University of Twente

Graduation assignment
master Industrial Engineering & Management

Production and Logistics Management

Improving retailers' service level: incorporating the order line size distribution and preventing unsellable stock levels

---

Author:
J.M. (Thijs) Veldhuizen

Supervisors:
Dr. M.C. (Matthieu) van der Heijden
Dr. ir. A. (Ahmad) al Hanbali
ir. J. (Jan) van Oostrum

17 May 2017

This page is left blank intentionally.

Management summary

This research project has been carried out at the University of Twente, in cooperation with *Slimstock*, the European market leader in inventory optimization. Currently, there is a considerable potential of keeping *useless inventory* which will hardly deplete, as the probability of a customer wanting this amount - e.g. three mugs - is very low. These different order line sizes are not accounted for in the standard (R,s,nQ) inventory model we research. This leads to the frequent occurrence of "left-over" inventory being above the reorder level, at which a replenishment order is triggered. Data from retail client *Company X* are used to tackle the following problem statement:

"An unknown but significant large set of retailers' SKUs reaches an unsellable low inventory level, leading to under-performances regarding their volume fill rates."

93.59% of the SKUs sold by company X, have a specific order line distribution, for which a minimum number of stocks on shelf - chosen by the planner - is applied, in addition to further automatic calculations. Thus, we formulate the research goal:

"to improve the inventory policies by determining more suitable reorder points, ultimately leading to more SKUs with typical order line sizes meeting their target service level."

Relevant methods from literature for calculating safety stock include among others *target fill rate* (percentage of demand directly met from stock) approaches, that assume demand during a period of time to be normal distributed. Next, the *compound Poisson distribution* is used. Finally, we research methods (we call the literature-based approach the "Add undershoot"-method) including undershoot (difference of the inventory position and the reorder point when ordering).

A theoretical standard model that resembles one of the inventory models of *Slim4* - Slimstock's software package for demand forecasting and inventory optimization - is compared with methods from literature and own-derived heuristics. In the final experiment we test the standard approach, the literature-based methods, and three own-derived safety stock calculations. These last methods directly incorporate the empirical distribution for the order line size. By doing so, we aim to set safety stock in such a way at least one extra customer can be served from stock right before replenishment. Ultimately, this would abolish the aforementioned useless inventory. These heuristics - the "Overwrite"-, "Max"- and "First moment β^*"-method - are based on the expected undershoot (either overwriting the standard safety stock, or taking the maximum of the undershoot and the standard safety stock), and the in-service order line size (an order line size the retailer wants to able to sell). The latter comprises the SKU's first moment - the average - of the order line size multiplied by the SKU's target fill rate and the customer arrival rate, making an amount per order line that the company wants to sell from average stock. We test the methods using a dataset from Company X, containing both *multi-modal (MM) order line size SKUs* and *unimodal (UM) order line size SKUs.*

We developed a classification scheme, based on the five most often occurring order line sizes, through which multi-modal order line size SKUs and unimodal order line size SKUs can be identified. E.g.: the former would have gaps in frequencies for succeeding order line sizes like '3' and '7', and for the latter would hold: the bigger the order line size, the lower its frequency of occurrence. This characteristic is used by calculating the 'distances' between ranks of successive order line sizes. Currently (assuming lost sales) in our own model, 75.68% of the multi-modal SKU do not meet their target fill rate on shop-level, versus 63.58% of the other type of SKUs. Note: in the latter case there are three SKUs showing unrealistically bad performance.

We measured performance through volume fill rate and average stock on hand, and we differentiate between lost sales and full back-ordering configurations. The compound Poisson method performs best for multi-modal order line size SKUs, as with this method the average volume fill rate increases with 8.31%-point, against 47.12% more stock on hand, resulting in 84.79% of the SKUs meeting their target fill rate. The "Add undershoot"-method

performs best for unimodal order line size SKUs, as the average volume fill rate increases with 1.98%-point, against 13.16% more stock on hand, resulting in 71.97% of the SKUs meeting their target fill rate. Other methods are denoted in Table 1, and in Figure 1 we depict the differences in performance in a graphical way. Here we see that, when desiring higher volume fill rates, multi-modal order line size SKUs can best handled with a compound Poisson method, and unimodal order line size SKUs show a slight improvement when adding undershoot.

		Standard	**"Add undershoot"**	**Compound Poisson**	**"Overwrite"**	**"Max"**	**"First moment β^*"**
MM SKUs	*Lost sales*	67.61% (9.15)	78.47% (10.89)	84.79% (12.66)	66.93% (9.17)	71.48% (9.79)	80.12% (12.22)
	Back-ordering	67.61% (9.15)	78.51% (10.89)	84.90% (12.68)	66.81% (9.13)	71.67% (9.78)	80.07% (12.22)
UM SKUs	*Lost sales*	64.15% (22.31)	71.97% (25.07)	64.95% (18.72)	50.46% (16.03)	65.07% (22.55)	55.38% (17.94)
	Back-ordering	64.52% (22.38)	72.24% (25.13)	64.67% (18.72)	50.69% (16.02)	65.07% (22.59)	55.75% (17.90)

Table 1: *Percentages of SKUs meeting their target fill rates, and between brackets the average stock levels.*

In Table 1 we underlined the best performing (mainly in terms of effectiveness) alternative heuristics. We also ran a sensitivity analysis, which led to the conclusion that unimodal order line size SKUs's volume fill rates could also be improved (on average 94.94%) by increasing their target fill rates to 99%. This yields lower stock levels (on average 22.78) than the "Add undershoot" method, and no new heuristic should be implemented. Hence, based on our experiment, we propose the following recommendations for Slimstock:

- Introduce the classification scheme, and analyse first results for other retail clients' data, for verification and fine-tuning of the method;
- Implement the "compound Poisson"-method for multi-modal order line size SKUs in Slim4. This method works better, especially for SKUs with relative high median order line sizes, less than 30 customers per year, and small chances of selling per one or two items. 8.31%-point more demand and 6.19%-point more customers can be served, against 47.12% more stock.
- Do not alter the inventory model for unimodal order line size SKUs yet. Although in "Add undershoot" method the undershoot is correctly considered, leading to less occurrences of 'useless inventory', we can also improve this type of SKU's performance by increasing the standard model's target fill rate. As this is more convenient for implementing, we recommend to further research the "Add undershoot" method. This could lead to 1.98%-point more demand and 1.49%-point more customers to be served, against 13.16% more stock.
- Our recommendations touch upon the core of an inventory model, so we recommend a long-term pilot, to research the impact of aspects such as trends, seasonality, and promotion.

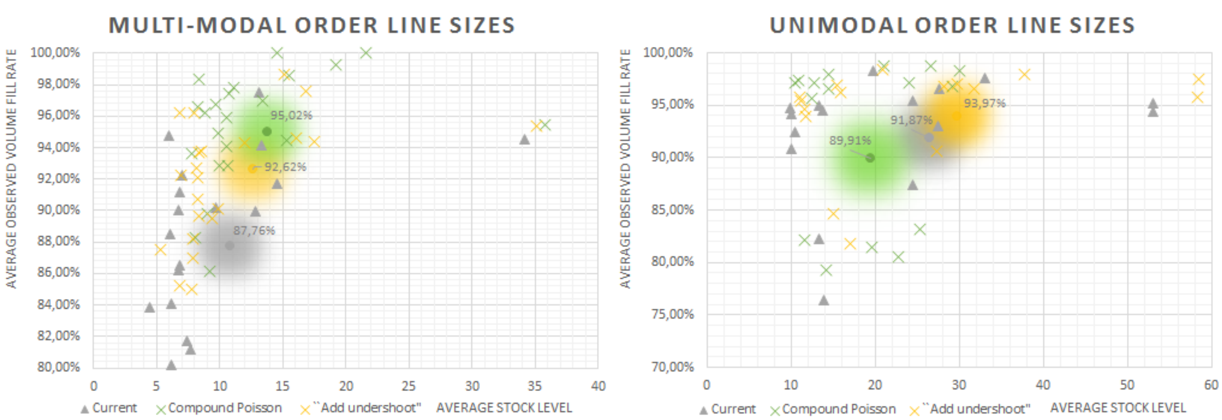

Figure 1: *The 22 multi-modal and 16 unimodal order line size SKUs and their performance for the two best-performing methods and the standard approach. The clouds represent the weighted average performance of each method.*

Acknowledgements

After an exciting, intense and interesting period of eight months, it is time to wrap up my graduation assignment. After being recruited at the *National Inventory Day* (in Dutch: "Nationale Voorraaddag") for the *Young Professional Programme*, I was given the chance to do my graduation assignment at Slimstock as well.

First of all, I like to thank Jan van Oostrum. Jan is developer at Slimstock and he has been my daily supervisor, helping me in the first place with all ins and outs concerning Slim4. But most of all, he and Bart van Gessel, also developer at Slimstock, contributed a lot with their high-end knowledge on inventory management. I would also like to thank all other colleagues at Slimstock who helped me.

Of course this all would not have been possible without the disciplined but fair supervision by my first supervisor Matthieu van der Heijden. Through several intense meetings we really were able to solve issues at hand quick and effective. Furthermore, I'd like to thank my second supervisor Ahmad al Hanbali for his second opinions and further in-depth knowledge.

Last but certainly not least I would like to thank my girlfriend for her support and help, and proofreading; and my friends and family for their support as well. All in all, I have learned a lot during the past eight months, both about inventory management in specific application areas and practical implementation in Excel - aided by SQL - and about my personal development.

Thijs Veldhuizen
Deventer, 17 May 2017

Contents

List of Tables

List of Figures

List of Abbreviations and Symbols

Abbreviation	*Definition*
BMAP	Batch MAP
BO	Back-Ordering
DC	Distribution Centre
EM	Expectation Maximization
EOQ	Economic Order Quantity
IOQ	Incremental Order Quantity
IP	Inventory Position
KPI	Key Performance Indicator
MAP	Markovian Arrival Process
MLE	Maximum Likelihood Estimate
MOQ	Minimum Order Quantity
MSCAD	Mixture SCAD
NL	Normal Loss
SCAD	Smoothly Clipped Absolute Deviation
SKU	Stock Keeping Unit
SOL	Sold Order Line
VBA	Visual Basic for Applications

Symbol	*Definition*
β	Chap. 3: hyper-parameter
β	Observed volume fill rate
β^*	Target volume fill rate
γ	Order line fill rate
$\delta_{x,x+1}$	Absolute differences between the ranked successive order line sizes
$\delta_{\%}$	Percentage of differences $\delta_{x,x+1}$ larger than 1
δ^*	Threshold for $\delta_{\%}$ to exceed for classification as multi-modal
χ^2	Chi-square (for testing)
λ	Arrival rate of customers in a certain period of time
μ	Average demand
ϕ	Standard normal distribution
Φ	Inverse cumulative standard normal distribution
σ	Standard deviation
X	Order line size
D	Demand during period
$J_u(k)$	Special function from the standard normal distribution for Tijms &Groenevelt
k	Safety factor
k	Chap. 3: Number of customer
L	Lead time
n	Batches of order quantity
n	Chap. 3: total number of orders within the lead time
p	Chap. 3: probability of success, i.e. fraction of items that were of a certain SKU
Q	Order quantity
R	Review period
r	Chap. 3: parameter for negative binomial distribution
s	Reorder point
S	Base stock level
SoH	Stock on Hand
ss	Safety stock
w	Chap. 3: weights (for mixture modelling)
z	Chap. 3: allocation

Glossary

- **Client**: Customer of Slimstock that uses Slim4 for its inventory optimization and ordering process.
- **Customer**: End-consumer of Slimstock's clients.
- **In-service order line size**: The SKU's average order line size multiplied by the SKU's target fill rate, making an amount per order that the company wants to sell from stock (meaning 'in-service).
- **Order Line Fill Rate**: We distinguish between *Target Fill Rate* and *Observed Fill Rate*. A target order line fill rate is an input parameter indicating the fraction of order lines that should be immediately fulfilled from stock. This equals the fraction of total orders that can be satisfied from inventory without shortages (Larsen & Thorstenson, 2008). Observed fill rates on the other hand are performance indicators resulting from experiments, and ideally they should be greater than or equal to target fill rates.
- **Volume Fill Rate**: Again, we distinguish between *Target Fill Rate* and *Observed Fill Rate*. A target volume fill rate is an input parameter indicating the fraction of demand that should be immediately fulfilled from stock. This equals a regular volume fill rate, defined as the fraction of total demand that can be satisfied from inventory without shortages (Silver, Pyke, & Thomas, 2017).
- **Multi-modal order line size**: A demand pattern that - on order line level - has more than one local maximum in its order line size distribution plot.
- **Order line**: Part of a customer's order at a client's shop, consisting of one Stock Keeping Unit (SKU), and information about the amount in which it is sold.
- **Replenishment order**: An order that is placed at the client's supplier for the corresponding SKU. When we use 'order' or 'order line' we mean the customer order (line), which represents demand from the customer, and we explicitly state 'replenishment order' when we mean an order to replenish the client's inventory.
- **Shop**: A client's location where both inventory is held, and sales take place.
- **Stock Keeping Unit**: An SKU represents a single, specific product (for example: a mug, with article code 1234) from a client, for which inventory is held.

Introduction

The first step of this thesis is to define the exact problem that has to be solved. For this we need a brief background on the company of research, and an introduction in its challenges. We introduce the need for research, our definition of the problem, and we formulate corresponding research questions, each accompanied by a brief methodology. Next, the scope is defined, where the boundaries of our research are set, and we conclude this chapter with a reading guide, addressing how the report is structured.

