρ_i, the probability that a customer either finds their desired product or a substitute in stock, given a specific set of inventory levels at the beginning of the cycle q for the entire product group.

$$\rho_i(q) = \Psi_i(q_i|\lambda_i,...,\lambda_N,\ q_1,...,q_N) + q_i \sum_{d=q_i+1}^{\infty} \frac{\psi(d|\lambda_i,...,\lambda_N,\ q_1,...,q_N)}{d} \tag{23}$$

In this equation, $\Psi_i(q_i|\lambda_i, ..., \lambda_N,\ q_1, ..., q_N)$ is the cumulative demand distribution of product i with substitution given the initial demand rates and a set of all expected inventory levels at the beginning of the cycle. However, this equation only denotes the fill rate per product, but the weighted average should be taken for the entire product group. The weights of the fill rates are based on the Λ_i computed using Equation 22 by dividing Λ_i by the sum of Λ in the entire product group. This equation is used in the heuristic explained in Section 4.4.2.3.

$$\rho(q) = \sum_i \frac{\Lambda_i}{\sum_j \Lambda_j} \rho_i(q) \tag{24}$$

4.4.2.2 | Inventory Value Computation

The inventory value in this instance will be computed for an (R, S)-model, which is consistent with the model currently used by Slimstock as stated in Section 2.1. This can be done by multiplying the expected on-hand inventory by the purchasing price of a particular product. From this, the following formula arises.

$$V_{i,\ inventory} = E_i(OH) * p_{i,purchasing} \tag{25}$$

Silver et al. (2016) provide a formula for computing the expected on-hand inventory given a certain distribution for an (R, S)-model. This formula comprises the minimum required inventory level and the average order quantity in stock. Assuming that the inventory decreases linearly over time during a cycle, which holds on average, the latter can be denoted by $\frac{\Lambda_i * R}{2}$. This is the demand during the review period divided by two. For the minimum required inventory level given a certain order-up-to level S_i, this can be denoted by $S_i - \Lambda_i * (R+L)$. Silver et al. use $\hat{x}_{R+L}$ and $\hat{D}_R$ instead of $\Lambda_i * (R+L)$ and $\Lambda_i * R$ respectively for the average demand during the $R+L$ and R period, but for consistency regarding demand including substitution, Λ_i is used here. This constitutes the following formula.

$$E_i(OH) = S_i - \Lambda_i * (R+L) + \frac{\Lambda_i * R}{2} \tag{26}$$

Combining this with Equation 25, yields the following formula to be used in Section 4.4.2.3 for scenario comparison.

$$V_{i,\ inventory} = \left(S_i - \Lambda_i * (R+L) + \frac{\Lambda_i * R}{2} \right) * p_{i,purchasing} \tag{27}$$

4.4.2.3 | Iterative Order-up-to Level Optimisation

Taking into account Sections 4.4.2.1 and 4.4.2.2, an iterative heuristic approach can be employed to obtain a set of order-up-to levels that are (nearly) optimal. The use of a heuristic method is preferred due to the scale of the problem. While it might be feasible to examine the set of possible solutions individually and determine the optimal set for one or two products, the solution space becomes undesirably large for more than that. If one attempts to assess all possible solutions, the computation time required will increase exponentially, making solving the problem in a reasonable amount of time impractical.

Therefore, for this specific optimisation, a greedy heuristic will be used. A greedy heuristic is a type of approach that starts from a particular starting point and, from that point, evaluates all neighbouring states towards a complete solution. The state that yields the most significant benefit will be chosen, and the new state is set as the starting point. This cycle continues until a stopping criterion is met. Applying that to the problem at hand comes to the following. The heuristic starts with a state in which all order-up-to levels are zero. The possible steps to be taken towards a complete solution are if the order-up-to level of a particular product is increased by one. Therefore, a possible step is that all order-up-to levels are kept at zero, except for one, which is increased by one. Then, the change in fill rate and inventory value of the entire assortment are computed. These measurements are then combined to form the change in fill rate per extra euro. After that, the order-up-to level is again decreased by one, and the algorithm evaluates the same for the next product. After all products are evaluated, the algorithm selects the product with the most significant change in fill rate per extra euro. It increases that order-up-to level by one for good. This procedure is repeated until the stopping criterion of reaching an overall target fill rate is reached. This all constitutes the following pseudo-code for the greedy heuristic. This algorithm eventually returns the near-optimal set of order-up-to levels that reaches the minimum target fill rate. If the inventory value could be neglected, it would be beneficial to increase the order-up-to levels even further as the fill rate will then increase. However, the inventory value cannot be neglected, meaning the set that just reaches the target fill rate will be best.

4.5 | Summary

Altogether, this section explains the method used for estimating the parameters for substitution and the initial demand rate and the usage of these estimations for determining the order-up-to levels. These parameters are firstly estimated using the method proposed by Karabati et al. (2009). They state that POS data can be used to determine these parameters by optimising the problem formulated at Equation 5. However, this is not the only method used to estimate the parameters. Anupindi et al. (1998) proposed another approach using MLE, which is stated in Equation 12. Maximising the log-likelihood function in Equation 10 yielded the lambdas for all

Algorithm 1 Iterative Order-up-to Level Greedy Heuristic.

$OrderUpToLevel[i] \leftarrow 0,\ \forall i \in Products$
while fill rate < target fill rate **do**
  for $i \in Products$ **do**
    $Oldq \leftarrow Updateq$
    $OldFillRate \leftarrow EvaluateFillRate$
    $OldInventoryValue \leftarrow EvaluateInventoryValue$
    $OrderUpToLevel[i] \leftarrow OrderUpToLevel[i] + 1$
    $Newq \leftarrow Updateq$
    $NewFillRate \leftarrow EvaluateFillRate$
    $NewInventoryValue \leftarrow EvaluateInventoryValue$
    $StateResults[i] \leftarrow \frac{NewFillRate - OldFillRate}{NewInventoryValue - OldInventoryValue}$
    $OrderUpToLevel[i] \leftarrow OrderUpToLevel[i] - 1$
  end for
  $ProductIncrease \leftarrow max_i(StateResults)$
  $Orderlevel[ProductIncrease] \leftarrow OrderUpToLevel[ProductIncrease] + 1$
end while

products and all inventory states. From these results, the substitution parameters could be computed by using equations like 15.

The results from these two methods were evaluated based on MAD, MAPE and MAX, of which the equations are denoted in Equations 16, 17, 18 respectively. The method with the best performance measurements would be used as the best suitable method for estimation for that group of products, which is then further used in Section 4.4.

This section continues by describing the approach of determining the near-optimal set of order-up-to levels. For this, the fill rate and the inventory value are required to be used in a greedy heuristic. Regarding the fill rate, this is computed based on the demand distribution explained in Section 4.4.1. This distribution depends on the target fill rate as a lower bound and accounts for the partial substitution effects. This is then used in Equation 23 to find the fill rate per product. The weighted average is then taken according to Equation 24 to find the product group fill rate. Next to that, the inventory value is computed with Equation 27. This formula consists of the average on-hand inventory during a cycle and the costs of purchasing an item.

Finally, the greedy heuristic explained in Section 4.4.2.3 is used. This heuristic states that in every iteration, all order-up-to levels of the products in the assortment are sequentially increased by one, for which the change in fill rate and the change in inventory value are determined. After that, the change in fill rate per euro is determined and saved. The product with the most significant change in fill rate per euro will have its order-up-to level increased by one permanently. This is repeated until the overall fill rate reaches a predetermined target fill rate.

5 | Results

This section presents the results of this research. It applies the complete model created in Sections 4.3 and 4.4 to a specific product group of the data set. Then, several analyses are performed using theoretical and real-life data in Section 5.2. These results are mainly evaluated based on performance measures of the model as also used in Section 4.3.3.

5.1 | Method

Regarding the approach for the results of this research, it is done in two phases. First, because the parameters α_{ij} and λ_i are unknown, the model's actual performance cannot be determined by simply analysing real-life data. Therefore, theoretical data generated with known parameters will be used first to assess the validity of the results. Several analyses will be performed on this generated data, and because the input is known, the performance can be determined correctly. First, an accuracy analysis of the inventory state approximation is performed. Due to the fact that this is a rather vital part of the model created and some precision can be lost due to false determinations, it is crucial that these stock outs are accurate.

Second, the accuracy of the model's ability to estimate several substitution parameters will be determined. For this, a product group with two products will be used for which the demand will be stochastically level. However, several experiments will be run for which the substitution parameter will vary from 0 up to and including 1 both ways. These will provide insight into how the accuracy develops with changing substitution parameters. They will also be performed with varying demand, as the hypothesis is that higher demand causes the difference between stochasticity in demand and substitution effect to be more apparent for the model as there is a more considerable absolute contrast.

Next, the added benefit of accounting for substitution compared to not accounting for it will be determined. This will again be done for several substitution parameter values, but the substitution only goes one way in this analysis. This adjustment is made for a more straightforward algorithm behaviour overview. Finally, the created model is run with real data obtained from Slimstock's customer. The result is compared to the original order-up-to levels to portray the added benefit for a selected product group.

5.2 | Inventory Model with Partial Substitution

In this section, both the results of the theoretical as well as real data will be presented. Some initial summaries will be made, but actual conclusions can be found in Section 6.

5.2.1 | Theoretical Data

Several data sets are constructed for which known values of α_{ij}, λ_i, T and the stock out percentage (SOP) are used. These data sets are utilised for the analyses performed on theoretical data. For these analyses, the model is divided into three main parts. First, the accuracy of the inventory state approximation is studied, followed by the evaluation of the parameter estimation performance. Finally, the effect of accounting for substitution in the order-up-to levels will be determined.