1.1 Background Slimstock

Founded in 1993, *Slimstock* has become the market leader in inventory optimization, with more than 650 clients all over the world. Its main software package, named *Slim4*, contains forecasting, demand planning, and inventory management, helping their clients (customers of Slimstock that use Slim4) to get the right inventory to the right place at the right time. Besides software solutions, Slimstock also offers project-based support and professional services, including coaching, analyses, and interim professional support. Slimstock can offer assistance to help reduce inventory while at the same time increasing the service level, thereby increasing efficiency and generating insights for the planners and management. So turnover increases, while costs decrease for Slimstocks clients. The company is organized in several departments, of which the most important for this scholar is *Development*, as we work on improving the software package, which is the responsibility of this department.

1.2 Introduction assignment

This research is conducted as part of the graduation project for the master's programme Industrial Engineering and Management (specialization: Production and Logistics Management) at the University of Twente. The project's duration was eight months, and it took place at Slimstock's head office in Deventer, where we use one of the inventory models from Slimstock's software package Slim4 as a starting point. This package is capable of optimizing both *single stock point inventories* and *multi-echelon inventories*. Its inventory model is based on a universal *(R, s, nQ) policy* from which parameters can be set as such that each of the four configurations from Table 1.1 can be obtained.

	Periodic review	Continuous review
Fixed order quantity	*(R,s,Q)*	*(s,Q)*
Order-up-to-level	*(R,s,S)*	*(s,S)*

Table 1.1: *All possible inventory model configurations within Slim4*

Slim4 aims to fulfil the *Target Volume Fill Rate* (fraction of demand that should be immediately fulfilled from stock), accounting for either *back-orders* or *lost sales* in case of insufficient inventory. These settings are configured during the implementation stage, and can manually be altered by the client's key users (Slim4-developer, 2016a). Some Stock Keeping Units (SKUs; a single, specific product from a client, for which inventory is held) of Slimstocks clients are very often demanded in certain amounts, such as mugs. For example, customers (end-consumers of Slimstock's clients) buy one mug as replacement for a broken one, or because they did not buy enough mugs in the first place, and they buy two, four, or six mugs when they go for a new series of dining equipment like plates, cutlery and mugs. One hardly ever buys three or seven mugs. An example of this pattern is depicted in Figures 1.1a for a low-budget mug and 1.1b for a breakfast plate.

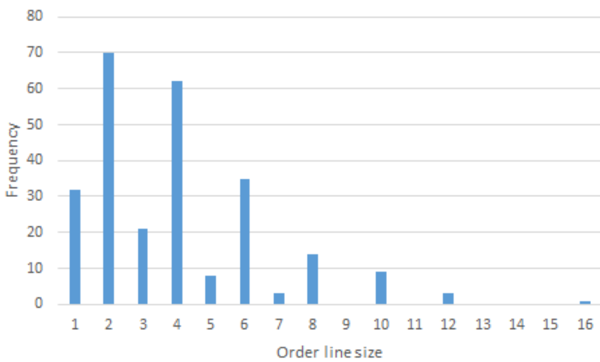

(a) *A low-budget mug clearly show that most orders are of even-numbered sizes*

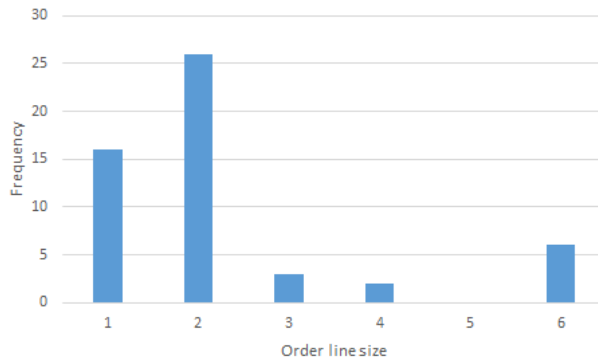

(b) *A breakfast plate clearly shows a gap in order line sizes 3-5, and an inflate in order line size 2*

Figure 1.1: *Exemplary multi-modal order line size SKUs*

These order line (part of a customer's order at a client's shop) sizes are not explicitly accounted for in the demand characteristics or reorder parameters within the inventory model. Replenishment order (an order that is placed at the client's supplier for the corresponding SKU) advices are - among others - determined based on the target volume fill rate (β^*). Stock is not be replenished if only three mugs - which sounds intuitively reasonable - are on stock, in case the reorder point is below three. Yet, when a customer willing to buy four mugs - yielding an observed volume fill rate of 75% - will not accept three mugs. When configured as such, sales are lost. This could be improved by adopting new methods that take into account order line size distributions of SKUs for which these *certain order line sizes* are relevant. Finally, as the true demand is unobservable in most lost sales environments, Slimstock uses a *ready rate*(α) performance measure: *fraction of time during which the stock on hand is above zero*. Although this is an intuitive measure, it can be misleading, due to the fact that the inventory level is likely to be unsellable (in case it is three in this example), despite having a very high ready rate.

1.3 Need for research

The research topics of multi-modal (demand pattern that has more than one local maximum) order line demand probability distributions (in our examples having peaks at even amounts of mugs, or at 2 or 6 breakfast plates), and implementations of order line size distributions in inventory control systems are both found in scientific literature. Several papers combine the research of order (line) size probability distributions and inventory modelling, and there is a great urge in the industries for it (Slim4-developer, 2016a).

To give empirical urgency to our topic of research, we conducted an illustrative brief analysis on the demand patterns of different types of laminate (strictly speaking, we assumed that sales are equal to true demand) for one-and-a-half year. We did this together with an expert panel consisting of the external supervisor and a colleague. First of all, only products having had more than five orders are considered, which was regarded to smallest sample size in order to draw significant conclusions. The laminate SKUs are roughly split into two categories, slow-moving SKUs with a low amount of small sized order lines, and fast-moving SKUs that are sold both often and in larger quantities. We aim to find out how many SKUs show a gap in their order line size figure, as shown in Figure 1.1a or Figure 1.1b. This would suggest a multi-modal demand distribution, and it would function as proof of the existence of this pattern on a large scale, making it interesting for research within this thesis. Ultimately, recognizing this pattern could lead to an improvement of the reorder configuration and the realized service level. This standard classification method has a rather rough approach, only functioning to provide general insights on the size of potential improvement.

Based on this brief, 1-shop laminate analysis (see Appendix A for our approach), we can conclude that 31 out of 69 SKUs that were sold for more than five times showed a potential multi-modal order line size distribution, and thereby are interesting for further research. This equals 44.93% of the SKUs that were reviewed. Although we admit that this is a quick and dirty method, we think this result is expandable to other SKUs within our focus as it classifies both slow-movers and fast-movers from a retailer company. Since the expert panel played a big role in verifying these results, which is not possible in real-life scales, it would be very interesting to further do research on classifying these SKUs through this thesis. Demand forecasts - and thereby replenishment order advices - for multi-modal order line size SKUs ought to improve as a result of incorporating the *order line size* distribution, instead of only looking at the demand during lead time. Figure 1.2 illustrates this issue, with the bold line representing the inventory level, with the dashed black line the inventory position, and the box-plot indicating that there is only variability in determining the total expected demand during the review period. The red box indicates the amount of time in which customers demanding the in-service order line size (the SKU's average order line size multiplied by the SKU's target fill rate, making an amount per order that the company wants to sell from stock) cannot be served, thus the inventory level is 'useless'. In the new situation, demand is forecast on customer arrival rate level and on order line size level, as depicted through the box-plots. Reorder parameters here account for the order line distribution. Note: in the new situation one occasional extra customer demanding the median order line size can be served, depicted by the grey box-plot, although we later expand this assumption to the expected number of customers yet to be served within the planning horizon. We chose the median, as this is a clear measure for often occurring order line sizes, taking in mind that a non-integer average makes no sense in terms of frequently sold order line sizes.

Hence, all order and customer-related information is known, but only expected demand during lead time, and average demand and average number of orders is used when calculating the demand forecasts, all is averaged out. In general terms, we would like to separate these customers arrival and order line size dimensions, by incorporating the transaction data that is already available, see Figure 1.2b. As we discussed earlier, around 44.93% of the SKUs within our preliminary dataset are indicated to have an median order line size (indicated by the vertical box plots in Figure 1.2b). So, this is a highly relevant future scenario.

1.4 Problem definition

After thorough discussions with my supervisor at Slimstock, we conclude on two main streams of challenges (Slim4-developer, 2016a). We have depicted these in Figure 1.3, and further address them below.

Core issue: A number of SKUs sold by retailers are showing under-performance on Key Performance Indicators (KPIs) like observed volume fill rate - an often-used objective for retailers. Hence, when there are only three mugs left, above a reorder point of two, there is an unsellable amount of stock, but Slim4 does

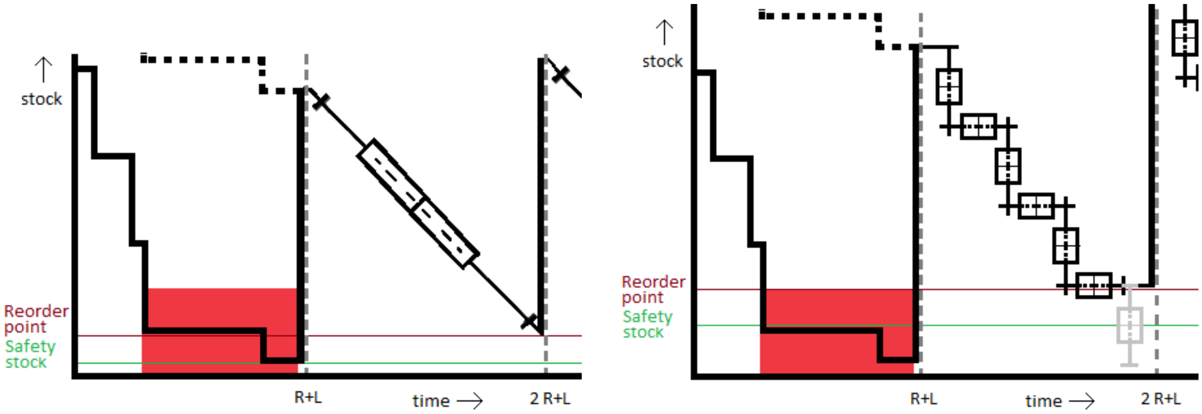

(a) *Old situation only considering total demand during lead time* (b) *Desired situation including order line information*

Figure 1.2: *Comparison of old and new situation of handling inventory*

not yet replenish. We estimated earlier in Section 1.2 that around 45% may show potential for improvement in terms of fill rate, as a result of having typical order line sizes. The direction for improvement lies in the replenishment of these SKUs. The optimal replenishment policy is hard to calculate, since there are also customers who like to buy only one mug (e.g. if one of his/her mugs is damaged), which results in a multi-modal order line size distribution. This makes often used distributions like the *Poisson* or *Normal* distribution not applicable, as they cannot account for multi-modal distributions.

1.a: Slimstocks clients have difficulty identifying the SKUs this 'multi-modal demand' effect applies to, as it is still up to the planner's experience whether or not a SKU has often-occurring order line sizes. They do not have information about which SKUs are usually (i.e.: how often?) sold in certain quantities (i.e.: in what quantity?).

1.b: The knowledge gap in 1a exists due to a lack of validated decision rules, through which an inventory model would be able to automatically classify SKUs.

1.c: After conducting a preliminary literature search for topics like compound Poisson processes with empirical compounding distributions, and service levels, we can conclude that there are only few papers that combine the classification of SKUs based on their particular demand distributions with a service-level approach for improving their performance within inventory modelling software. Nevertheless, there are plenty of papers on the separate topics, so the main challenge lies in combining and applying the right papers.

2.a: Performance is calculated based on sales that in reality take place, instead of the true demand. Next, lost sales are in general not reported to the shop resulting in unobservable lost sales. If three mugs are on stock, replenishment would possibly not take place, and selling is not be possible if a customer demands four mugs. When the reorder point lies below a certain median order line size, stock on hand turns useless as there are hardly customers who are willing to buy the amount that is on stock.

2.b: Data on order line sizes are available yet unused, see Figure 1.2a. We need to consider the order line size distribution, instead of just the demand during the lead time and review period.

Hence, the problem is two-fold, and should be tackled as such. Before adjusting the reordering policy, one needs to be able to automatically identify the relevant SKUs or SKU groups. On the other hand, solely

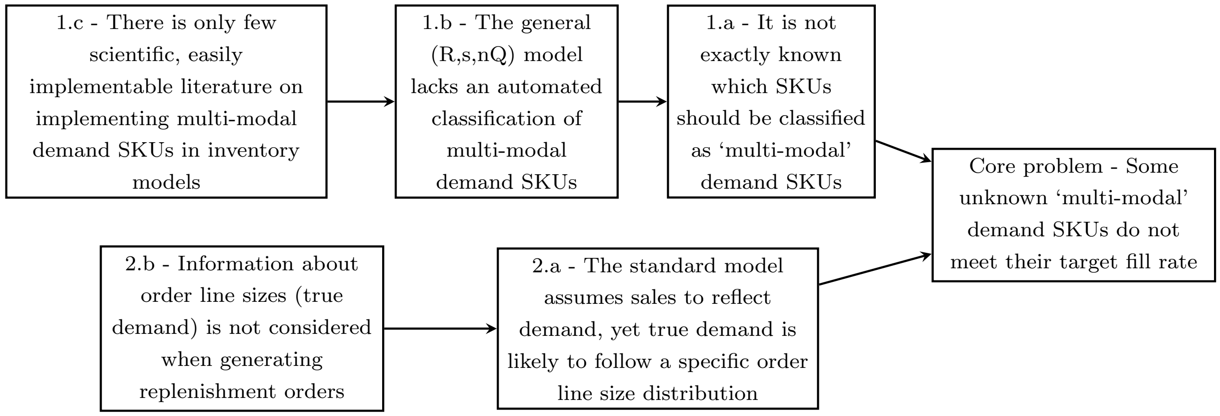

Figure 1.3: *Problem cluster for multi-modal order line size SKUs*

identifying the relevant SKUs would make no sense, both research-wise and for practical use. Therefore, the core problem with its root problems 1 and 2 should be tackled in this sequence. Henceforth, we can conclude with the following problem statement:

"An unknown but significant large set of retailers' SKUs reaches an unsellable low inventory level, leading to under-performances regarding their volume fill rates."