5.2.1.1 | Inventory State Approximation

First, the accuracy of the inventory state approximation will be looked into. Recall that since no factual information is known about whether certain products are in or out of stock per day, an estimation is made based on the POS data and the assumed Poisson distribution. The entire reasoning behind this can be found in Section 2.2.

For λ_i varying from 5 to 5000 items per day, data sets are created with different substitution parameters. The higher values of this range are unrealistic in this setting but are nevertheless chosen to validate the method. The substitution parameters differ from 0 to 1 with increments of 0.05. To do this, the inventory states are created randomly with a SOP of 20%. This SOP indicates the percentage of periods an arbitrary product is out of stock. After this, sales data is generated using the Poisson distribution that includes the substitution effect based on the λ_i and α_{ij} set for the data set. Once the set is created, the inventory approximation method is applied independently from the actual inventory states. These two sets of inventory states are compared to find the number of wrongly classified periods. Once performed for all data sets, the results are aggregated and shown in Figures 8 and 9.

As can be seen, the number of wrongly classified periods is rather level for all α_{ij}, with an average of 0.64 periods, suggesting that α_{ij} does not significantly affect the accuracy of the inventory state determination. The λ_i, however, gives an entirely different view. It can be noted that the inaccuracy starts relatively high at around 40 periods wrongly classified at a low λ_i. However, it does decrease quickly. Excluding the high value at a low λ_i and zooming in, results in Figures 10 and 11.

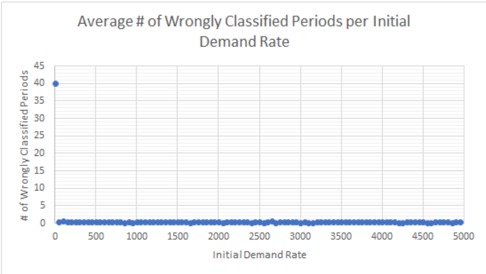

Figure 8: Average # of Wrongly Classified Periods per λ_i. ($T = 200$, $SOP = 20\%$)

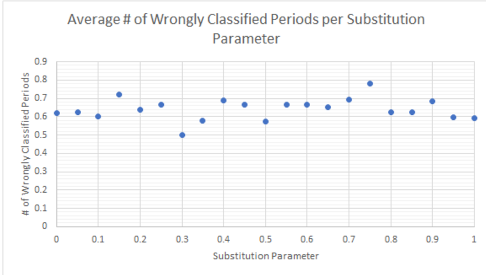

Figure 9: Average # of Wrongly Classified Periods per α_{ij}. ($T = 200$, $SOP = 20\%$)

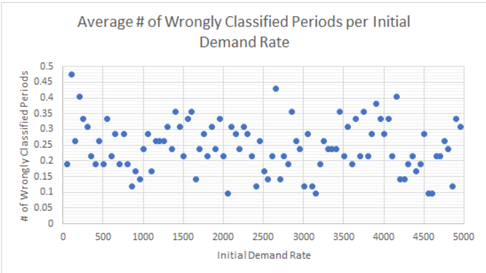

Figure 10: Average # of Wrongly Classified Periods per λ_i Excluding Low λ_i. ($T = 200$, $SOP = 20\%$)

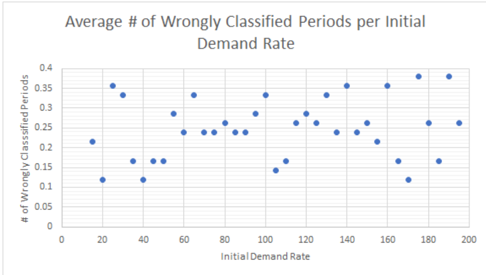

Figure 11: Average # of Wrongly Classified Periods per λ_i Excluding Low λ_i and Zoomed in. ($T = 200$, $SOP = 20\%$)

In both figures, the data points seem somewhat randomly level, indicating that once λ_i surpasses a certain level, an increase in λ_i does not necessarily lead to a change in the accuracy. Both have an average number of 0.24 wrongly classified periods. In reality, lower values are more likely to occur, creating a cause to also zoom in on that part of the range. The results of this analysis can be found in Figure 12. From this figure, it can be seen that it is relatively high at around 40 wrongly classified periods. This is not illogical as that is 20% of T, precisely the SOP mentioned earlier. After $\lambda_i = 10$, it decreases and settles at about 0 for $\lambda_i = 12$. Therefore, it is safe to assume that a minimum initial demand rate of 12 is needed for the classification to be accurate.

This same analysis has been made concerning the SOP to see if that value affects the inaccuracy of the inventory state determination. For this, the SOP is varied between 0% and 100% with increments of 2.5%. This allows insight into whether this percentage's bounds exist to accurately determine the inventory states for further analysis. The results can be seen in Figure 13.

The SOP seems to have a positive effect on the inaccuracy. However, with higher values of this percentage, it is significantly higher than for the earlier analyses. In

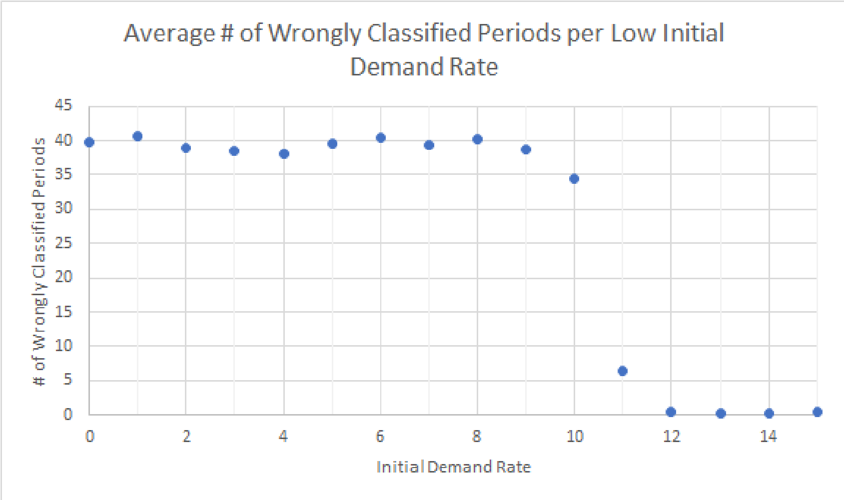

Figure 12: Average # of Wrongly Classified Periods for Low λ_i. ($T = 200$, $SOP = 20\%$)

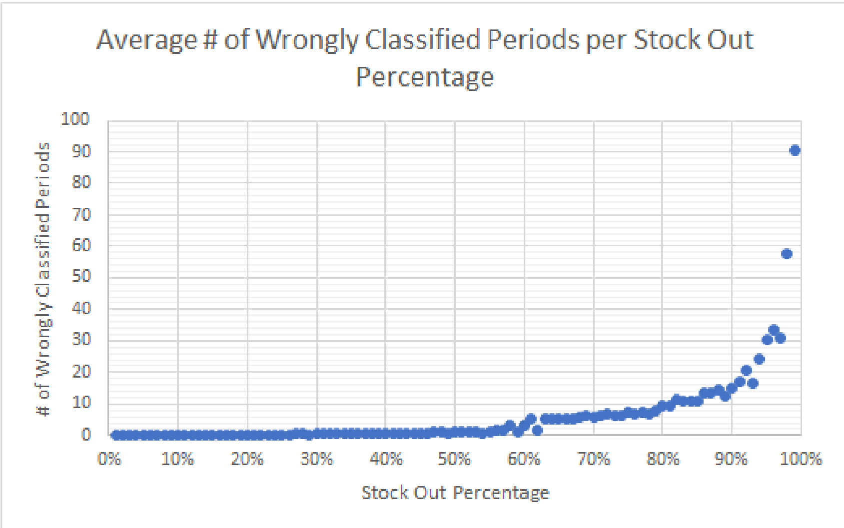

Figure 13: Average # of Wrongly Classified Periods per Stock Out Percentage. ($\lambda_i = 15$, $T = 200$)

order to achieve adequate performance, Figure 13 along with a zoomed-in version denoted in Figure 14 show that a maximum SOP of 55% may be present as the inaccuracy increases aggressively after that. This percentage is excessive in reality, meaning that the SOPs are likely lower than that.

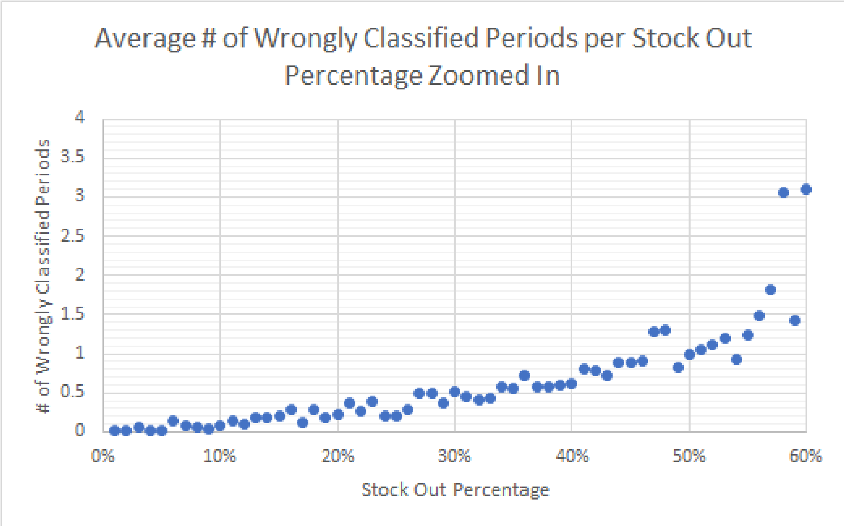

Figure 14: Average # of Wrongly Classified Periods per Stock Our Percentage Zoomed In. ($\lambda_i = 15$, $T = 200$)