To this extent we formulate the research goal as:

"to improve the inventory policies by determining more suitable reorder points, ultimately leading to more SKUs with typical order line sizes meeting their target service level."

1.5 Research questions

As a result of the earlier introduced need for research and our problem statement, we can develop research questions. These research questions comprise sub questions, and we discuss these one per chapter, as follows:

1. How does the current way of generating order policies using a (R,s,nQ) model at Slimstock look like?
   a. *How do the standard methods as used within Slim4 work?*
   b. *How big are the opportunities for improvement of handling multi-modal order line size SKUs?*
   c. *Under which circumstances are these opportunities highly relevant?*
   d. *What are Slimstock's requirements the solution should adhere to?*

   For the first research question we combine the insights obtained from working with Slim4 (more in-depth knowledge is gained during a training week), with brief, informal meetings with the stakeholders at Slimstock. My daily supervisor and his colleagues are the main source of knowledge in this part. Furthermore, cases from a client are used, based on the scope as described in Section 1.6. Analyses in following chapters are partly based on the insights we gain by zooming into a (R,s,nQ) model.

2. What can we learn from literature about optimization of multi-modal order line size SKUs inventory?

    a. *How can we identify multi-modal order line size SKUs the best way?*

    b. *How can we incorporate the 'multi-modal order line size effect' into inventory management?*

    For our literature research we start by a brief review on the specific type of inventory model we investigate. Next, we need to transform historical sales data into true demand, taking into account unobservable lost sales. The demand distribution should be modelled in an effective way, for which we search interesting literature. Finally, we address research on different methods for replenishment ordering policies within inventory models, taking into account the order line size distribution.

3. How can the inventory management at best incorporate multi-modal order line demand?

    a. *What are the different possibilities, and which alternative performs best?*

    b. *How can the inventory model be validated and verified?*

    c. *Which SKUs are classified and how is this done?*

    At research question 3, we construct a compact inventory model, which can be configured in all scenarios as found in the standard methods, literature, and possibly own derivations. This model is validated and verified throughout the whole process, and here we describe which measures were taken to ensure a valid model.

4. Which heuristic performs best on both the products classified and not classified as 'multi-modal order line size SKUs'?

    a. *Which data are needed and how do we gather them?*

    b. *Which of the methods and heuristics for generating ordering policies performs best?*

    The fourth research question contains an experiment, using the developed inventory model. We link the need for data to the data available, and we select the classified SKUs for our experiment. Finally, the experiment is designed and conducted, and results on performance are obtained. We conclude by analysing the results.

    Hence, the deliverable of this master's assignment consists of several parts: knowledge on classifying SKUs on their order line size distribution and how to implement this at best; an Excel model that supports this piece of knowledge; and finally a brief elaboration on recommendations about the implementation of the results in an inventory management and forecasting system like Slim4.

1.6 Research scope

For our research we limit ourselves to the following characteristics for various reasons. The assumptions are denoted in arbitrary order. Note: assumptions regarding the analysis are specified in Section 4.1, and model assumptions for running the experiment are provided in Section 5.2.

- **Business to Consumer retailers**: the problem of a multi-modal order line size distribution mainly arises at retailers in all kinds of markets such as non-food retail. Moreover, we work with data from a retailer from this market. These chains also tend to have an online web shop. Yet, as these web shops' inventories operate in a very different way (supply and demand are often aggregated for a complete country), we leave online web shops out of this research.

- **First Come First Serve**: in a physical shop (client's location where both inventory is held, and sales take place) customers arrive with varying demand during the day, and they are typically served First Come First Serve. There is (hardly) no prioritizing possible between customers.
- **Purchase to Stock**: Inventory management systems are able to handle both stocked (Purchase to Stock) and non-stocked (Purchase to Order) items. As we investigate an inventory issue, it speaks for itself that we focus on stocked items.
- **Positive demand**: on transaction level negative demand occurs in case of SKUs being returned to the shop. For now we assume the customer demand to be independent from the number of returns, so we abolish all the negative order line sizes. Finally, there are no order lines of zero items, so we restrict ourselves to purely positive demand.
- **Single stock point**: in this thesis we focus on performance on the 'shop-side' of a supply chain, abolishing for example Distribution Centres (DC) and lateral shipments. Hence, stock is only considered when it is in the shop or in the pipeline to the shop.

1.7 Reading guide

The remainder of this thesis is structured as follows. In Chapter 2 we investigate how the standard situation looks like, setting a base line for improvement for later chapters. We touch upon the technical details of the (R,s,nQ) model, and we set requirements for the solution. We conclude by answering research question 1 and its sub questions. We begin Chapter 3 by providing the theoretical perspective of the research. We start with literature on transforming sales to demand, followed by modelling the demand distribution. Next topic is the classification of SKUs based on their demand distribution, and we conclude by literature on implications of these classes and demand distributions on inventory models. Thereby we answer research question 2 and its sub questions. In Chapter 4 and 5 we make the core analyses of this research and answer research question 3 and 4 respectively, and their sub questions. We develop different models both safety stock and reorder point calculations, validate our approach, and we run an experiment. In this experiment, we test the different methods on a dataset, and analyse the results in Section 5.7. In *Conclusions and recommendations* we present the conclusions on the research questions, and we propose managerial recommendations, both concerning with the solution to be implemented itself, and how this should be done. In the final chapter of this thesis - *Discussion and future research* - we discuss this thesis process-wise, and we present topics for future research.

Current situation

For the first research question we combine the insights obtained from working with Slim4, with brief, informal meetings with its stakeholders at Slimstock. My daily supervisor and his colleagues are the main source of knowledge in this part. In the first section of this chapter we discuss Slim4's inventory model and we provide calculations for the two main parameters of this configuration: safety stock and the reorder point. Furthermore, cases from a client are used, based on the scope as described in 1.6, and some SKUs are zoomed into. In order to compare later performance, we set a baseline performance using standard methods. Finally, we detail out requirements our solution should adhere to, and finish the chapter with conclusions on the standard state of Slim4. The literature research gives a follow-up on the outcome of this chapter, and analyses in later chapters are partly based on the insights we gain by zooming into the (R,s,nQ) model. Throughout chapter *Current situation* we answer the following research question and sub questions:

1. How does the standard way of generating order policies using a (R,s,nQ) model at Slimstock look like?

   a. *How do the standard methods as used within Slim4 work?*

   b. *How big are the opportunities for improvement of handling multi-modal order line size SKUs?*

   c. *Under which circumstances are these opportunities highly relevant?*

   d. *What are Slimstock's requirements the solution should adhere to?*

2.1 Current standard inventory model

We focus on the inventory policy (R,s,nQ). In this policy, which best suits the scope as defined in Section 1.6, Slim4 is capable of dynamically calculating reorder parameters and generating replenishment orders. These parameters include the *safety stock* (ss) and the *reorder point* (s).

At the beginning of *review period* (R), the model compares the standard *inventory position* (IP), which equals the stock on hand plus pipeline inventory minus unfulfilled back-orders, with the reorder point. If this inventory position is smaller than or equal to the reorder point, the model generates a replenishment order advice. The *replenishment order quantity* that should at least be covered, i.e. the reorder point minus the inventory position right before ordering, consists of the safety stock and the *expected demand* [E(D)] during the cover period (*lead time*(L) + review period), abbreviated by $E(D_{L+R})$. The safety stock depends on the standard deviation of demand during the cover period $\sigma_{(D_{L+R})}$ and the *safety factor* k, which results from a target volume fill rate β^*, set by the client. The replenishment order quantity is transformed into a replenishment order advice, which consists of a number of batches with a minimum order size. Namely, when placing a replenishment order in Slim4, this has to be at least of the size of the *replenishment minimum order quantity* (e.g.: a full pallet). In case a bigger replenishment order is required in order to exceed the reorder point, n batches of the size of the *replenishment incremental order quantity*(e.g. a box) can be added to the total order quantity. Both are often equal and based on supplier contracts and/or profitability analyses (Slim4-developer, 2016a).
Slim4 is also capable of computing the replenishment order quantity based on the *Economic Order Quantity* (EOQ) and incorporating price discounts and logistic quantities (pallet, full truckloads, etcetera), instead of the difference between the reorder point and inventory position. However, we leave this out of our scope

for research, as we focus on inventory control with service level constraints. In Chapter 4 we develop our own mathematical inventory model, which is based on literature, yet it closely resembles a heavily simplified model of Slim4, as described above. Note: as this is a strongly abstracted version of Slim4, this does not reflect Slim4's true performance, but it follows a similar pattern. Hence, we do not research Slim4 itself, but we propose recommendations regarding this (R,s,nQ) model for Slimstock, which can in the end be implemented in Slim4.

2.1.1 Demand classes

Slim4 currently classifies SKUs based on their historical sales per period of time - not on transaction level. This research's opportunities for improvement occur in all of the demand classes, although we expect the most issues to rise in slow-moving SKUs with irregular-sized order lines. Little demand generally leads to few replenishments, making it more critical to order the right amount of SKUs at the supplier at the right moment (i.e.: the right inventory position). Furthermore, multi-modal order line size SKUs by their nature are demanded in non-unit and not strictly successively occurring order line sizes.

Slim4 divides SKUs in several demand classes, reaching from fast-moving SKUs to slow-moving SKUs, and everything in between. This is done based on the number of orders containing the SKU in a certain amount of time. Slimstock sets the threshold for fast-moving SKUs on at least 26 order lines per year (Slim4-developer, 2016a). These demand classes are automatically updated when the demand level changes. Demand is forecast while also taking into account trends, promotions, seasonal patterns, and exponential smoothing, if necessary.

2.1.2 Fast-moving SKUs

Our standard (R,s,nQ) model's method for determining safety stock for smooth demand (fast-moving and stationary) is the fill-rate safety stock determination. This equation accounts for either lost sales (Equation 2.1) or back-ordering (Equation 2.2) in case of insufficient inventory. This method incorporates the normal loss function as widely described in literature (Winston, 2003). Next, it ignores *undershoot* (the difference of the inventory position at the moment of ordering and the reorder point), since it employs the aforementioned batch replenishment quantity for exceeding the reorder point after replenishment. Later in Chapter 5, we research the effect the incorporation of undershoot has on the model's performance.

$$NL\left(\frac{s - E(D_{L+R})}{\sigma_{(D_{L+R})}}\right) = \left(\frac{1-\beta^*}{\beta^*}\right) \cdot \frac{Q}{\sigma_{(D_{L+R})}} \tag{2.1}$$

and

$$NL\left(\frac{s - E(D_{L+R})}{\sigma_{(D_{L+R})}}\right) = (1-\beta^*) \cdot \frac{Q}{\sigma_{(D_{L+R})}} \tag{2.2}$$

Let β^* be the target volume fill rate, Qr be the *replenishment lot size*: $\max(R * E(D), MOQ)$, and $\sigma_{(D_{L+R})} = \sigma_D \cdot \sqrt{L+R}$. The standard normal loss function restricts us to, as the name already reveals, modelling the demand during the cover period with the normal distribution. The standard *normal loss function* NL(k) for safety factor k is given in Equation 2.3.

$$NL(k) = \phi(k) - k \cdot (1 - \Phi(k)) \tag{2.3}$$

In Equation 2.3 we define $\phi(k)$ as the probability density function of the standard normal distribution, and $\Phi(k)$ as its cumulative density function. Safety stock is calculated by Equation 2.4.

$$ss = k \cdot \sigma_{(D_{L+R})} \tag{2.4}$$

Consequently, the reorder point as used in the (R,s,nQ) model is denoted in Equation 2.5.

$$\text{Reorder point} = ss + E(D_{L+R}) \tag{2.5}$$

2.1.3 Slow-moving SKUs

Slow-moving SKUs generally have a small number (26 order lines annually) of order lines to base a demand distribution on, which furthermore leads to the standard deviation of demand during a certain period being unknown. In order to deal with these characteristics, a specific heuristic is developed by Slimstock (Slim4-developer, 2016b). Next, there are also demand classes within Slim4 that work with discrete probability distributions like the Poisson distribution - in accordance with literature such as Silver, Pyke and Thomas (2017), which are not within the scope of this research. In our simplified model we do not differentiate between fast-moving and slow-moving SKUs. This furthermore has the advantage of equal comparisons, and further programming issues are avoided. For both fast-moving and slow-moving SKUs, the order line size distribution is clearly not included in the inventory model, as mentioned in Section 1.4.

2.2 Case study

For the analysis on the standard performance of the (R,s,nQ) model, we perform a case study in which we zoom in on how a client of Slim4 currently deals with the multi-modal order line size SKUs. To this extent we include the company whose data we also use in our experiment. This client is a big non-food retailing client of Slimstock, which we further refer to as *company X*. This company has more than 550 physical shops and 5,000 SKUs. Their partnership with Slimstock originates from 2009, and currently this client uses Slim4 for their distribution centre, web shop and partly for their shops. In the near future all of their shops will follow. We also use the data from a subset of the most-selling shops - that already use Slim4 - of this company in our data analysis in Chapter 4.

Company X has raised the issue we are tackling in this thesis. They regard to the issue from a very practical point of view, having much experience in non-food retail markets. In order to stay connected with this practical vision on the issue raised in Chapter 1, we investigate the standard performance of among others the multi-modal order line size SKUs. This is done through several meetings with company X's consultant/developer at Slimstock (Slim4-developer, 2016c) and a small test using a small dataset and inventory model.

Company X sells many SKUs that they know to have a multi-modal median order line size. They stick to this exact formulation of *knowing*, as it is currently not possible within the inventory model to automate the incorporation of this information. Company X currently resolves this issue by ensuring their stock on hand level to be at least a representative amount they want to show on the shelf, through a so called *insurance inventory*. Three different types of insurance inventory are used by company X at this moment: maximum of the calculated safety stock and insurance inventory (i.e.: not letting the proposed safety stock getting below the manually set insurance inventory level), overwriting safety stock, and overwriting the reorder point; the latter two speak for themselves. Company X uses insurance inventory for 93.59% of their SKU-shop combinations (a SKU-shop combination represents the record of SKU x at shop y), making this issue highly relevant for company X.

For the insurance inventory policy of company X: as this is a manually inserted value, it results in potentially bad performance. The client's inventory planner manually estimates - solely based on his own experience of previous sales and his/her preference on what amount looks nice on shelf - this insurance inventory level per product group. There are no calculations involved. In addition, the lack of automatically taking into account order line sizes, decreases the quality of demand forecasting, and consequently it results in worse performance. By Slimstock this is measured in more out-of-stock situations than desirable (although their own norm is not specified, company X expects an improvement) due to an insufficient stock level. Furthermore, the range of SKUs having insurance inventory is limited by the knowledge of the planner, and not purely data-driven. Henceforth, SKUs that might by multi-modal are ought not have insurance inventory, if the planner does not know about its multi-modal nature. Finally, manually calculating this insurance stock is time-consuming and therefore for large clients already a problem on its own.

Hence, there are opportunities for improvement of the inventory model as used by Slim4, in order to make it fit standard practice of SKUs that are frequently sold in the same quantities. Finally, when insurance inventory settings are 'freed' from usage as extra buffer stock for multi-modal order line size SKUs, these insurance stocks can be used again for the purpose they originally serve, making sure a user-controlled number of products is on shelf in order to show the product to the customer.

2.3 Performance standard mdoel

We tested a small dataset (the same we use for the final experiment) from company X, for which we first elaborate on the standard baseline performance. Results are presented in Table C3 in Appendix C.

2.3.1 Observed volume fill rate

Recall that this method includes a target fill rate approach for the safety stock, and by adding the expected demand during the lead time plus review period, the reorder point is calculated. As we see in Table 2.1, the problem of the standard method arises in both types of SKU. In Table C1 in Appendix C we provide information on the mentioned SKUs. However, both types of SKUs are very different in terms of median order line size, number of orders per period of time. But most of all, both types of SKUs differ in their nature, since the multi-modal order line size SKUs are very often demanded in particular order line sizes, whereas unimodal order line size SKUs lack this property. Henceforth, they need different strategies for improvement. Multi-modal order line size SKUs show an under-performance of on average 10.30%-point and 10.35%-point for lost sales and back-ordering respectively.

			Lost sales		*Back-ordering*	
SKU	β^*	**Number of order lines**	β	$\beta - \beta^*$	β	$\beta - \beta^*$
MM SKU 1	98.03%	6.73	91.71%	-6.32%	91.80%	-6.22%
MM SKU 2	96.17%	2.68	97.51%	1.34%	97.51%	1.34%
MM SKU 3	96.94%	1.78	94.12%	-2.82%	94.12%	-2.82%
MM SKU 4	97.77%	7.06	86.22%	-11.54%	86.83%	-10.94%
MM SKU 5	96.98%	5.79	88.52%	-8.46%	87.88%	-9.09%
MM SKU 6	95.96%	3.63	94.77%	-1.19%	94.53%	-1.43%
MM SKU 7	97.04%	4.83	90.04%	-6.99%	89.96%	-7.07%
MM SKU 8	97.12%	3.62	86.52%	-10.60%	87.40%	-9.73%
MM SKU 9	97.94%	9.99	94.57%	-3.37%	95.63%	-2.31%
MM SKU 10	96.78%	3.58	92.26%	-4.52%	90.65%	-6.13%
MM SKU 11	97.87%	3.61	89.93%	-7.94%	90.76%	-7.11%
MM SKU 12	97.48%	3.60	84.05%	-13.43%	84.55%	-12.93%
MM SKU 13	96.59%	2.90	79.10%	-17.49%	79.21%	-17.39%
MM SKU 14	96.93%	3.74	91.19%	-5.74%	90.90%	-6.02%
MM SKU 15	96.62%	3.40	80.21%	-16.41%	78.64%	-17.98%
MM SKU 16	98.17%	5.75	81.72%	-16.45%	81.18%	-16.99%
MM SKU 17	97.35%	4.95	79.14%	-18.21%	80.77%	-16.58%
MM SKU 18	97.70%	3.56	81.16%	-16.54%	79.11%	-18.59%
MM SKU 19	98.78%	12.97	90.18%	-8.60%	89.66%	-9.12%
MM SKU 20	97.22%	3.70	77.68%	-19.53%	78.68%	-18.53%
MM SKU 21	97.48%	3.15	78.99%	-18.48%	79.73%	-17.75%
MM SKU 22	97.19%	2.21	83.89%	-13.30%	82.82%	-14.37%
Avg. MM	97.28%	4.69	86.98%	-10.30%	86.92%	-10.35%
UM SKU 23	98.66%	17.26	93.02%	-5.64%	94.56%	-4.10%
UM SKU 24	99.00%	75.30	76.45%	-22.55%	77.42%	-21.58%
UM SKU 25	98.33%	10.87	94.11%	-4.22%	93.87%	-4.45%
UM SKU 26	98.59%	53.32	97.55%	-1.04%	97.63%	-0.96%
UM SKU 27	98.98%	40.62	95.47%	-3.51%	96.09%	-2.90%
UM SKU 28	98.96%	35.48	92.45%	-6.51%	92.42%	-6.54%
UM SKU 29	97.61%	11.80	96.60%	-1.01%	96.49%	-1.12%
UM SKU 30	98.07%	58.96	87.42%	-10.66%	87.52%	-10.55%
UM SKU 31	98.88%	64.72	82.22%	-16.66%	82.71%	-16.17%
UM SKU 32	98.78%	22.68	94.68%	-4.11%	94.64%	-4.14%
UM SKU 33	97.54%	42.44	94.34%	-3.20%	94.87%	-2.67%
UM SKU 34	97.95%	40.98	95.19%	-2.76%	95.08%	-2.87%
UM SKU 35	96.57%	3.30	90.89%	-5.68%	91.75%	-4.82%
UM SKU 36	98.97%	50.73	94.54%	-4.44%	94.75%	-4.22%
UM SKU 37	98.94%	24.34	94.92%	-4.02%	95.20%	-3.75%
UM SKU 38	98.23%	13.33	98.29%	0.06%	98.19%	-0.04%
Avg. UM	98.38%	35.38	92.38%	-6.00%	92.70%	-5.68%
Difference	97.83%	-86.74%	89.68%	-8.15%	89.81%	-8.02%
Avg. total	-1.12%	20.04	-5.85%	71.82%	-6.23%	82.27%

Table 2.1: *Differences between target and volume fill rates per SKU*

Unimodal order line size SKUs on the contrary under-perform on average 6.00%-point and 5.68%-point for the two out-of-stock settings. Thus, under-performance occurs in both types of SKUs, but the standard model performs worst on multi-modal order line size SKUs. Concluding - when also considering the SKU information from Table C1 -, compared to unimodal order line size SKUs, multi-modal order line size SKUs from the first category:

- are ordered at the supplier in MOQ/IOQs (for retailers, MOQ and IOQ are equal in practice) that are on average 51.69% smaller. Only 22.73% of the multi-modal order line size SKUs have a MOQ/IOQ larger than 6, compared to none of the unimodal order line size SKUs having a MOQ/IOQ less than 10;
- have a median order line size that is on average 2 times bigger. Only two of the unimodal order line size SKUs have a median order line size bigger than 1;
- were 86.74% less often sold (number of order lines of the SKU) during the performance measurement part. Only one of the multi-modal order line size SKUs has an average number of orders during four months of more than 10, and only one of the unimodal order line size SKUs has an average number of orders during four months of less than 10;
- were given target fill rates that were on average 1.12% lower. Note: average target fill rates for multi-modal order line size SKUs is still very high, namely 97.28%;
- have lower cumulative probabilities for the first two order line sizes *1* and *2*, which follows directly from their classification scheme. $P(X \leq 1)$ and $P(X \leq 2)$ amount 26.26% and 57.70% on average for multi-modal order line size SKUs, whereas these values equal 56.33% and 81.43% respectively for unimodal order line size SKUs.

Based on this comparison, we can conclude that the multi-modal order line size SKUs are characterized through being ordered at the supplier in small batches (here in at most packs of six), have high median order line sizes (here on average 2.50), and are slow-moving (here with not more than 30 customer orders per year). The small batches make sense due to the fact that the standard methods tend to overshoot more for larger MOQ/IOQ, resulting in higher service level in general. Next, the high median order line sizes lies in the very nature of these products, customers demand relative constant, high amounts of this SKU. The number of orders is a result of the fact that inventory control - in terms of attaining a specified service level - is more critical for slow-movers, as there are not much opportunities for replenishment, as there are few moments of demand.

2.3.2 Zooming in

We now dive into some SKUs that function as example for the issue we are facing, in order to gain more insights on the size of potential improvement. We present three SKUs, each having a different order line size distribution, and descriptive statistics and input parameters for the underlying statistical distributions. The other SKUs from our dataset containing roughly follow one of these three SKUs patters. Next to the categories mentioned below, there is a fourth category, which we regard to as *not interesting*, in which the order line size distribution is highly compact centred around one particular order line size. As this category would provide hardly no difference with standard lead time demand calculations, this category is not interesting in our research. Moreover, a SKU's heavily centred order line size distribution would provide no more information than regular average order line size and customer arrival rate or even lead time demand characteristics.

We test two out-out-stock assumptions, namely lost sales and full back-ordering, and use corresponding formulae. The first SKU is a low-budget mug, with an average target fill rate of 98.03% and an average observed volume fill rate of 91.71% for lost sales and 91.80% for back-ordering, of which the order line size distribution is depicted in Figure 2.1a.

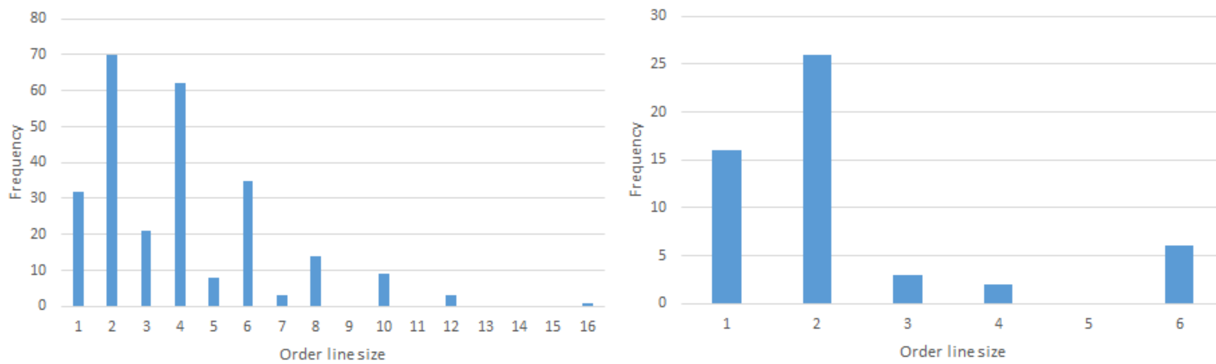

(a) *Figure of multi-modal demand mug* (b) *Figure of multi-modal demand breakfast plate*

Figure 2.1: *Multi-modal order line size figures*

In Figure 2.1a, we can clearly distinguish between even and odd order sizes, both follow a very smooth pattern, and the even order sizes having the biggest right hand side. However, combined they construct a multi-modal distribution. This roughly implies that the empirical probability of having an order line of six mugs, is a bit bigger than an order line of one mug, and considerably bigger than order line sizes of all other order sizes except two and four.

Next, we focus on breakfast plates (which come from a series of dining room equipment, although we do not research multi-item relations), of which the order line size distribution is depicted in Figure 2.1b. This specific SKU resembles another sign of multi-modality, which surely cannot be modelled through a single probability distribution. Having a large gap in the order line sizes between 3-5, it can be stated that inventory levels of 4 and 5 are useless, as they are insufficient for order lines of size 6, and they can fulfil 2 order lines of size 2 and 0 or 1 order line of size 1 respectively, which is quite an overkill in this setting. The target volume fill rate of this SKU amount 96.62%, and its observed volume fill rates measure are 80.21% and 78.64% for lost sales and back-ordering respectively. Thus, some multi-modal order line size SKUs perform extremely worse than desired through the standard methods.

Finally, we describe a third branch, for which we take the example of a certain cutlery tray, of which the order line size distribution is depicted in Figure 2.2.

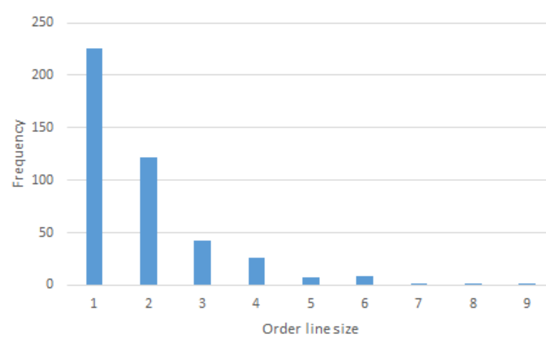

Figure 2.2: *Order line size figure of unimodal order line cutlery tray*

The majority of SKUs sold by retail companies comprises SKUs that are most often sold per one, and only sometimes order line sizes of more than 1 occur (other solely occurring order line sizes are also possible of course). Furthermore, the order line size distribution is (approximately) monotonically decreasing, meaning $p(X = x) \geq p(X = x + 1)$, with x being an arbitrary order size. In general, this pattern is unimodal. We take this type of SKU in account during the remainder of our research, as we want to know how they perform under other heuristics. The target volume fill rate of this SKU amounts 98.33%, and its observed volume fill rates equal 94.11% and 93.87% respectively.