5.2.1.2 | Parameter Estimation

Subsequently, the performance of the model regarding the estimations of α_{ij} and λ_i is determined. This analysis is performed by generating several data sets using known values of α_{ij} and λ_i. Then, these sets are analysed by the model, assuming the inventory states are known for the purpose of solely measuring this method's performance, and of the resulting $\hat{\alpha}_{ij}$ and $\hat{\lambda}_i$ the MAD, MAPE and MAX are determined. This is performed for two products for which only the values for α_{ij} and α_{ji} are changed per data set, which are not necessarily the same. The results are then averaged per α_{ij} and shown in Figures 15, 16 and 17 respectively. These figures

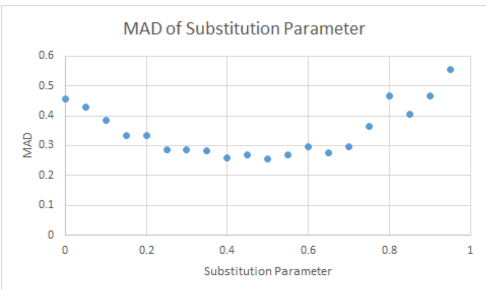

Figure 15: MAD per Substitution Parameter. ($\lambda_i = 15$, $T = 200$, $SOP = 25\%$)

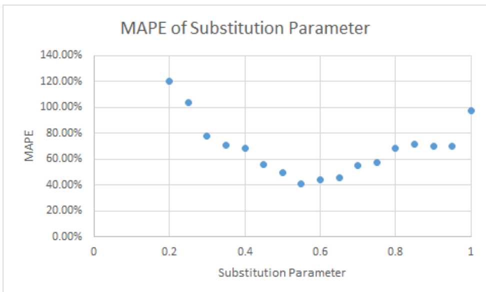

Figure 16: MAPE per Substitution Parameter. ($\lambda_i = 15$, $T = 200$, $SOP = 25\%$)

show that overall the performance measures are rather high, which indicates that the model is unable to estimate the substitution parameters accurately. On top of that,

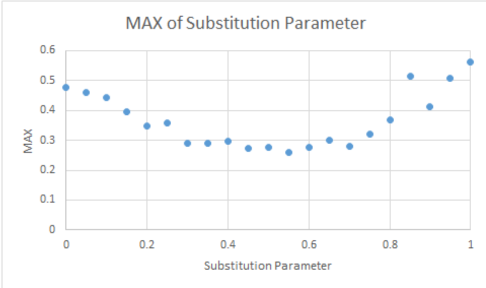

Figure 17: MAX per Substitution Parameter. ($\lambda_i = 15$, $T = 200$, $SOP = 25\%$)

the performance is worse when the substitution parameter is below 0.25, after which an increase in the substitution parameter does not influence the MAPE, unlike the MAD and MAX. The latter two follow a similar pattern with worse performance for low and high values of α_{ij}. This bad performance is likely due to an insufficient number of periods. Therefore, another experiment is run with increasing T and a fixed substitution parameter of 0.5 to see if the performance improves. The results can be seen in Figures 21, 22 and 23. This is also done for the performance of λ_i and can be seen in Figures **??**, **??** and **??**. It can be seen in the figures that the

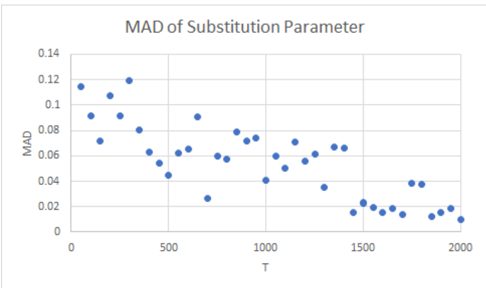

Figure 18: MAD for Varying Values of T. ($\lambda_i = 15$, $\alpha_{ij} = 0.5$, $SOP = 25\%$)

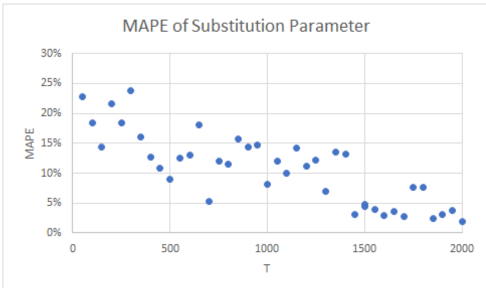

Figure 19: MAPE for Varying Values of T. ($\lambda_i = 15$, $\alpha_{ij} = 0.5$, $SOP = 25\%$)

performance definitely improves for increasing values of T. They even show that the performance converges from a value for T of around 1500 on for α_{ij} and even a value of T of 500 on for λ_i to a good performance. However, this is performed for an SOP of 25%, which is extremely high in practice. Therefore, to improve the applicability of this requirement to real cases, the experiment is also performed using an SOP of 5%. Due to the logic that reducing the SOP will increase the number of periods needed, as the substitution parameter is estimated based on periods in which products are out of stock, this experiment is performed starting from 1500 periods. The results can be seen in Figures 24, 25 and 26. It is interesting to see that for an SOP of

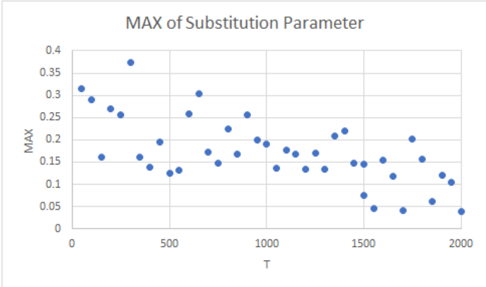

Figure 20: MAX for Varying Values of T. ($\lambda_i = 15$, $\alpha_{ij} = 0.5$, $SOP = 25\%$)

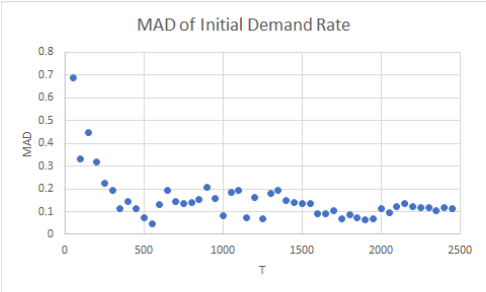

Figure 21: MAD for Varying Values of T. ($\lambda_i = 15$, $\alpha_{ij} = 0.5$, $SOP = 25\%$)

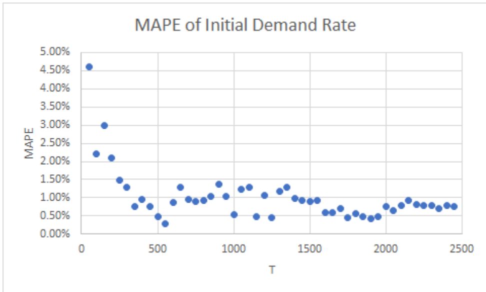

Figure 22: MAPE for Varying Values of T. ($\lambda_i = 15$, $\alpha_{ij} = 0.5$, $SOP = 25\%$)

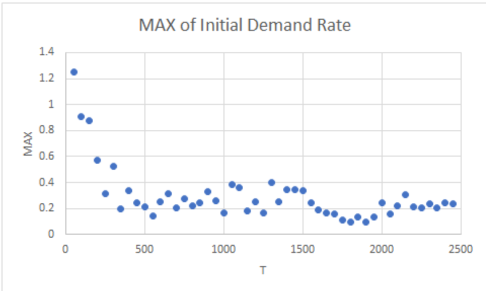

Figure 23: MAX for Varying Values of T. ($\lambda_i = 15$, $\alpha_{ij} = 0.5$, $SOP = 25\%$)

5%, increasing T above 1500 will not improve the performance drastically. It seems there is a negative trend in all three figures, but it is not as extreme as it was for the case in which there was a SOP of 25%. In the beginning, however, it can be seen in these figures that there is a more considerable decrease in the performance measures, indicating that an increase in T does increase the performance. Therefore, it would be safe to say that for an SOP of 5%, at least 2800 periods should be present as that is the point at which the decrease flattens with relatively fair performance.

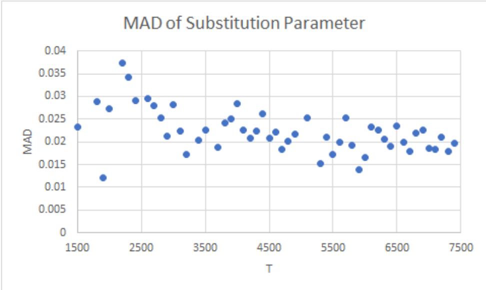

Figure 24: MAD for Varying Values of T. ($\lambda_i = 15$, $\alpha_{ij} = 0.5$, $SOP = 5\%$)

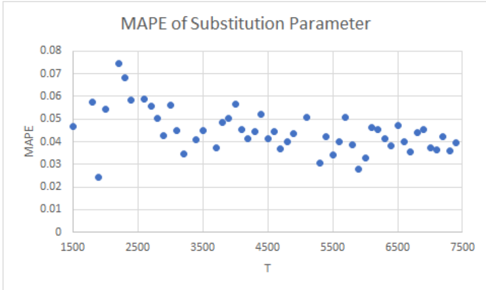

Figure 25: MAPE for Varying Values of T. ($\lambda_i = 15$, $\alpha_{ij} = 0.5$, $SOP = 5\%$)

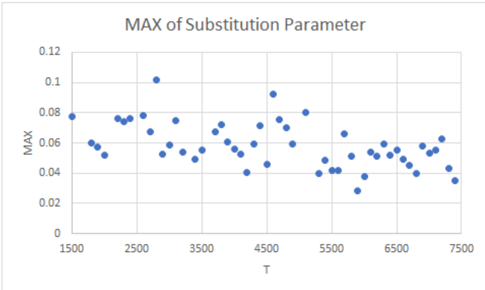

Figure 26: MAX for Varying Values of T. ($\lambda_i = 15$, $\alpha_{ij} = 0.5$, $SOP = 5\%$)