2.4 Conditions for implementation

Next, we investigate under which circumstances, given the scope in Section 1.6, the improvements are most relevant. We restrict ourselves to non-technical, quality-related requirements, as technical requirements are irrelevant for the scope this research (Slim4-developer, 2016b). In other words, we do not take into account the computation time, this is something Slimstock should look into when implementing our solution. Furthermore, it is important that the user-friendliness of Slim4 is not damaged by implementing our solution (Slim4-developer, 2016b).

The most important condition for implementation is as effective as it is intuitive, namely that the proposed solution should at least improve Slim4's performance on handling inventory for all SKUs. Two things are important here, that its performance should not be worse, and preferably it should be improved, although there is no a priori improvement threshold. We quantify performance as the percentage with which the volume fill rate is improved, the average stock is altered (as much as decrease as possible, or the smallest increase), and the order line fill rate as third KPI (fraction of order lines that should be immediately fulfilled from stock). Second aspect is that this condition holds for all SKUs for which the client holds stock and places replenishment orders. For example, an improvement for 95% of the SKUs affected, and a decrease of performance for the other 5% of the SKUs, is not a preferable solution, although future inspection of separate under-performing SKUs is possible in the future.

2.5 Conclusions

Now that we have dived into Slim4 and the current way of working with this standard inventory management system, we can answer the sub questions and the research question:

a. *How do the standard methods as used within Slim4 work?*
Slim4 has many different configurations and methods for calculating safety stock and setting reorder points. We focus on the general (R,s,nQ) model, and we use its calculations for both fast-moving and slow-moving SKUs through the normal distribution. Out-of-stock situation can either be handled as lost sales or as back-ordering.

b. *How big are the opportunities for improvement of handling multi-modal order line size SKUs?*
In our preliminary problem size analysis of Section 1.3, we found that around 45% are potentially multi-modal. This group has a potential of 8.31%-point or 8.42%-point higher volume fill rates for lost sales and back-ordering respectively, lie in breaking total demand per period of time into a customer arrival rate and an order line size component.

c. *Under which circumstances are these opportunities highly relevant?*
We found that most issues with standard methods often arise in SKUs that are ordered at the supplier in small batches (for company X in at most packs of six), have high median order line sizes (for company X on average approximately 2.5), are slow-moving (for company X with not more than 30 customer orders per year), have small probabilities of selling less than 1 or 2 item(s) per order line, and in general have slightly lower target fill rates.

d. *What are Slimstock's requirements the solution should adhere to?*
The following requirements for the solution are set:

a. The observed volume fill rate should have the highest volume fill rate, with the smallest increase (or highest decrease) of observed average stock;

b. The under [a.] mentioned improvement should hold for as most SKUs as possible;

c. The user-friendliness of Slim4 is not damaged.

Now that we have answered the three sub questions of research question 1, it is time to conclude on what the standard situation within a (R,s,nQ) model looks like.

How does the current way of generating order policies using a (R,s,nQ) model at Slimstock look like? Slim4 takes into account both the total demand during a period of time, and the customer arrival rate, but no information on order line size is included. However, a large number of SKUs have such a specific order line distribution, hence it would be interesting to include this element into the inventory model.

This brings the following implications for our literature review. We look for a solution that yields improvement for as many SKUs as possible, this should also be reflected in the literature review. Order line sizes can follow a quite specific, irregular path, which might be well fitted through regular probability distributions, so we need to look for creative ways to deal with the typical order line size distributions.

Literature

We start by a brief review on the specific type of inventory model we investigate. Next, we need to transform historical sales data into true demand, taking into account the lost sales. In addition, we need scientific literature on possibilities for accounting for empirically-modelled order line size distributions, through a set of heuristics for calculating safety stock and the reorder point. We conclude this chapter by combining the two themes of true customer demand, and methods to implement these issues in an inventory model, and we briefly elaborate on performance measurement and sensitivity analyses. This literature review is concluded by answering the research questions below:

2. What can we learn from literature about optimization of multi-modal order line size SKUs inventory?

   a. *How can we identify multi-modal order line size SKUs the best way?*

   b. *How can we incorporate the 'multi-modal order line size effect' into inventory management?*

3.1 Transforming sales into demand

As lost sales are unobserved, we have to deal with so called *censored data*. In practice this yields in worse performance as sales data on big orders miss as these are likely to be lost due to stock-out. Assuming sales to equal true demand could lead to insufficient inventory levels.

There are however some methods to deal with these censored data. One way is to enrich the right hand side of the order line size distribution (assuming type II censoring, (Johnson, Kemp, & Kotz, 2005)), with a theoretical distribution. To this extent, the *exponential distribution* is employed (Bratley, Fox, & Schrage, 1986), as well as the *Weibull distribution* (Bisi, Dada, & Tokdar, 2011) - a practical spreadsheet for Weibull right hand side enrichment was also found -, and the geometric distribution (Axsäter, 2006). Next, two authors discuss a method in which they transform sales into demand through an *empirical sales distribution* (Lau & Lau, 1996). They assume that when inventory is zero after an order took place, the true demand must be greater than or equal to the sales quantity. Furthermore, they only assume larger order sizes (right hand side) to be censored. As we do not have access to inventory level data, nor is it valid in our setting to only accommodate for larger order sizes, we do not go further into this line of work.

3.2 Order line size modelling

In order to further characterize the demand distribution, we investigate manners to analyse this demand distribution. We highlight some of the various methods, coming from different fields of research (some from pure statics, or biology). As we mentioned in Section 1.3, a main issue of our research is to go from a general lead time demand perspective to a division of the demand process into the estimation of demand frequency and order line size functions, like in Equation 3.1 (coming from compound Poisson distribution theory).

$$\begin{aligned} E(\text{Demand during lead time}) = E(\text{Number of customer arrivals during lead time}) \cdot E(\text{Order line size}) \\ E(D_L) = E(\lambda_L) \cdot E(X) \end{aligned} \tag{3.1}$$

The empirical reasons are provided in Section 1.3, but also research underlines this, e.g. it works better than applying a *regular exponential smoothing procedure* (Willemain, Smart, Shockor, & DeSautels, 1994).

3.2.1 Compound Poisson process

There are two main possibilities for divisions, both comprise modelling the demand process as stochastic process, either as discrete Bernoulli process, or Poisson process. The differences between both stochastic processes are given in Table 3.1.

	Bernoulli process	**Poisson Process**
Arrival time	Discrete	Continuous
Customer arrival rate	p per trial	λ per unit of time
Inter-arrival time	Geometric	Exponential
Distribution of number of arrivals	Binomial	Poisson

Table 3.1: *Differences between Bernoulli process and Poisson process*

The fact that the latter is more often used within inventory management research, see Table 3.2, is a advantage for finding literature, and it serves as verification for using this process. Although transaction data is aggregated on day level (that is, the exact time of a transaction is unknown), we expect that this does not violate the continuous exponential distribution for inter-arrival time. Moreover, inter-arrival time between customers arriving at the shop for buying a certain SKU can still be regarded memoryless, enabling the exponential distribution. Thus, it is convenient to prefer this over the Bernoulli process.

"Inventory" AND (management OR control) AND "Poisson process"	255 hits	Scopus, 9-3-2017
"Inventory" AND (management OR control) AND "Bernoulli process"	3 hits	Scopus, 9-3-2017

Table 3.2: *Search queries comparison*

As the Poisson process can constitute a compound Poisson distribution, we consider modelling the demand as combination of two parts, namely an arrival process of customers, and the order line size on the other hand. This has implications for the rest of our literature research, as we focus on modelling the order line size distribution.

3.2.2 General order line size distributions

order line sizes of fast-moving SKUs can be modelled with a continuous probability distribution - such as the Normal distribution (Croston, 1972) or *lognormal* (Willemain et al., 1994) - yet they should be approximated by discrete probability distributions as well. Other authors however found that order line size can also be well modelled by discrete distributions like the geometric distribution, or its generalization, the negative binomial distribution (Hadley & Whitin, 1963), (Johnston, Boylan, & Shale, 2003). In addition, we consider the empirical distribution for the order line size, as this would very intuitively forecast the typical multi-modal historical pattern as demand forecast. However, there is a big drawback connected to the empirical distribution: order line sizes that by coincidence did not occur in the past sales, can never occur in the future forecast, as their relative frequency is set to zero. More generally, as relative frequencies also remain the same for the forecast, this way of modelling is quite static, which is a potential problem for reorder policies. Rare order line sizes that by coincidence occurred in the past time, are thereby directly incorporated in the forecast, potentially leading to bad performance compared to a theoretical distribution (Law, 2015). This mainly is relevant for SKUs with rather few historic sales and/or a multi-modal demand pattern, as the more continuous demand pattern are all right with empirical order line size distributions, as they are already smooth. Nevertheless, the empirical distribution being extremely understandable and accurate for stable order line sizes (with which we deal with, as multi-modal SKUs like mugs are always sold through roughly the same order line size distribution), yields a big advantage. Henceforth, we include it in our further analysis.

3.2.3 Package Poisson distribution

We can also regard to the order line size as constant *packages*, e.g.: when demand occurs, it always equals 4 items (Friend, 1960). This *package Poisson distribution* yielded "satisfying results" compared to separate Poisson distribution (for slow-moving SKUs), Normal distribution (for fast-moving SKUs), stuttering Poisson distribution (for lumpy SKUs) - in which the arrival process of demand is Poisson-distributed, and order line size is geometric-distributed, and a stuttering Poisson distribution with Normal-distributed order line sizes (Vereecke & Verstraeten, 1994). This package Poisson distribution could be useful as replacement for a set of different probability distributions, as these authors also state that problems arise when SKUs change from classes, as their demand probability distribution changes.

3.2.4 Hurdle, zero-inflated, and zero-truncated distribution

Another possibility for modelling order line size within a Poisson process, involves *hurdle*, *zero-inflated* and/or *zero-truncated* probability distributions. First of all, both hurdle and zero-inflated models increase the probability of having a 'zero' in a dataset, yet, they differ in the origin from this zero. A zero-inflated model is a mixture distribution comprising a zero-generating function, and a regular Poisson or Negative binomial counting distribution. In Section 3.2.5, we discuss mixture distributions in more detail. Henceforth, zeros can come from both the *only-zero* and *regular* distribution (Hu, Pavlicova, & Nunes, 2011). A hurdle distribution on the other hand is a piecewise function, consisting of a zero-generating function and a zero-truncated distribution (i.e. a distribution that cannot generate zeros). Hence, all zeroes are coming from one source. A great example is given by Hu, Pavlicova & Nunes (2011), about a cocaine survey, in which a secondary result is the number of cigarettes that were smoked in the last month. If one is considered a smoker, it is by its nature impossible to smoke zero cigarettes, and vice versa.

"zero-inflated" AND ("Poisson" OR "Negative binomial")	1,369 hits	Scopus, 20-12-2016
"general inflated" AND ("Poisson" OR "Negative binomial")	8 hits	Scopus, 20-12-2016

Table 3.3: *Search queries comparison*

As can be seen in Table 3.3, it is safe to assume more research is done about zero-inflation. Following the big amount of order line sizes of 1, see figures A2, and A3, and A1, it would be interesting to investigate the performance of a zero-inflated (transform all order line size by subtracting exactly 1, in order to yield the well-studied zero-inflation) probability distribution. As the Poisson distribution is known for its equal mean and variance, this is not the best distribution for our demand data. The negative binomial distribution, has a variance-to-mean ratio of 1/p (with p being the probability of success), thereby always being larger than 1. This makes the negative binomial distribution attractive for modelling demand for retailing applications, which tend to have high variability (Agrawal & Smith, 1996). One final advantage of the negative binomial distribution is its extendibility to the geometric distribution - in fact, the negative binomial distribution is a generalization of the geometric distribution - and the Poisson distribution, as approximation. Furthermore, it is related to the Binomial, *Beta*, and *Gamma* distribution (Cook, 2017).

3.2.5 Mixture models

A more elaborated way of modelling our sales (which we assume to equal future demand as mentioned earlier) through a demand distribution, is by using finite mixture models. This would lead to a *mixture* of probability distributions, combined through some mixture weights. A very complete overview of this topic is given by Lindsay (1995), in which mixture models are described as follows: "The mixture model arises when one samples from a population that consists of several homogeneous sub-populations, which we call the components of the population"(Lindsay, 1995). This also applies to our research, as the multi-modal demand distribution is caused by strictly differing reasons to buy a specific amount (either for replacing one broken mug, or replacing all dining equipment), resulting in sub-populations. The difficulty in this finite (stands for a finite number of components) is finding a satisfying amount of components (i.e.: identifying the different sub-populations). Modelling too many components usually over-fits the data, resulting in too high order line size forecasts, and it is hardly intuitive if the order line size distribution is modelled with a mixture consisting of a large number of models (Xu & Chen, 2015). Considering too few on the other hand could lead to a poor description of the underlying model (Figueiredo & Jain, 2002).

Trade-off within order selection

Xu and Chen(2015) state that this process of selecting the right number of components (i.e.: order selection) can be seen as a "trade-off between the model complexity and the goodness of fit". Furthermore, the number of modes does not necessarily equal the number of models within a mixture distributions (Chen & Khalili, 2009). We later return to the specific topic of multi-modality.

Different methods for order selection

Henceforth, *order selection* is a topic for research on its own. Figueiredo and Jain (2002) give a great overview of both deterministic criteria, including several Bayesian and information criteria, and stochastic criteria, including Markov Chain Monte Carlo methods and Cross-validation. For these Markov chains, the following actions are possible (Richardson & Green, 1997):

a. updating the weights w;
b. updating the parameters;
c. updating the allocation z;
d. updating the hyperparameter β;
e. splitting one mixture component into two, or combining two into one;
f. the birth or death of an empty component.