Furthermore, Figures 27, 28 and 29 show the same performance measures but for the initial demand rate. All three performance measures for the initial demand rate

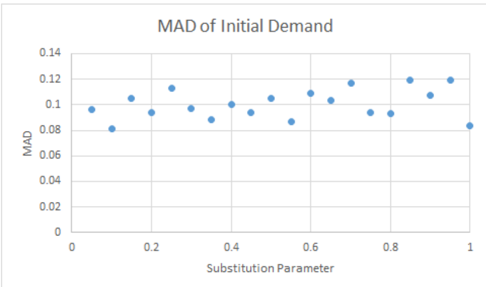

Figure 27: MAD for Initial Demand Rate with varying Substitution Parameter. ($\lambda_i = 15$, $T = 1500$, $SOP = 25\%$)

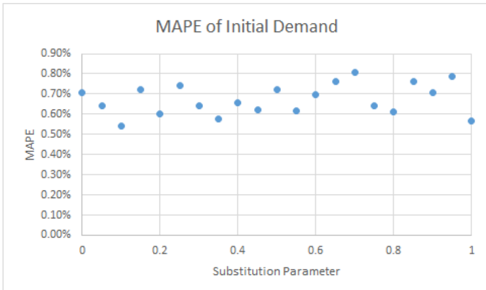

Figure 28: MAPE for Initial Demand Rate with varying Substitution Parameter. ($\lambda_i = 15$, $T = 1500$, $SOP = 25\%$)

are similar. Whereas the MAPE of α_{ij} averaged at 3.0% with T surpassing 1500, λ_i does so at 0.70%. Also, the other measures provide a similar view, enforcing the idea that the model is able to estimate λ_i more accurately. Moreover, there is no clear fluctuating pattern to be seen as there was regarding α_{ij}. Here, it seems more

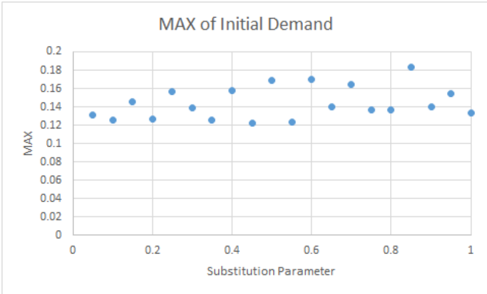

Figure 29: MAX for Initial Demand Rate with varying Substitution Parameter. ($\lambda_i = 15$, $T = 1500$, $SOP = 25\%$)

level, indicating that α_{ij} does not influence the ability to estimate λ_i, which can be estimated with relative accuracy.

Next, the bias of the estimation of both parameters is determined. In Figures 30 and 31, this is denoted such that a positive bias means that the estimation is generally higher than the real value, and a negative bias means that the estimation is generally lower than the real value. It can be seen from the figures that the substitution

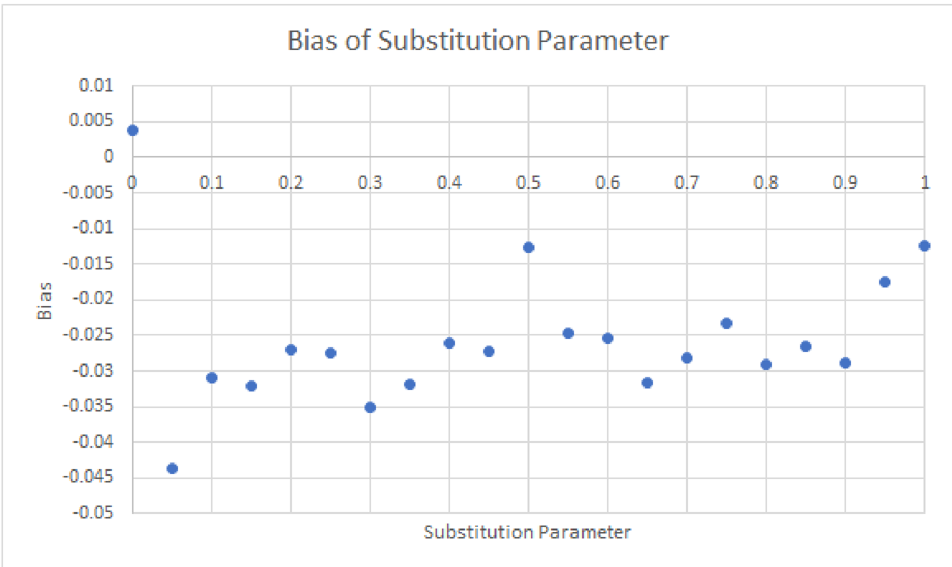

Figure 30: Bias of the Substitution Parameter. ($\lambda_i = 15$, $T = 1500$, $SOP = 25\%$)

parameter has a negative bias, meaning that it is generally underestimated. On the other hand, the initial demand rate revolves more around zero, with more positive values in the lower region of the substitution parameter and more negative values in the upper region. This means that the initial demand rate is generally overestimated when the substitution parameter is lower and underestimated when it

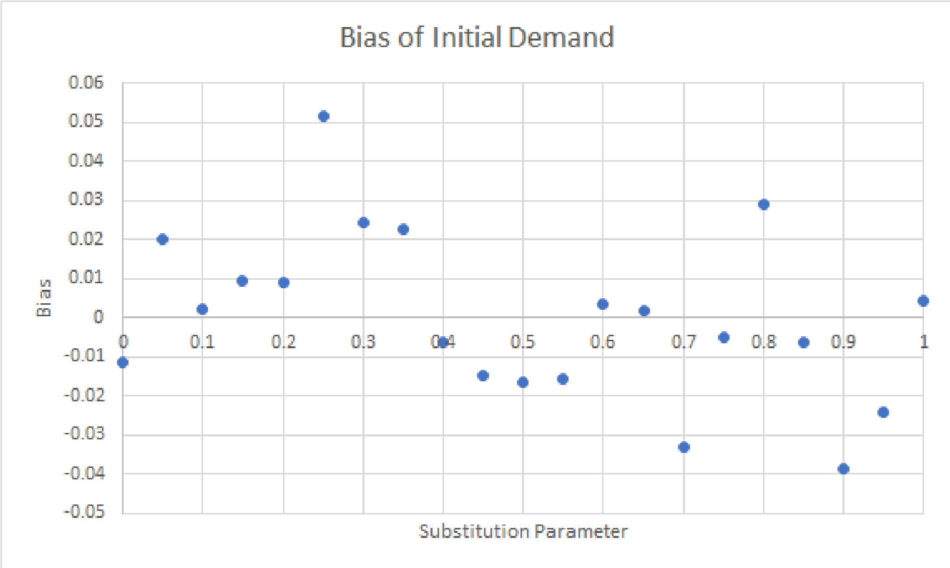

Figure 31: Bias of the Initial Demand Rate. ($\lambda_i = 15$, $T = 1500$, $SOP = 25\%$)

is higher. Overall, this means that the substitution effect is assessed as a more minor phenomenon than it actually is, causing the effect on the order-up-to levels to be smaller as well. This results in a decreased benefit in inventory levels when using the model at hand, whereas it could be more.

Furthermore, as Karabati et al. (2009) state that not one of the two parameter estimation methods performs best in all situations, the subsequent analysis is done. This determines the inventory states and estimates α_{ij} and λ_i. These estimates from the two methods are evaluated, and the best one is chosen. For the analysis, nothing is done with the estimates, but which is selected is just noted. This is done for several values of α_{ij} used as an input parameter for the generation of the data set, and the results are shown in Figure 32.

From this figure, it is clear that the method proposed by Anupindi et al., which used MLE to estimate the parameters, is chosen significantly more often than the one by Karabati et al., which is based on the minimisation of error terms. This contradicts the finding of Karabati et al. mentioned earlier. However, with a mean of 7.2% for the method by Karabati et al. and 92.7% for the one by Anupindi et al., the 95% confidence intervals are [4.9%, 9.5%] and [90.4%, 95.1%] respectively. These intervals mean that with 95% certainty, the method by Anupindi et al. performs better than that of Karabati et al.

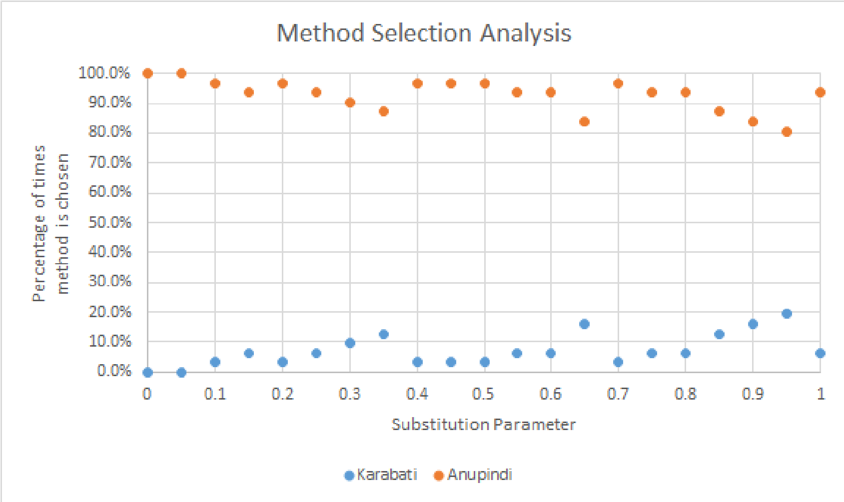

Figure 32: Method Selection Analysis. ($\lambda_i = 15$, $T = 1500$, $SOP = 25\%$)

5.2.1.3 | Accounting for Substitution in the Order-up-to Levels

Finally, the comparison is made between accounting for substitution and determining the order-up-to levels assuming no substitution exists. This is again done for data sets created using a varying α_{ij} and keeping λ_i, T and the SOP constant, assuming that the inventory states are known. This analysis investigates the added benefit of accounting for substitution in multiple severities of it. In both situations, a target fill rate of 95% is set, such that both approaches can be compared evenly. The two approaches taken are as follows. For the first method, the substitution is taken into account. This is done by using the input parameters for the theoretical data sets as α_{ij} and λ_i to see the effect of the proposed greedy heuristic. The second approach negates the effect of substitution in cases where there is actually substitution present in the theoretical data set. For this, α_{ij} is set to 0 for both combinations and an average is taken for λ_i. Taking a simple average suffices because no trend, seasonality, or any other significant disturbance is present, as it is theoretical. Furthermore, the prices of the products are noteworthy as these are a substantial part of Algorithm 1. For these theoretical data sets, a price of €30.− is used for product 1 and €25.− for product 2. Finally, a review period of one day and a lead time of seven days are assumed based on the data found in the data set provided by Slimstock, meaning these assumptions are logical. In this experiment, the substitution is only one way, meaning that customers may opt for product 2 upon a stock out of product 1, but customers will never opt for product 1 upon a stock out of product 2. In the latter case, all demand for that period is considered lost sales. The results of the order-up-to levels are shown in Figures 33 and 34. When only an orange dot can be seen, it should be noted that there is also a blue one, but the two have precisely the same value, placing the orange dot on top of the blue one.

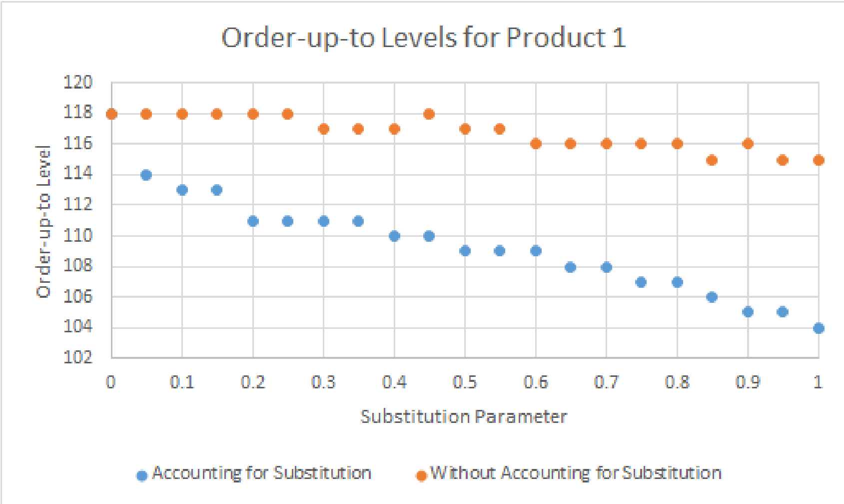

Figure 33: Order-up-to Levels of Product 1. ($\lambda_i = 15$, $T = 1500$, $SOP = 25\%$)

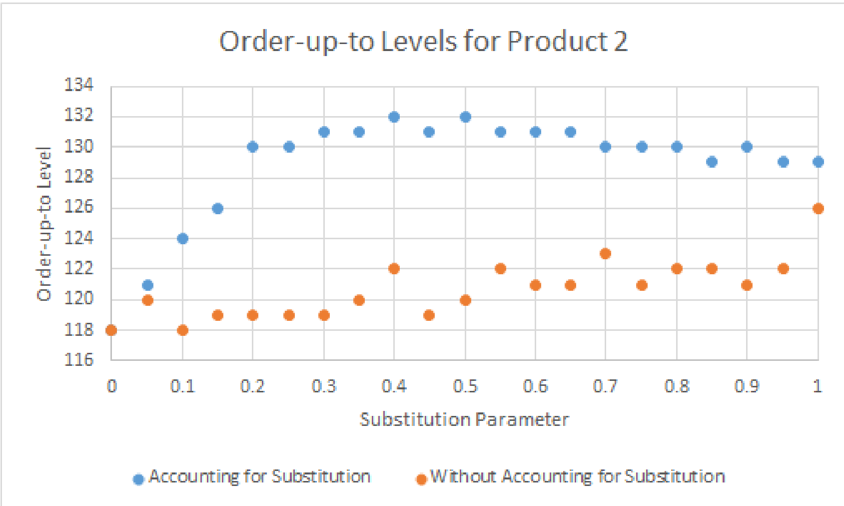

Figure 34: Order-up-to Levels of Product 2. ($\lambda_i = 15$, $T = 1500$, $SOP = 25\%$)

Looking at the figures, it should be noted that accounting for substitution does not always influence the order-up-to levels equally strongly. The order-up-to levels start the same in the cases with and without accounting for substitution. This makes sense because, in those cases, the situations are identical. However, as the substitution parameter increases, the order-up-to levels of product 1 are consistently lower than when substitution is not accounted for. This difference even grows as the substitution parameter increases. This notion is intuitively sound as less and less demand is considered lost sales upon a stock out, it becomes beneficial to lower the order-up-to level as it will be covered by the other product.

Product 2 shows another picture. Here, the order-up-to levels first increase rapidly, after which it decreases slowly. For lower substitution parameter values, the extra uncertainty due to substitution causes the order-up-to levels to increase significantly. As the substitution parameter grows large enough, the uncertainty decreases and it can be mitigated, resulting in decreasing order-up-to levels. Figure 34 shows this happens from a substitution parameter of 0.4. The expected inventory values are determined for both cases to get the full extent of the situations. These can be seen in Figure 35.

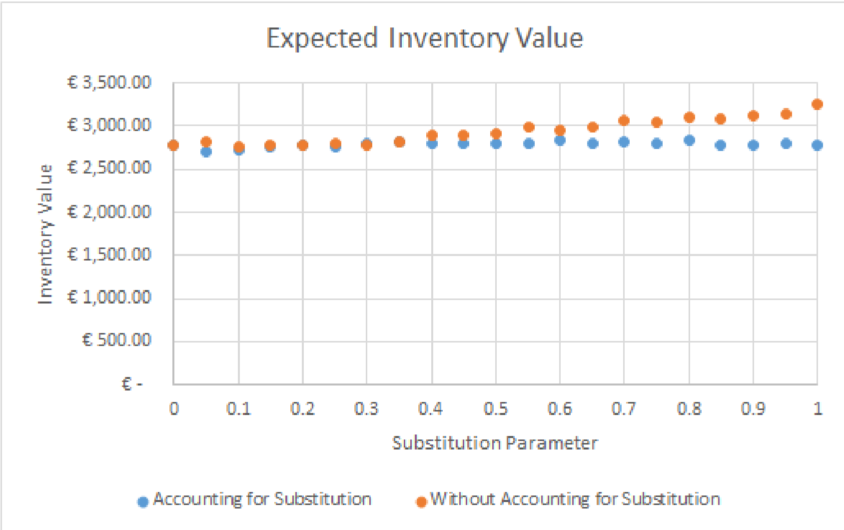

Figure 35: Order-up-to Level Value for Different Substitution Parameters. ($\lambda_i = 15$, $T = 1500$, $SOP = 25\%$)

This figure shows a different pattern than the earlier figures but for the same reason. Due to the fact that the order-up-to levels of product 1 decrease steadily but those of product 2 increase rapidly for low substitution parameters, this roughly cancels each other out. It can be seen in figure 35 that it is somewhat identical to the expected inventory values of the case when substitution is not accounted for. However, from a substitution parameter of 0.4, the order-up-to levels of product 2 decrease due to less uncertainty, which means that the expected inventory value turns out lower than the case in which substitution is negated.

Lastly, Figures 33, 34 and 35 have one oddity. It would make sense that, as the substitution parameter approaches 1, the order-up-to levels of product 1 approach 0. This is because product 1 is more expensive to hold, and it is known that customers opt for product 2 every time, meaning there are no lost sales. However, this is not the case, as seen in the figures, because of how the fill rate is computed. Recall from Section 4.4.2.1 that the fill rate is first computed per product, after which a weighted average is taken to achieve an overall fill rate. Using this method will cause the overall fill rate to never surpass 50% in the case of an order-up-to level of 0 for product 1 if the weights were equal. Even if that is the more desirable choice, this option would never be considered as the end solution since the target fill rate is not reached. Should those options need to be considered, another service measurement should be chosen to use in the greedy heuristic. Currently, this falls outside the scope of this thesis but could be beneficial for future research.

5.2.2 | Real Data

As earlier mentioned, a product group is selected for a case study of the created model. This group is one according to the set minimum and maximum requirements following from Section 5.2.1. These are that a minimum initial demand rate of 12 is needed along with a maximum SOP of 55%, and the data set should contain a minimum of 1500 periods. Although multiple product groups meet the second requirement, the first and last one make selecting one difficult. Since none of the product groups has more than 1500 periods, as that is a little over 4 years and SLIM4 only holds the data in this data set for 2 years, this requirement cannot be met. Therefore, this has to be foregone, which will decrease the results' accuracy. Furthermore, the product groups with more periods available, which is helpful for the accuracy, generally have lower sales per period in the provided data set. In an attempt to mitigate these concerns, product group NL06-00 is chosen with a SOP of 6.30%.

To fully use the model, two things should be created. First, a list of all the transactions over the past two years should be made. Retrieving the input parameters for product group NL06-00 constitutes Table 7.

Product	**Target fill rate**	**Sales price**	**Lead time**	**Review time**
Product 1	90.00%	€53.70	1	7
Product 2	90.00%	€29.18	1	7

Table 7: Input Parameters for the Selected Product Group.