Xu and Chen (2015) furthermore give an overview of even more methods, namely distance measure-based approaches and hypothesis testing-based approaches. Often, focus is put on two very popular information criteria, namely the Akaike Information Criterion (Akaike, 1973)) and Bayesian Information Criterion

(Schwarz, 1978). In addition, we highlight a recently developed penalized likelihood approach called that is combination of the modified likelihood (Chen & Kalbfleisch, 1996), and the variable selection method called Smoothly Clipped Absolute Deviation [SCAD] (Fan & Li, 2001), resulting in a Mixture SCAD (MSCAD). The latter is compared by Chen and Kahili (2008) with most of the briefly mentioned other criteria, and found to yield better performance. To be more precise: it converged to the right number of Poisson distributions within the mixture, both more often and with fewer iterations than AIC and BIC - and several other aforementioned approaches (Chen & Khalili, 2009). Unfortunately, this article is not quite known to colleague-researchers, as depicted by 196 views on its journal - the Journal of the American Statistical Association - and 22 citation on Scopus (dated 20-12-2016). Furthermore, no implementation in any software package can be found, nor a pseudo-code algorithm, making it difficult to re-create this algorithm within the scope of our research.

Parameter estimation of mixture distribution

Once an adequate order selection method is selected, this can be implemented in the *Expectation Maximization*(EM) algorithm. This is a widely-known, powerful, two-step algorithm which originates from the problem of unobserved data (Dempster, Laird, & Rubin, 1977). Through finding the *Maximum Likelihood Estimate* (MLE), it seeks a correction, which can also be considered as a mixture of data-generating probability distributions. Implementations for mixing probability distributions are found in both archaic tools, Java applets, and data mining tools.

3.2.6 Other methods

Finally, other methods for modelling demand distribution include the following:

- *Batch Markov Arrival Process*(BMAP): a point process characterized by Markov-modulated batch arrivals of random size. This makes it a generalization of many processes, including the *Markovian arrival process* (MAP), the Poisson process, and the Markov-modulated Poisson process (Cordeiro & Kharoufeh, 2010). These authors provide a very extensive overview of these processes, in which they for example elaborate that the assumption of exponentially-distributed inter-arrival times is dropped in BMAP, and by its nature it enables the possibility of having differing batch sizes - hence, we do not need a compounding order line size distribution any more. Finally, both EM-algorithms and estimations are known (Breuer, 2002), (Breuer & Kume, 2010), and (Lucantoni, 1991).

- Kernel density estimation: a non-parametric way of modelling data, through a specific *kernel*. Combining these kernels, each with a bandwidth of values it should model, resembles the empirical data. A possible probability distribution is the Normal distribution, yet other models are also possible, along with all other geometrically based kernels (Tsybakov, 2008).

3.3 Classification of multi-modal order line size distribution

We now return to the topic of multi-modality. Modes can be discovered through so called *Gaussian mixture model-based parametric bump hunting methodologies* (Mukhopadhyay, 2017). The author also provides a description of the algorithm developed for discovering modes within a dataset, by the name of “LPMode”. Unfortunately, to the best of our knowledge, there is no public implementation available. In addition, demand would be assumed to be normally distributed, which is a too limiting restriction for our research. Also, no other classification schemes that really incorporate the order line size distribution were found.

3.4 Inventory management: setting safety stock and reorder point

As mentioned in Section 3.2 and in Section 3.2.2 in specific, we model demand either in conventional methods like normal or compound Poisson, or as compound Poisson process with a certain arrival rate of the customers, and compounding empirical distribution for the order line size. Henceforth, we look for methods to incorporate these demand distributions in inventory models, in order to improve performance of the multi-modal order line size SKUs. For this research we stick to the restrictions as set in 1.6. This model is better known in scientific literature as a single-echelon, periodic-review, batched order quantity model with constant lead times and service level constraints, in a purchase-to-stock environment. Focus is put on the lost-sales variant, in case of insufficient stock, although we also test performance for back-ordering and a hybrid own-developed policy.

We provide a brief recall from Section 2.1. Products in our (R,s,nQ) model are handled per piece, taking into account seasonal and promotional effects (yet we focus on stationary demand in our analysis), and a First Come First Serve allocation rule for customers is employed. Demand is assumed to be strictly positive, as negative demand implies SKUs being returned to the shop, and zero demand is not recorded for it is no transaction. Thus, non-positive demand is not included in setting reorder parameters or generating replenishment orders. Furthermore, the assumption of being in a retail business-to-consumer segment of the market is included.

Literature mainly focusses on back-ordering systems, as they are easier to analyse, and mainly used by the industry (Bijvank & Vis, 2011). For a great overview of literature working with the back-ordering assumption we refer to (Williams & Tokar, 2008). In addition we refer for inventory modelling in general to books such as Axsäter (2006); Silver, Pyke, and Thomas (2017); and Winston (2003).

3.4.1 Safety stock calculation

We address a number of different methods for calculating safety stock and reorder point, in a service-level, periodic review situation. We identify *Type I* and *Type II* service measures, respectively the cycle service level and fill rate (Nahmias, 1993). For all formulae it holds that the demand distribution is assumed to be known, so we consider averages and standard deviations of demand during a period, instead of forecasting errors.

Cycle service level

The first method exploits the cycle service level, which is defined as the "probability of having no stock-out within a replenishment cycle" (Silver et al., 2017). Winston provides a standard equation, assuming normal demand during the cover period, not differing between back-ordering or lost sales, see equations 3.2 and 3.3.

$$P(E(D_{L+R}) \leq s) = P\left(\frac{E(D_{L+R}) - \mu}{\sigma_{(D_{L+R})}} \leq \frac{s - \mu}{\sigma_{(D_{L+R})}}\right) = \Phi(z) = \alpha \tag{3.2}$$

$$\text{Safety stock} = z \cdot \sigma_{(D_{L+R})} \tag{3.3}$$

Let α be the target cycle service level, $E(D_{L+R}) = E(D) \cdot (L + R)$, $\sigma_{(D_{L+R})} = \sigma(D) \cdot \sqrt{L + R}$, and Φ the inverse cumulative standard normal distribution. This method works with the (R,s,Q) model, although it is also easily transformed into its continuous review counterpart.

Target volume fill rate

The next equation comprises a target volume fill rate, as defined in the glossary. Demand during cover period is assumed to be normal again, the function differs per out-of-stock setting, Equation 3.4 for the lost sales case, and Equation 3.5 for the back-ordering case.

$$NL\left(\frac{s - E(D_{L+R})}{\sigma_{(D_{L+R})}}\right) = \left(\frac{1-\beta^*}{\beta^*}\right) \cdot \frac{Q}{\sigma_{(D_{L+R})}} \tag{3.4}$$

$$NL\left(\frac{s - E(D_{L+R})}{\sigma_{(D_{L+R})}}\right) = (1-\beta^*) \cdot \frac{Q}{\sigma_{(D_{L+R})}} \tag{3.5}$$

Let β^* be the target volume fill rate, and Qr be the replenishment lot size: $\max(R * E(D), MOQ)$ (van Donselaar & Broekmeulen, 2013). The standard normal loss function restricts us to, as the name already reveals, modelling the demand during the cover period with the normal distribution. The standard normal loss function NL(k) for safety factor k is given in Equation 3.6.

$$NL(k) = \phi(k) - k \cdot (1 - \Phi(k)) \tag{3.6}$$

Safety stock is then calculated by Equation 3.7.

$$\text{Safety stock} = k \cdot \sigma_{(D_{L+R})} \tag{3.7}$$

This method works with the (R,s,Q) model, although it is also easily transformed into its continuous review counterpart.

Approximation Tijms and Groenevelt

The third method incorporating either the normal distribution or Gamma distribution (the latter is not incorporated in our research) for demand during the cover period we have found in literature for determining the safety stock is coming from renewal theory (Tijms & Groenevelt, 1984), including the same back-ordering to lost sales transition of $1 - \beta^*$ to $\dfrac{1-\beta^*}{\beta^*}$ as in Section 3.4.1. This method works with the (R,s,S) model, and we transform S-s into Q, for our (R,s,nQ) model, hence limiting to a (R,s,Q) model. See Equation 3.8 for the back-ordering case, and Equation 3.9 for the lost sales case.

$$J_u(k) = \frac{1}{\sigma^2_{(D_{L+R})}} \left[2 \cdot (1-\beta^*) \cdot E(D_{L+R}) \cdot \left[S - s + \frac{\sigma^2_{D_R} + E(D_R)^2}{2 \cdot E(D_R)} \right] \right] \tag{3.8}$$

$$J_u(k) = \frac{1}{\sigma^2_{(D_{L+R})}} \left[2 \cdot \frac{(1-\beta^*)}{\beta^*} \cdot E(D_{L+R}) \cdot \left[S - s + \frac{\sigma^2_{D_R} + E(D_R)^2}{2 \cdot E(D_R)} \right] \right] \tag{3.9}$$

Assumptions for using this formula are smooth demand, a target volume fill rate of at least 90%, the review period should not be too small compared to the lead time. Finally, the normal distribution is only reasonable when the coefficient of variation of demand during the cover period ($CV_{D_{L+R}}$) does not exceed 0.5. In case the $CV_{D_{L+R}}$ does exceed 0.5, modelling demand with a Gamma distribution (not further taken into account in this thesis) is more appropriate due to a big variability leading to a high probability of negative demand with the normal distribution (Silver et al., 2017). The $J_u(k)$ is another special function from the standard normal distribution (like the normal loss table), originating from the seminal book by Hadley and Whitin (1963). We provide the formula in Equation 3.10, with safety factor k.

$$J_u(k) = (1 + k^2) \cdot [1 - \Phi(k)] - k \cdot \phi(k) \tag{3.10}$$

Yet, it can also be approximated by a rational function as provided by Tijms and Groenevelt (1984). As this approximation has an maximum expected error smaller than 2.3 x 10^{-4}, for -4 $\leq k \leq$ 4, and it runs faster than looking up k-values in a table, we employ these rational functions, which are depicted in Equation 3.11.

$$k = \frac{a_0 + a_1 z + a_2 z^2 + a_3 z^3}{b_0 + b_1 z + b_2 z^2 + b_3 z^3} \tag{3.11}$$

, with $z, a_0, a_1, a_2, a_3, b_0, b_1, b_2, b_3$ obtained through:

for $0 \leq J_u(k) \leq 0.5$

$$\begin{array}{l|l} z = \sqrt{ln(1/J_u(k)^2)} & \\ a_0 = -4.18884136 \times 10^{-1} & b_0 = 1 \\ a_1 = -2.5546970 \times 10^{-1} & b_1 = 2.1340807 \times 10^{-1} \\ a_2 = -5.1891032 \times 10^{-1} & b_2 = 4.4399342 \times 10^{-2} \\ a_3 = 0 & b_3 = -2.6397875 \times 10^{-3} \end{array}$$

for $J_u(k) > 0.5$

$$\begin{array}{l|l} z = J_u(k) & \\ a_0 = 1.1259464 & b_0 = 1 \\ a_1 = -1.3190021 & b_1 = 2.8367383 \\ a_2 = -1.8096435 & b_2 = 6.5593780 \times 10^{-1} \\ a_3 = -1.1650097 \times 10^{-1} & b_3 = 8.2204352 \times 10^{-3} \end{array}$$

Note: we took the newer, and thereby more accurate rational function of Silver, Pyke, Thomas (2017), hence the earlier mentioned maximum error rate is even smaller. The safety stock *ss* ultimately is given in Equation 3.12.

$$\text{Safety stock} = k \cdot \sigma_{(D_{L+R})} \tag{3.12}$$

Expected service levels with compound Poisson demand

Another method of calculating safety stock is by modelling our demand as compound Poisson distribution. The most realistic configuration for our needs is having the demand process modelled as compound Poisson process, henceforth with Poisson number of customers per time interval, e.g. during the cover period, and the order line size distribution used as compounding distribution. Axsäter (2006) provides a very tractable method for determining safety stocks through both volume fill rate and ready rate constraints for continuous review (s,Q) models with back-ordering. The continuous review model can be extended to a periodic review model, by adding the review period to the lead time when modelling demand. This model works with partial back-ordering, so when we introduce this method in a lost sales environment, we must keep in mind that the real performance might deviate a bit. Furthermore, since we assume that only complete order lines are delivered, the resulting reorder points might be too low compared to a full back-ordering model. Due to time restrictions of our project, we refrained from adjusting this model to a full back-ordering setting. However, we propose a direction for further research on this adjustment. We derive reorder point from calculations in the same manner as described in Equation 3.26.