As can be seen in this table, the lead time and the review time are all the same for these products. They differ mainly in the sales price, which will be essential in determining the order-up-to level using Algorithm 1. After running the created model, it resulted in the λ_i and α_{ij} as shown in Tables 8 and 9. Note that almost the same methods are used as in Section 5.2.1.3, meaning that the inventory states and parameters α_{ij} and λ_i are estimated using the methods from Section 4. However, this is different from the case in which substitution is negated. In that case, a simple average is taken as an estimation for λ_i and α_{ij} is set to zero.

Product	$\boldsymbol{\lambda_i}$ **using the Model**	$\boldsymbol{\lambda_i}$ **using a Simple Average**
Product 1	0.524	0.691
Product 2	0.240	0.465

Table 8: Determined λ_i for the Selected Product Group.

	Product 1	**Product 2**
Product 1	0.00	0.029
Product 2	0.008	0.00

Table 9: Determined α_{ij} for the Selected Product Group.

Now that the parameters are estimated, the order-up-to levels can be determined. For this, the algorithm created in Section 4 is used for both cases. This way, it is clear what the effect of using the model is in this real case. The result can be found in Table 10.

Product	**Order-up-to Level with Substitution**	**Order-up-to Level without Substitution**
Product 1	5	6
Product 2	2	5
Fill rate (based on Equation 24)	**91.04%**	**93.41%**

Table 10: Order-up-to Levels with and without Substitution for the Selected Product Group.

Note that the effect of using the model created is much more significant for such low values of α_{ij} than resulting from Section 5.2.1.3. This is mainly because the model's parameter estimation is skewed due to too few periods in the data set. This causes lower parameters as Section 5.2.1.2 denotes a negative bias, which results in lower order-up-to levels. To mitigate the effect of the bias, the λ_i, when the substitution is not accounted for, will also be estimated using the model created, but with α_{ij} forced to 0. With newly found $\lambda_1 = 0.564$ and $\lambda_2 = 0.281$, the order-up-to levels are determined and shown in Table 11. The effect shown is more in line with the findings earlier in this section. Since the substitution parameters are small, the order-up-to level of product 2 can only be decreased by 1. Note that this differs

Product	Order-up-to Level with Substitution	Order-up-to Level without Substitution
Product 1	5	5
Product 2	2	3
Fill rate (based on Equation 24)	**91.04%**	**92.43%**

Table 11: Order-up-to Levels with and without Substitution for the Selected Product Group.

from lower values in Section 5.2.1.3, but that experiment assumed that only one-way substitution is possible. Contrary, this case has a substitution opportunity in both ways. Looking at the expected inventory value for the entire cycle, it has somewhat decreased from €182.85 in the case when substitution is negated to €166.90 in case the substitution is accounted for. Again, this difference of €15.95 is minimal, but that is intuitively sound as the substitution effect is also small.

5.3 | Summary

All in all, this section described several analyses of the created model. It starts with the analysis of theoretical data, such that several scenarios can be created to test the model's performance. The analyses regard the inventory state determination inaccuracy by varying α_{ij}, λ_i and the stock out percentage. This provides an overview under what conditions the model can accurately determine when a product is out of stock. Furthermore, these found conditions are used to create data sets in which only α_{ij} changes. Using those, the ability of the model to accurately estimate α_{ij} and λ_i is measured.

Moreover, the added benefit of using the created model is computed by comparing the determined order-up-to levels to those found using SLIM4's method. Then, the section turns to the analysis of real data as it determines the α_{ij} and λ_i to be used in the heuristic for finding the near-optimal order-up-to levels. This is also compared to the situation in which substitution is not accounted for. The conclusions to be drawn from these results will be presented in Section 6.2.

6 | Conclusion, Discussion and Recommendation

This section will finalise the research done in this thesis. It will start by providing a brief summary of that what was described earlier. After that, the conclusions to be drawn from the presented results will be given, and the research question of this thesis will be answered. Furthermore, possible shortcomings and limitations of this research will be addressed to state how the results can be interpreted. This section will end with recommendations for further research and improvement of what was studied in this thesis.

6.1 | Brief Summary

This thesis introduces a model that can estimate and account for partial inventory-based substitution in the determination of order-up-to levels. This is started by discussing two methods for estimating the substitution parameter and the initial demand rate. The first method is based on Karabati et al. (2009) and tries to minimise weighted squared errors. The second method based on Anupindi et al. (1998) uses MLE. The results of both methods are evaluated based on MAD, MAPE, and MAX criteria, and the method with the best performance is selected for further use in the model. Next, the approach for determining near-optimal order-up-to levels is outlined. This involves using the fill rate and inventory value in a greedy heuristic by iteratively increasing the order-up-to levels until a target fill rate is reached. The fill rate is calculated per product group, accounting for partial substitution effects. The inventory value is computed using average on-hand inventory and purchasing costs.

Finally, the results are presented. There are three primary analyses performed consisting of multiple experiments. First, the ability of the model to correctly approximate the inventory substates is analysed. The effects of varying the substitution parameters, the initial demand rates and the stock out percentage are determined. Second, the model's performance in estimating the substitution parameters and the initial demand rates is measured. This is done based on earlier explained performance measures MAD, MAPE and MAX. Thirdly, the effect of accounting for substitution when determining order-up-to levels given a varying substitution parameter is computed. Unlike earlier analyses, in which the substitution was allowed two ways, it is only permitted one way for the sake of a clear demonstration of the results. Here, not only the order-up-to levels are determined, but also the expected inventory value to get an insight into the monetary side of accounting for substitution in the determination of order-up-to levels.

6.2 | Conclusion

Here, the conclusions will be presented. This is done based on the results presented. Finally, an answer will be provided for the research question, as stated before. Firstly, the model is rather accurately able to estimate the inventory substates based on different situations. Although the substitution parameters do not significantly influence the wrongly classified inventory substates as the results seem patternless but consistently low, varying the initial demand rates does share valuable insights. The number of wrongly classified periods is relatively high for low values of the initial demand rate. Upon closer inspection, it can be seen that this is the case for initial demand rates smaller than 12. It is shown that for values more considerable than that, the number of wrongly classified periods seems patternless but constantly low, as they did for the experiment concerning varying values of the substitution parameters. On top of that, the stock out percentage of the product group significantly affects the model's ability to approximate the inventory substates correctly. The number of wrongly classified periods is relatively low for low stock out percentages, but this increases aggressively for values larger than 55%.

Secondly, the created model is adept at estimating the initial demand rates for all substitution parameters. However, this is not necessarily true for the substitution parameters. The initial experiment was performed using 200 periods, but the performance was relatively poor. Therefore, this total number of periods T was iteratively increased, resulting in better findings for the model's performance in estimating the substitution parameters according to the MAD, MAPE and MAX. It results that at least 1500 periods are needed to provide an accurate estimation with a stock out percentage of 25%, which is a large stock out percentage. Therefore, the same experiment is performed with a stock out percentage of 5%, from which it can be concluded that a minimum of 2800 periods are needed for accurate estimations. Note that this minimum number of periods increases as the stock out percentage decreases as there will be fewer data points to provide accurate estimations. These estimations are also subject to bias. From the analyses performed on that, it can be noted that the substitution parameter is generally underestimated. The initial demand rate, however, is somewhat overestimated for lower values of the substitution parameter and underestimated for higher values of the substitution parameter. Especially the bias for the substitution parameter is proportionally large, resulting in a notable effect on the results. Next, it can be concluded from the experiment regarding estimation method selection that the model chooses the estimation method based on MLE most for all values of the substitution parameter, meaning that the method that tries to minimise weighted squared errors can be negated in future research.

Thirdly, the experiment regarding the heuristical determination of the order-up-to levels shows that these levels decrease for the product that experiences substitution to another product as the substitution parameter increases. Next to that, the product that receives substituted demand will have higher order-up-to levels as it should most likely deal with more demand. This increase in order-up-to level is more rapid for lower substitution parameter values and decreases after about 0.4. After this point, the uncertainty becomes smaller and is easier to mitigate with lower order-up-to levels. However, it stays larger than if the substitution is negated.

The total expected inventory value begins similarly to the case when substitution is negated. Even with the rapid increase in the order-up-to level of the second product, the decrease in the order-up-to level of the first product cancels that out. After the same threshold of a substitution parameter of 0.4, the total expected inventory value starts to be lower than the case for which substitution is negated. This results in the conclusion that although an effect can be seen directly when substitution is accounted for, it becomes financially beneficial when the substitution parameter increases past 0.4. Smaller values of the substitution parameter come with relatively large uncertainty, causing the order-up-to levels to compensate for that. If the substitution parameter is large enough, this uncertainty can be mitigated and the order-up-to levels can decrease steadily. The case study performed shows a similar view. Here, the substitution parameters were small, and the effect was small as well. Although it did not increase any order-up-to levels rapidly, this was because substitution was possible in two directions. Therefore, the effect was cancelled. It could decrease the total expected inventory value a little, making it even more beneficial for lower substitution parameter values in this case.

Recall that the research question stated at the beginning of this thesis is as follows.

"What is the effect of accounting for partial substitution on the order-up-to level advice for inventory models?"

As elaborated earlier in this section, the effect of accounting for substitution with a substitution parameter is that the order-up-to levels of the product experiencing substitution to another product can be decreased to attain a similar fill rate. Next, comparing the proposed model to a case in which substitution is negated results in a lower total expected inventory value when the model is used, and the substitution parameter is higher than 0.4. Below that level, the total expected inventory values are similar. This constitutes the notion that accounting for substitution has a positive effect, provided that the requirements of the parameters explained at the beginning of this section are met as these influence the accuracy of the results, and the substitution level is higher than 0.4.

6.3 | Discussion and Limitations

Below, an overview of the possible shortcomings and limitations will be displayed. Most of the points to be discussed are regarding the parameter estimation, so those will be the start. First, the current parameter estimation approach regards every transaction as equally important for the parameters. However, this is not necessarily the case. The product group of the real data set most likely includes several processes that introduce more variability into the data. Think about a trend over the course of two years, possible seasonality of products and promotions introduced that boost sales. SLIM4 accounts for situations like these by computing these effects and normalising the data. This is not done for the current model, introducing more inaccuracy into the results.

Second, as the number of products increases, the number of inventory states that need to be present in the data set for an accurate estimation increases exponentially. As stated earlier in this thesis, the number of possible states equals 2^N. However, the state where all products are out of stock does not need to be analysed, resulting in $2^N - 1$ states needing to be present in the data set. While this may be achievable for a small number of products, it will become increasingly more difficult as this number grows, requiring a significant amount of additional data once that happens.

Third, there are some shortcomings in the inventory state approximation. One of them regards the method of estimating the demand for the Poisson distribution. This is currently based upon all nonzero POS intervals. However, that assumes that all other intervals are out of stock, while this is not necessarily true. There may very well be POS intervals in which the product was in stock, while there were no sales, especially for slow movers. This is taken into account later in the method, but for the demand estimation, this is negated. This causes the demand to be overestimated and as a result, too many intervals to be classified as out of stock.

This comes to the fourth point, being a limitation on the classification of out of stock intervals. Using the current method, once the probability of a series of l consecutive intervals is smaller than the threshold of 10^{-4}, the whole series is considered as out of stock. Due to the fact that, especially for products with a low demand rate, intervals that are in stock can have no sales, this may be incorrect. It could be the case that the product went out of stock somewhere within this series, causing this classification to be false.

Fifth, the manner of generating the data sets could prove difficult to analyse by the inventory state approximation. Due to the randomness added in creating the sets, out of stock periods of one or two POS intervals are more likely to occur. When the demand is low, these intervals will then not be considered out of stock due to the Poisson distribution, while in truth they are. In practice, these cases might not occur that much, but in the theoretical experiments, it does happen. That could be the reason why the out of stock inventory states are almost all wrongly classified for low values of the initial demand rate.

Sixth, the current determination of the order-up-to level includes the creation of a new demand distribution based on the Poisson distribution. This includes the parameter $A_i(q_i)$, which is the probability of product i being in stock given that the inventory level at the beginning of the cycle is equal to q_i. This parameter can be approximated by subtracting the demand during the $R+L$ period from the order-up-to level. However, this is a pessimistic approximation as the time between reaching the order-up-to level and the arrival of the order is most likely shorter than $R+L$. This is because once the inventory position is reviewed and it is less than the order-up-to level, this happened somewhere between the previous review moment and the current one. To improve this accurately, the expected undershoot, which is the expected number of products sold below the order-up-to level at the beginning of the cycle, should be used instead of the demand during the $R+L$ period.

Seventh and finally, the results do not incorporate the possibility of more than two products substituting each other. The model currently created can estimate and account for situations like this, but the validity and accuracy of that need to be researched. The addition of products in the product group could influence the accuracy of the estimations as well as the added benefit of the order-up-to level determinations. This was outside this thesis's scope and should, therefore, be researched in further research. However, this notion does mean that the conclusion cannot, with definite certainty, be generalised for more than two products.

6.4 | Recommendations

Following the earlier presented conclusions and discussion, several key recommendations emerge, providing ideas on further improvement of the model and possible implementation. Currently, this model looks only at one product group. However, looking at the assortment as a whole could be beneficial. This might be less relevant for a paint store, but if, for example, the little blocks of cheese are out of stock in stores like supermarkets, customers might buy more chips as a snack. These two products are in different product groups but may have some substitution effects that can be considered. This may increase the added benefit of accounting for substitution as not only the intra-product group but also the inter-product group effects are taken into account.

Next, as explained before, the model works for two or more products. However, in the results presented before, the decision was made, for the sake of clear demonstration, that only product groups of two products would be tested. This is not realistic in practice, so more research should be done into the performance of this model on more than two products in a product group.

Furthermore, there is a general recommendation regarding the estimation of the parameters. As the model currently does not reckon that more recent periods have more value for future demand, it should be looked into that these periods are valued as more important. Next to that, the model should normalise for seasonality, trend and promotions to work with cleaner data. This will improve the quality of the data and, with that, the output quality.

Moreover, further research is needed to accurately determine the order-up-to level regarding the incorporation of the undershoot. Currently, an approximation of the inventory level at the beginning of the cycle is made by using the expected demand during the $R+L$ period. However, this is a pessimistic view of the situation as the time period between reaching the order-up-to level and the beginning of the cycle is most likely shorter than $R+L$. Incorporating the expected undershoot considers this and will improve the quality of the resulting order-up-to levels.

Additionally, the results presented earlier show that some inaccuracy is introduced at every step. Altogether, this may have a significant effect on the outcomes, which is currently not taken into account when doing the experiments. Therefore, it is good to critically look at the experiments and the methodology used to find out what the effect of the inaccuracy is on the results of the model. This is not performed for the current analyses but should be done in further research.

Next to that, the model is currently focusing on minimising expected inventory value while reaching a target fill rate. However, it is more likely that one of Slimstock's customers wants to focus on expected revenue or, even better, expected profit. Therefore, it could be beneficial to look into adapting the heuristic to optimise these objectives as well. This also goes with the notion made in Section 5 that possibly beneficial combinations are not reached due to the current computation of the fill rate. Therefore, looking critically at the optimisation objective of the greedy heuristic is integral to the model's improvement.

In addition, the discussion describes three shortcomings of the inventory state approximation method. To solve this there are two main possibilities. On the one hand, the method should be looked at critically and the data set generation should be improved. After this is done, the performance should be experimented on again for more accurate results. On the other hand, Slimstock should hold accurate data about inventory levels for a significant period of time. If this is done, the inventory state approximation can be eliminated from the methodology completely, which results in more accurate outcomes. The latter one is simpler as well as more desirable, so that is primarily recommended.

Finally, the model is currently a separate program made using Python. For Slimstock to fully integrate this into SLIM4, it should be adapted to fit SLIM4's infrastructure. However, this can only be done if the model is completely improved using the aforementioned points. Therefore, it is recommended first to improve the model. Afterwards, look into what should be adapted and in what manner, such that this model gets ready to be integrated into Slimstock's software and can be used by their customers to improve their inventory management further.

References

Anupindi, R., Dada, M., & Gupta, S. (1998). Estimation of Consumer Demand with Stock-Out Based Substitution: An Application to Vending Machine Products. *Marketing Science*, *17*(4), 406–423. https://doi.org/10.1287/MKSC.17.4.406

Bassok, Y., Anupindi, R., & Akella, R. (1999). Single-Period Multiproduct Inventory Models with Substitution. *Operations Research*, *47*(4), 632–642. https://doi.org/10.1287/OPRE.47.4.632

Boylan, J. E., Syntetos, A. A., & Karakostas, G. C. (2008). Classification for forecasting and stock control: a case study. *Journal of the Operational Research Society*, *59*(4), 473–481. https://doi.org/10.1057/palgrave.jors.2602312

Conlon, C. T., & Mortimer, J. H. (2013). Demand Estimation Under Incomplete Product Availability. *American Economic Journal: Microeconomics*, *5*(4), 1–30. https://doi.org/10.1257/mic.5.4.1

Deflem, Y., & Van Nieuwenhuyse, I. (2013). Managing Inventories with One-way Substitution: A Newsvendor Analysis. *European Journal of Operational Research*, *228*(3), 484–493. https://doi.org/10.1016/J.EJOR.2013.02.003

Drezner, Z., Gurnani, H., & Pasternack, B. A. (1995). An EOQ Model with Substitutions Between Products. *Journal of the Operational Research Society 1995 46:7*, *46*(7), 887–891. https://doi.org/10.1057/JORS.1995.120

Dutta, P., & Chakraborty, D. (2010). Incorporating One-way Substitution Policy into the Newsboy Problem with Imprecise Customer Demand. *European Journal of Operational Research*, *200*(1), 99–110. https://doi.org/10.1016/J.EJOR.2008.12.011

Honhon, D., Gaur, V., & Seshadri, S. (2010). Assortment Planning and Inventory Decisions Under Stockout-Based Substitution. *Operations Research*, *58*(5), 1364–1379. https://doi.org/10.1287/opre.1090.0805

Honhon, D., & Seshadri, S. (2013). Fixed vs. Random Proportions Demand Models for the Assortment Planning Problem Under Stockout-Based Substitution. *Manufacturing & Service Operations Management*, *15*(3), 378–386. https://doi.org/10.1287/msom.1120.0425

Hsu, V. N., Li, C. L., & Xiao, W. Q. (2005). Dynamic Lot Size Problems With One-way Product Substitution. *IIE Transactions (Institute of Industrial Engineers)*, *37*(3), 201–215. https://doi.org/10.1080/07408170590899607

Huang, D., Zhou, H., & Zhao, Q. H. (2011). A Competitive Multiple-product Newsboy Problem with Partial Product Substitution. *Omega*, *39*(3), 302–312. https://doi.org/10.1016/J.OMEGA.2010.07.008

Karabati, S., Tan, B., & Öztürk, Ö. C. (2009). A Method for Estimating Stock-out-based Substitution Rates by Using Point-of-sale Data. *IIE Transactions*, *41*(5), 408–420. https://doi.org/10.1080/07408170802512578

Khouja, M., Mehrez, A., & Rabinowitz, G. (1996). A Two-item Newsboy Problem with Substitutability. *International Journal of Production Economics*, *44*(3), 267–275. https://doi.org/10.1016/0925-5273(96)80002-V

Kök, A. G., & Fisher, M. L. (2007). Demand Estimation and Assortment Optimization Under Substitution: Methodology and Application. *Operations Research*, *55*(6), 1001–1021. https://doi.org/10.1287/opre.1070.0409