The equations are as provided in equations 3.13, and 3.14. The first two equation provides us the probability of the stock on hand being j, with order line sizes k.

$$P(SoH = j) = \frac{1}{Q} \sum_{x=max\{s+1,j\}}^{s+Q} P(E(D_L) = x - j) \qquad j \leq s + Q \tag{3.13}$$

We transform this equation to a periodic review (R,s,nQ) model by replacing $E(D_L)$ with $E(D_{L+R})$. By doing so, we let go the assumption that the inventory position is uniformly distributed over [s+1;s+Q] as proposed by Axsäter. Next step is to incorporate a service level constraint, and thereby setting the reorder point equal to the first reorder point for which the expected volume fill rate exceeds the target volume fill rate, see Equation 3.14.

$$E(\beta) = \frac{\sum_{x=1}^{\infty} \sum_{j=1}^{\infty} min(j, x) \cdot f_x \cdot P(SoH = j)}{\sum_{x=1}^{\infty} x \cdot f_x} \tag{3.14}$$

In Equation 3.14, E(β) is the expected volume fill rate, which should exceed the target volume fill rate β^*, $min(j, x)$ being the delivered quantity, f_k the probability of order line size x, and $P(SoH = j)$ as denoted in Equation 3.13. As stock on hand level j in practice can never be larger than reorder point s + replenishment order quantity MOQ (although ignoring undershoot due to non-unit order line sizes, and in our case multiple batch of IOQ + 1$\cdot$ MOQ can occur), the summation in practice stops at $s + Q$, $P(SoH = j)$ already converges to zero before $j = s + Q$. Furthermore, we propose a transformation into a full back-ordering model, as follows: replace the partial delivery element $min(j, x)$ by a full delivery (and thereby full back-ordering/lost sales) if-then structure:
if $x > j$
then 0
else x
This leads to Equation 3.15:

$$\begin{cases} E(\beta)' = \dfrac{\sum_{x=1}^{\infty} \sum_{j=1}^{\infty} x \cdot f_x \cdot P(SoH = j)}{\sum_{x=1}^{\infty} x \cdot f_x} & \text{, for } x \leq j \\ E(\beta)'' = \dfrac{\sum_{x=1}^{\infty} \sum_{j=1}^{\infty} 0 \cdot f_x \cdot P(SoH = j)}{\sum_{x=1}^{\infty} x \cdot f_x} & \text{, for } x > j \end{cases} \tag{3.15}$$

We reformulate Equation 3.15 into Equation 3.16, and combine these in Equation 3.17. In addition, we reintroduce the upper-bound (replacing the infinity in the summation) for inventory level j of s+Q, as earlier mentioned.

$$\begin{cases} E(\beta)' = \sum_{j=1}^{s+Q} P(SoH = j) & \text{, for } x \leq j \\ E(\beta)'' = 0 & \text{, for } x > j \end{cases} \tag{3.16}$$

$$E(\beta) = E(\beta)' + E(\beta)'' \tag{3.17}$$

Including undershoot

As mentioned in Section 2.1.1, in the standard methods the undershoot is ignored, as it is in for example the compound Poisson method of Section 3.4.1. However, we could incorporate this in our model. This undershoot effect comprises of two components caused by two aspects: the order line size distribution and the time until the next review moment, the latter obviously only occurs in periodic review settings, which we work with (Hill, 1988). Hence, for this method we assume the (R,s,Q) model. For the expected values of undershoot of a SKU, we used the formula from Hill (1988). We distinguish between undershoot due to the order line size distribution $Under_X$ (Eqs. 3.18 and 3.19), and due to the time until the next review moment $Under_W$ (Eqs. 3.20 and 3.21). We introduce random variable X for the order line size, and N for the number of customers during a review interval respectively.

$$E(Under_X) = \frac{E(X^2)}{2 \cdot E(X)} - \frac{1}{2} \tag{3.18}$$

$$var(Under_X) = \frac{E(X^3)}{3 \cdot E(X)} - \frac{(E(X^2))^2}{(2 \cdot E(X))^2} - \frac{1}{12} \tag{3.19}$$

$$E(Under_W) = E(X) \cdot \frac{(E(N^2) - E(N))}{2 \cdot E(N)} \tag{3.20}$$

$$var(Under_W) = E(X^2) \cdot \frac{(E(N^2) - E(N))}{2 \cdot E(N)} + \frac{(E(X))^2}{12 \cdot E(N)} \cdot \left[4 \cdot E(N^3) - 6 \cdot E(N^2) - \frac{3 \cdot (E(N^2))^2}{E(N)} + 5 \cdot E(N)\right] \tag{3.21}$$

When independence between both components is assumed, these two components are to be combined in equations 3.22 and 3.23. Although Hill gives no reasoning whether this assumption on independence, for our multi-modal order line size SKUs this is rather reasonable, as both dimensions are not connected in whatever way, i.e. the number of items a customer buys is not related to the time until the next inventory review moment. For these SKUs customer want a specific number of items, not one less or one more.

$$E(Under_{total}) = E(Under_X) + E(Under_W) \tag{3.22}$$

and

$$var(Under_{total}) = var(Under_X) + var(Under_W) \tag{3.23}$$

In accordance with Hill (1988) we add the expected undershoot to the standard safety stock calculation. Next, we incorporate the variance of the undershoot components in the variance of the demand during the cover period. Furthermore, it includes the order line size distribution, through its first three moments, making it likely to provide different results for multi-modal order line size SKUs and unimodal order line size SKUs. Formally, we define for this “Add undershoot”-method the safety stock and the altered standard deviation of demand during the cover period - denoted by an accent - in equations 3.24 and 3.25.

$$\text{Safety stock} = \text{standard method} + E(Under_{total}) \tag{3.24}$$

and

$$\sigma'_{(D_{L+R})} = \sqrt{\sigma^2_{(D_{L+R})} + var(Under_{total})} \tag{3.25}$$

Power approximations Nenes et al.

Finally, we address a method of determining the safety stock through finding its base stock level. This method resembles the power approximation method proposed in 1979 (Ehrhardt, 1979). This heuristic provides predetermined power approximations for setting the base stock level S, given a review period and lead time, only needing the demand during the cover period (Nenes, Panagiotidou, & Tagaras, 2010), inferences on variance follow from the corresponding distribution and are thereby included in the power approximation. Hence, this method assumes a (R,s,S) method, which we transform into a (R,s,Q) model. Nenes' power approximations assume demand during cover period either to be Gamma distributed or approximated through a ‘package Poisson’ distribution - we discuss the latter in Section 3.2.3, as follows for R=L=1:
$safety\ stock = [3.403 \cdot E(D_{L+R}) - MOQ] \cdot E(D_{L+R})$ for $\beta^* = 0.80$
$safety\ stock = [4.275 \cdot E(D_{L+R}) - MOQ] \cdot E(D_{L+R})$ for $\beta^* = 0.90$
$safety\ stock = [5.509 \cdot E(D_{L+R}) - MOQ] \cdot E(D_{L+R})$ for $\beta^* = 0.95$
$safety\ stock = [7.670 \cdot E(D_{L+R}) - MOQ] \cdot E(D_{L+R})$ for $\beta^* = 0.99$

Hence, base stock levels are transformed into reorder points by subtracting the minimal replenishment order quantity. Next, we subtract the expect demand during the cover period from the newly derived reorder point in order to find the safety stock, see Equation 3.26.

3.4.2 Reorder point calculation

To the best of our knowledge, there is one well-known method for setting the reorder point to be found in literature, which entails a general, straight-forward combination of the safety stock, and the expected demand during the cover period. This reorder point determination is widely described in literature, such as Silver, Pyke, Thomas (2017), we provide Equation 3.26.

$$\text{Reorder point} = E(D_{L+R}) + \text{safety stock} \tag{3.26}$$

Safety stock is calculated using the cycle service level, target fill rate, the method of Tijms and Groenevelt, or any other method. Hence, this is a distribution-independent method, as long as a closed-form expression of the expected value of the distribution for demand during the cover period exists.

3.5 Conclusions

We finalize this chapter by first answering the sub questions, and we conclude with our research question for our literature study. Implications for next chapters are given.

a. *How can we identify multi-modal order line size SKUs the best way?*
To the best of our knowledge there is hardly any research done on identifying modes within order line size distribution, except for mixture models. We prefer the empirical distribution (see Section 3.2.2), over a mixture of theoretical distributions, and for mixing the empirical distribution would make no sense. Hence, we need to develop our own method for classifying SKUs on their order line size distribution.

b. *How can we incorporate the 'multi-modal order line size effect' into inventory management?*
The most important heuristics from literature for incorporating the 'multi-modal order line size effect' from SKUs classified as such, focus on incorporating the order line size distribution into safety stock and reorder points. Safety stock calculations are divided into either a direct service level method, an approximation by Tijms & Groenevelt, Poisson-based method, including undershoot, and a base stock level approximative method. The reorder point calculation methods involves a distribution-free method.

Finally, it is time to answer the research question:
What can we learn from literature about optimization of multi-modal order line size SKUs inventory?
Sales can be transformed into true demand by enriching data with theoretical distributions. However, this is not preferable, as we do not have access to daily inventory level data, nor is it valid in our setting to only accommodate for larger order line sizes and not on other unobserved lost sales. So, we assume demand to equal sales. The demand process can, as mentioned in the last paragraph of Section 3.2.1, at best be modelled via a compound Poisson process, therein focussing on the order line size distribution, which can be some mixture of zero-truncated, one-inflated, universal negative binomial distributions, or a single empirical distribution. Several heuristics exist for computing safety stocks.

The next step is to compare the standard methods (see Section 2.1) with literature-based methods (see Section 3.4), and own-derived heuristics that do include the empirical order line size distribution. We introduce these heuristics in Section 4.5. The literature-based methods mainly assume demand during the cover period to be normally, Gamma or compound-Poisson distributed. Our methods on the contrary further employ the compound Poisson process, with an empirical order line size distribution as compounding function for the compound Poisson distribution.

Inventory model

Based on the model description as provided in Section 2.1, we develop our inventory model in the following way. At first we provide our assumptions, and we describe our notation. Next, we determine the parameters and describe our different heuristics, and the verification and validation steps are discussed. With a small dataset we determine the best performing methods. Finally, in Chapter 5 we test our inventory model for a dataset from the company as described in Section 2.2. To this extent, we split transaction data into an initialization period for setting inventory parameters, and a testing stage in which performance of the methods is measured. We also develop a heuristic for the reorder point calculation, but it shows bad results, more details can be found in Appendix D. In this chapter we answer the corresponding research questions:

3. How can the inventory management at best incorporate multi-modal order line demand?

   a. *What are the different possibilities, and which alternative performs best?*

   b. *How can the inventory model be validated and verified?*

   c. *Which SKUs are classified and how is this done?*

4.1 Model assumptions

We start constructing our inventory model by setting the following model assumptions:

- **Single-item**: each item is handled on its own. Hence, there are no relations and dependencies between items, and we continue with this *single-item* characteristic. Products' demand distributions are assumed to be mutually independent.

- **Service level constraints**: within scientific papers there are two objectives possible for optimization, either cost-based, or with service-level constraints. We focus on the latter, as this is mostly used in this retail environment. To be more precise, our models strive to achieve a target volume fill rate β^*, with a as few average stock on hand as possible.

- **Full lost sales, conditional back-ordering, full back-ordering**: in case there is not enough inventory to fulfil a customer's complete order line, our model is capable of running three scenarios: the customer leaves the shop resulting in a lost sale, the customer places a back-order and receives his/her SKU when the SKU is replenished, or an intermediate option.
  In this intermediate option, the so called *conditional back-ordering*, the customer places this back-order if and only if he/she orders at least a threshold number of items or - in case of multi-modal order line size SKUs - 1 item of this SKU (as necessary replacement for example). In descending order of being realistic within our market of research, our own-developed conditional back-ordering setting (i.e. the in-between form) comes first, second is the lost sales setting, and least realistic is full back-ordering. For example, one does not return to a retailer to pick up the back-order for a homogeneous, one-at-a-time SKU like a bottle of soap. Instead, one simply goes to another retailer nearby. However, if he/she needs an extraordinary amount, say 10, bottles of soap, we hypothesize that he/she is likely to wait a small amount of time for the back-order to arrive (the rest of the review time + lead time).
  Hence, we set a back-order threshold for order line sizes that are bigger than *BO*-factor times the

average order line size. *BO*-factor equals a constant related to the extent to which customers are eager to back-order their demand because they are aware of the fact that they order a large amount of items. For multi-modal order line size SKUs we also account for replacement orders of order line size '1', by back-ordering order lines from these sizes. This is logical, for example as customers wanting their specific mug (of which they already have 3 non-broken mugs at home) being back-ordered, because they want this mug specific. It is a management decision whether to use lost sales or conditional back-ordering, where possible we used the right formula for either back-ordering or lost sales. We have developed the conditional back-ordering setting ourselves. Furthermore, it uses the lost sales formulae, as most of the orders (except for the ones with order line sizes more than a certain factor times the average order line size, or an order line size of '1') are lost in case of insufficient stock. Note: in the back-ordering case, we included a successfully delivered back-order in the observed fill rates. We expect that this increases the back-ordering's case performances only slightly, compared to the lost sales case.

- **Periodic review**: as most of Slimstocks clients make use of a periodic review model, especially in the environment of non-food retailers, it is logical to work with periodic review models.
- **Constant lead times**: retailers are stocked by their distribution centre in a tight delivering scheme, resulting in near-constant lead times. Hence, we model with deterministic, constant lead times.
- **Seasonality and promotion**: SKUs like garden chairs most probably have a strong seasonality effect, as they are hardly sold during the winter, and there are big peaks in sales during spring and summer. This effect should be accounted for. We start by leaving these - seasonality-affected - SKU-shop combinations out of our analysis. We hypothesize, based on the specific attribute of our 'issue SKUs' being most often demanded in certain order line sizes, that only the arrival rate of customers might change as a result of seasons. That is, more people still buy four garden chairs, and not five, and therefore we assume customer arrival rate and the order line size distribution to be independent. The same goes for promotions, which can be seen as ad hoc seasons.
- **Pre-emptive back-orders**: back-orders are ordered at the supplier at the beginning of the next review period, and replenished after exactly one lead time. Right after their delivery, the back-ordered SKUs are picked up by the back-ordering customers, before the regular demand of that day is fulfilled. Back-orders always trigger a replenishment order.
- **Stationary demand**: we focus on *stationary* demand per day and per week, since non-stationary demand results in all kinds of exceptions and ad hoc decisions to be implemented. This is due to the lack of business rules and exception-based control like there are in more complex systems.
- **No insurance inventory**: in the standard situation clients set an *insurance inventory* in order to ensure to have a safety stock of some manually set level. This setting, entailing three configuration influencing the safety stock (either having a minimum level of the insurance inventory and safety stock or overwriting the calculated safety stock) or the reorder point (by overwriting this reorder point), is not taken into account in our model, as we want to get rid of this way of dealing with multi-modal demand SKUs.