Mahajan, S., & Van Ryzin, G. (2001a). Inventory Competition Under Dynamic Consumer Choice. *Operations Research*, *49*(5), 646–657. https://doi.org/10.1287/OPRE.49.5.646.10603

Mahajan, S., & Van Ryzin, G. (2001b). Stocking Retail Assortments Under Dynamic Consumer Substitution. *Operations Research*, *49*(3), 334–351. https://doi.org/10.1287/OPRE.49.3.334.11210

McGillivray, A. R., & Silver, E. A. (1978). Some Concepts For Inventory Control Under Substitutable Demand*. *INFOR: Information Systems and Operational Research*, *16*(1), 47–63. https://doi.org/10.1080/03155986.1978.11731687

Parlar, M. (1988). Game Theoretic Analysis of the Substitutable Product Inventory Problem with Random Demands. *Naval Research Logistics (NRL)*, *35*(3), 397–409. https://doi.org/https://doi.org/10.1002/1520-6750(198806)35:3<397::AID-NAV3220350308>3.0.CO;2-Z

Pasternack, B. A., & Drezner, Z. (1991). Optimal Inventory Policies for Substitutable Commodities with Stochastic Demand. *Naval Research Logistics (NRL)*, *38*(2), 221–240. https://doi.org/https://doi.org/10.1002/1520-6750(199104)38:2<221::AID-NAV3220380208>3.0.CO;2-7

Rajaram, K., & Tang, C. S. (2001). The Impact of Product Substitution on Retail Merchandising. *European Journal of Operational Research*, *135*(3), 582–601. https://doi.org/10.1016/S0377-2217(01)00021-2

Rao, U. S., Swaminathan, J. M., & Zhang, J. (2004). Multi-Product Inventory Planning with Downward Substitution, Stochastic Demand and Setup Costs. *IIE Transactions (Institute of Industrial Engineers)*, *36*(1), 59–71. https://doi.org/10.1080/07408170490247304

Shah, J., & Avittathur, B. (2007). The Retailer Multi-item Inventory Problem with Demand Cannibalization and Substitution. *International Journal of Production Economics*, *106*(1), 104–114. https://doi.org/10.1016/J.IJPE.2006.04.004

Shin, H., Park, S., Lee, E., & Benton, W. C. (2015). A Classification of the Literature on the Planning of Substitutable Products. *European Journal of Operational Research*, *246*(3), 686–699. https://doi.org/10.1016/J.EJOR.2015.04.013

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). Inventory and Production Management in Supply Chains, Fourth Edition. *Inventory and Production Management in Supply Chains, Fourth Edition*, 1–781. https://doi.org/10.1201/9781315374406/Inventory-Production-Management-Supply-Chains-Edward-Silver-David-Pyke-Douglas-Thomas

Slimstock B.V. (2023). *Essentials of SLIM4* (tech. rep.).

Smith, S. A., & Agrawal, N. (2000). Management of Multi-Item Retail Inventory Systems with Demand Substitution. *Operations Research*, *48*(1), 50–64. https://doi.org/10.1287/OPRE.48.1.50.12443

Statista. 'Estimated value of the in-store and e-commerce retail sales worldwide from 2022 and 2026 (in trillion U.S. dollars)'. https://www.statista.com/statistics/1095969/retail-sales-by-channel-worldwide/. *Statista*.

Talluri, K., & van Ryzin, G. (2004). Revenue Management Under a General Discrete Choice Model of Consumer Behavior. *Management Science*, *50*(1), 15–33. https://doi.org/10.1287/mnsc.1030.0147

Vulcano, G., van Ryzin, G., & Chaar, W. (2009). OM Practice—Choice-Based Revenue Management: An Empirical Study of Estimation and Optimization. *Manufacturing & Service Operations Management*, *12*(3), 371–392. https://doi.org/10.1287/msom.1090.0275

Ye, T. (2014). Inventory Management with Simultaneously Horizontal and Vertical Substitution. *International Journal of Production Economics*, *156*, 316–324. https://doi.org/10.1016/J.IJPE.2014.06.017

Yu, Y., Shou, B., Ni, Y., & Chen, L. (2017). Optimal Production, Pricing, and Substitution Policies in Continuous Review Production-Inventory Systems. *European Journal of Operational Research*, *260*(2), 631–649. https://doi.org/10.1016/J.EJOR.2017.01.012

Appendices

A. Product Group Stock out Percentage Analysis

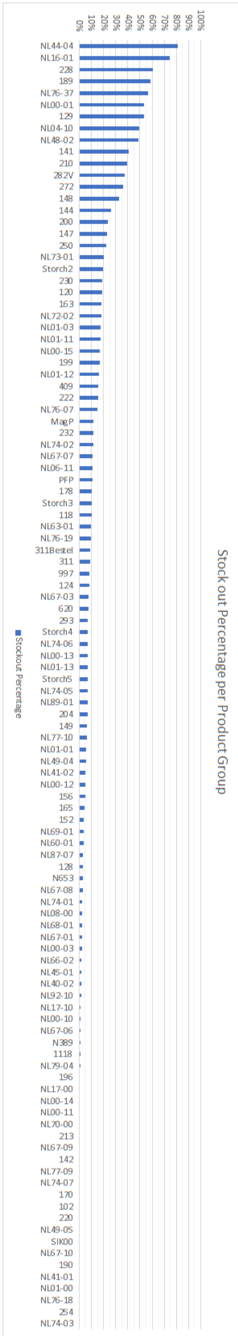

Figure 36: Product Group Stock out Percentage Analysis

B. Product Group Lost Sales Analysis

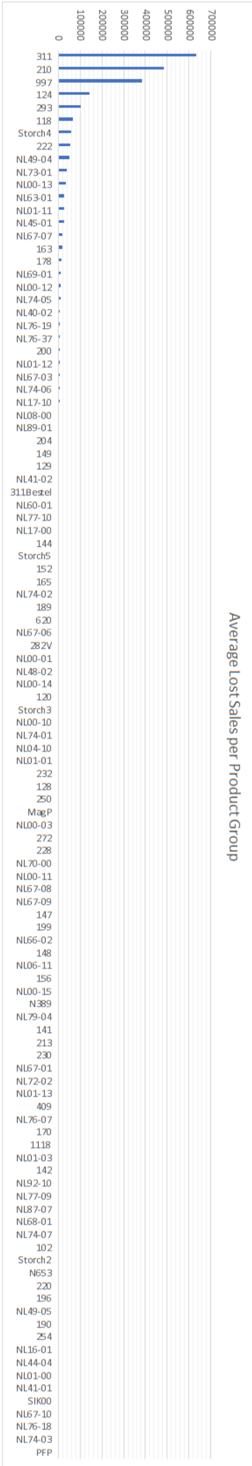

Figure 37: Product Group Lost Sales Analysis

C. Product Group Lost Revenue Analysis

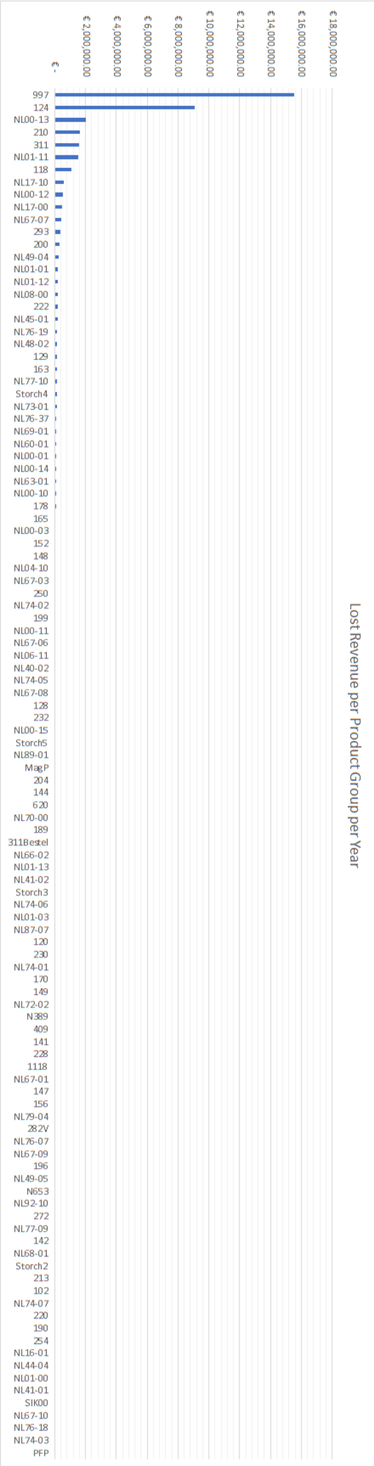

Figure 38: Product Group Lost Revenue Analysis

D. Product Group Substitution Parameter

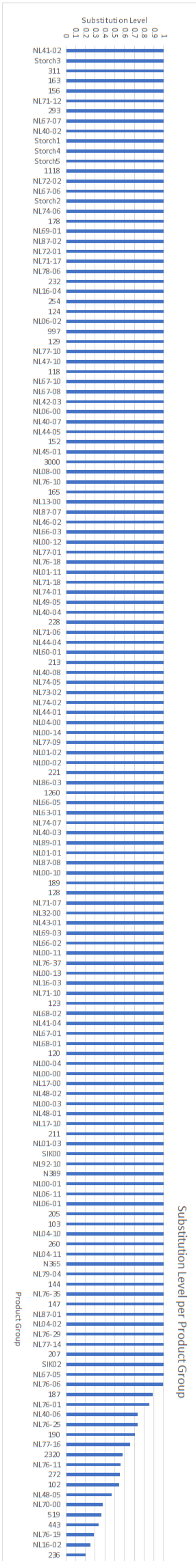

Figure 39: Product Group Substitution Parameter