4.2 Notation

In this section we introduce the notation, parameters, and variables which we use in our model.

4.2.1 Indices and input

We introduce the following indices:

- c for the c^{th} customer $\in$ C of a specific SKU i at a shop j on day d of review period t;
- d for the d^{th} day of review period t, d: 1, 2, 3, 4, 5, 6, or 7, corresponding with the weekdays starting on Monday;
- i for the SKU i $\in$ I;
- j for the shop j $\in$ J;
- k for the order line size k $\in$ K;
- t = mR for the m^{th} review period with m $\in$ M and t^{-} be the last review period of the initialization period;

Next, we gather input as in Table 4.1, based on Section 2.1 and Section 2.2.

Parameter	**Symbol**	**Source**
(Replenishment) Minimum Order Quantity	$MOQ_{i,j}$	Product information SKUs
(Replenishment) Incremental Order Quantity	$IOQ_{i,j}$	Product information SKUs
Target volume fill rate β^*	$\beta^*_{i,j}$	Product information SKUs
BO-factor	$BO\text{-factor}_i$	Product information SKUs
Sold order line	$SOL_{c,d,t,i,j}$	Transaction data

Table 4.1: *Input and notation*

Henceforth, we need two types of data: transaction data, containing at least the transaction data, SKU ID, and order line size; and information on the SKUs, such as the associated target service level, and the minimal and incremental order quantity (towards suppliers). We return to the gathering of data in Section E.

4.2.2 Parameters

The upcoming category consists of the inventory parameters which are pre-calculated once the aforementioned data is loaded, see Table 4.2. We restrict the general terms *mean of demand during a period* and *standard deviation of demand during a period* to weekly, for programming purposes. Furthermore, for the weekly demand standard deviation, we can employ the population standard deviation when we aggregate per SKU on all shops, as we have full information. Formally, for $|C| \leq 1$ we average on a all J shops selling SKU i. For the specific SKU-shop combination in which more than one order took place in the initialization period, we are restricted to using the sample standard deviation, since we take a sample (both time-wise and shop-wise). We hypothesize that the multi-modal order line size SKUs only show sales of quantities that customers would have wanted. I.e.: when a customer's true demand is four chairs, he/she is most unlikely to buy three chairs when there three chairs on stock. He/she either leaves the shop (lost sales) or places a back-order. As these lost sales are still unobservable, we have censored sales data in that sense, but we take the gap between sales and true demand as a result of unobservable lost sales for granted. Furthermore, we did not find methods in literature to transform this censored sales into true demand that would be easily implemented for our specific demand process. Thus, we choose to use sales directly as input for our demand. In Table 4.2, let back-order (BO) threshold be the average order line size of this SKU-shop combination multiplied by the general BO factor which is the same for all products, and strictly bigger than 1. In the conditional back-ordering setting, this order line size and all order line sizes above this one, are not taken into account when determining the weekly demand, since we want to back-order these order line sizes anyway.

Parameter	**Symbol**	**Formula**
Mean of weekly demand	$E(D)_{i,j}$	$\frac{1}{t^-} \cdot \sum_{t=1}^{t^-} \sum_{d=1}^{7} \sum_{c=1}^{\infty} SOL_{c,d,i,j,t} \quad , SOL_{c,d,i,j,t} < BO_{threshold,i}$
Aggregated mean of weekly demand	$E(D)_i$	$\frac{1}{J} \cdot \sum_{j=1}^{J} \left[\frac{1}{t^-} \cdot \sum_{t=1}^{t^-} \sum_{d=1}^{7} \sum_{c=1}^{\infty} SOL_{c,d,i,j,t} \right]$
Standard deviation of weekly demand	$\sigma_{D_{i,j}}$	$stdev.s(\sum_{d=1}^{7} SOL_{c,d,i,j,t})$
Aggregated standard deviation of weekly demand	σ_{D_i}	$\frac{1}{J} \cdot \sum_{j=1}^{J} \left[stdev.p\left(\sum_{d=1}^{7} SOL_{c,d,i,j,t} \right) \right]$
Mean of demand during cover period	$E(D_{L+R})_{i,j}$	$E(D_{i,j}) \cdot$ (L+R)
Standard deviation of demand during cover period	$\sigma_{(D_{L+R})_{i,j}}$	$\sigma_{D_{i,j}} \cdot \sqrt{L+R}$
Replenishment lot size	$Q_{r_{i,j}}$	$\max[R * E(D_{i,j}), MOQ_{i,j}]$ (van Donselaar & Broekmeulen, 2013)
Aggregated in-service order line size	$E(X)_i$	$\frac{1}{J} \cdot \sum_{j=1}^{J} \beta^*_{i,j} \cdot \left[\sum_{k=1}^{\infty} k \cdot p(SOL_{c,d,i,j,t} = k) \right]$ for t:1,...,$t^-, \forall d, c$
BO-threshold	BO_i	BO-factor$_i \cdot \left[\sum_{k=1}^{\infty} k \cdot p(SOL_{c,d,i,j,t} = k) \right]$ for t:1,...,$t^-, \forall d, c$

Table 4.2: *Inventory parameters and their notation*

In Table 4.2, let $x \in \mathbb{N}$ be an arbitrary order line size. Note: the BO-threshold does not depend on j, since we define it on SKU-level.

4.2.3 Variables

The next category consists of the inventory variables, calculated per week and day as denoted in Table 4.3. These variables generate the replenishment order advices. Next, as we assume in Section 4.1, sales are completely lost or back-ordered when there is not sufficient inventory available to fulfil the entire order line. Recall that back-orders have priority over regular orders (see Section 4.1 as well). Both two aforementioned aspects result in the calculation for stock on hand, which is calculated each day after each customer that is either served or not (resulting in either a back-order or lost sales). Since back-orders always trigger (larger) replenishment order if the inventory position is below the reorder point before the reviewing moment, and the replenishment arrives first thing in the morning, the stock on hand can never be below zero. On order is equal to the sum of the not yet delivered replenishment orders, and thereby equal to the amount that is to be delivered at the next moment of delivery. In Chapter 5 we elaborate on the assumption to set the review period equal to the lead time, both lasting one week, which is handy for example for these expressions to develop. As we deal with a rather typical version of a (R,s,nQ) model, in which n batches of Q (the incremental order quantity) are ordered if and only if at least the minimal order quantity is ordered at the supplier as well. For this reason, we introduce an auxiliary variable to determine whether or not we should order the minimal order quantity or nothing at all: *Ordering MOQ*. In case the inventory position is at or above the reorder point, no order is placed, so this variable remains zero. Next, we calculate the number of batches of IOQ we need, if we order any at all (hence, determining the maximum of 0 and n). Finally, we generate the replenishment order quantity, composed of the minimum order quantity and n times the incremental order quantity. Let *Backorders* be the set of orders that were back-ordered, and $\lceil x \rceil$ ensures rounding up to next integer, with $x \in \mathbb{R}$.

Variable	**Symbol** *for each tuple (i,j,t)*	**Formula**
Stock on Hand	SoH_d	$SoH_{d-1,i,j,t} - backorders_{i,j,t-L} + IOO_{i,j,t} - min[E(D_{i,j,t}), SoH_{d-1,i,j,t}]$, for $t > 0$
Ordering MOQ	$OMOQ$	$\begin{cases} 0 & , IP_{7,i,j,t-L} \geq s_{i,j}, \forall t \in T \\ MOQ_{i,j} & , IP_{7,i,j,t-L} < s_{i,j}, \forall t \in T \end{cases}$
Number of IOQs	n	$\begin{cases} 0 & , IP_{7,i,j,t-L} \geq s_{i,j}, \forall t \in T \\ max\left(0, \left\lceil \frac{(s_{i,j} - IP_{7,i,j,t-L}) - MOQ_{i,j}}{IOQ_{i,j}} \right\rceil \right) & , IP_{7,i,j,t-L} < s_{i,j}, \forall t \in T \end{cases}$
Replenishment order quantity	OQ	$OMOQ_{i,j,t} + n_{i,j,t} IOQ_{i,j}$
On order	IOO	$\sum_{t=t-L}^{t} OQ_{i,j,t}$, for $t > 0$
Inventory position	IP_d	$SoH_{d,i,j,t}$ + $IOO_{i,j,t}$ - $backorders_{i,j,t-L}$

Table 4.3: *Inventory variables and their notation*

All-in-all, the sequence of events is as follows:

1. A replenishment order for next period is generated;
2. A replenishment order is delivered;
3. Back-orders - if any - are picked-up by the back-ordering customers;
4. Regular customers of this day arrive and buy SKUs;
5. The process continues at [1.].

4.3 Demand distribution

As described in Section 2.1.1, Slim4 already classifies SKUs based on their demand during a certain period of time. Within our own heuristics, we distinguish between three different probability distributions. First of all we can model demand during the cover period with the normal distribution, which is often suitable for fast-moving SKUs (Silver et al., 2017). Next, we introduce the compound Poisson distribution, which suits slow-moving SKUs (demand of less than 10 items - so not customers - during the cover period), and we explain why the empirical distribution represents the order line size distribution. Finally, we develop heuristics that directly employ the first moment of the empirical probability distribution for order line size of each SKU-shop combination, since the standard methods (see Section 2.1) nor the methods from literature (see Section 3.4, only the Poisson methods indirectly use the order line size distribution) can directly include an empirical order line size distribution.

4.4 Classification SKUs

In order to develop insights on performance differences, we classify SKUs from our dataset on *multi-modal order line size* vs. *unimodal order line size*. We have refined our method from Section 1.3, in order to generate a more accurate, intuitive distinction. We found our method by comparing difference between three artificial, very different example SKUs, let's call them SKU X, SKU Y, and SKU Z, with order line size probabilities as denoted in Table 4.4.

Order line size x	1	2	3	4	5	6	7	8	9	10
$p(X = x)$	0.4	0.3	0.2	0.075	0	0.025	0	0	0	0
$p(Y = x)$	0.05	0.3	0.05	0.2	0.1	0.2	0	0	0	0.1
$p(Z = x)$	0	0.5	0	0.4	0.1	0	0	0	0	0

Table 4.4: *Three exemplary SKUs*

Obviously, SKU X has an approximate monotonically decreasing probability distribution, whereas SKU Y's probability distribution shows several modes on $p(Y = 2), p(Y = 4), p(Y = 5)$, and $p(Y = 10)$. Thus, the fact that the order line sizes are almost completely well ordered in a descending order of probability for SKU X, compared to the SKU Y's and SKU Z's order line sizes being not ordered at all, is interesting. If we order SKU Y and Z in an descending on the probability, we obtain Tables 4.5 and 4.6.

Order line size x	2	4	6	5	10	1	3	7	8	9
$p(Y = x)$	0.3	0.2	0.2	0.1	0.1	0.5	0.5	0	0	0

Table 4.5: *SKU Y's order line sizes ordered in descending order on probability*

Now we see that the sequence of order line sizes for SKU Y are not succeeding at all. Only two of the nine transitions (7 to 8, and 8 to 9) are successors and these all have 0 probability, and one of the nine transitions (6 to 5) is preceding. For SKU Z, 5 out of 9 transitions are succeeding and correspond to $p(Z = x) = 0$ (in total 7 order line sizes had zero probability). Since probabilities of zero are clearly no modes, we should restrict ourselves to a certain amount of most occurring order line sizes. When looking at the data, we found that looking at the 5 most often occurring order line sizes satisfied. Hence, when looking at the 5 most-often occurring order line sizes - so, four transitions - we found that for SKU Y three out of four transitions were not succeeding or preceding each other, a 75% score. For SKU Z this percentage amounts 75% as well. For SKU X, the correct top-5 order should be 1, 2, 3, 4 , 6, and thus the

Order line size x	2	4	5	1	3	6	7	8	9	10
$p(Y = x)$	0.5	0.4	0.1	0	0	0	0	0	0	0

Table 4.6: *SKU Z's order line sizes ordered in descending order on probability*

percentage of not-succeeding order line sizes is 25%. Hence, for these artificial SKUs we found that SKUs with a percentage of at least 75% are multi-modal. This small example sounds intuitive and effective, and when testing it with larger datasets it still provided accurate results, and we found a percentage (δ^*) of 50% to resemble classifications from our time-consuming expert panel approach from appendix A, which are considered accurate. So, formally, we define the following *classification scheme*, for each SKU:

0. Initialization: set threshold δ^*;
1. Put the order line sizes in descending order on their empirical probability $P(X = x)$. If two order line sizes have equal probabilities, the lowest order line size is ranked first;
2. Calculate the absolute differences $\delta_{x,x+1}$ between the ranks of the top-5 order line sizes. Henceforth, monotonically decreasing distributions generate mainly 1s;
3. Calculate the $\delta_{\%}$ percentage of differences $\delta_{x,x+1}$ being bigger than 1, for $x : 1, 2, 3, 4$;
4. If $\delta_{\%}$ is greater than or equal to threshold δ^*, this SKU is classified as multi-modal.

We illustrate our method by classifying the SKU from an arbitrary SKU, as follows in Table 4.7.

Ranked order line size x	2	1	4	6	3	5	8	9	24	7
Frequency	53	37	14	11	4	3	3	1	1	0
$\delta_{x,x+1}$	1	3	2	3	n/a					

Table 4.7: *Own-derived classification example*

Table 4.7 shows that three out of four $\delta_{x,x+1}$ are bigger than 1, leading to a $\delta_{\%} = 3/4 = 75\%$. This exceeds the threshold, i.e.: $\delta_{\%} \geq \delta^*$, hence we classify this product as multi-modal.

4.5 Heuristics for safety stock

We now discuss three different heuristics for determining the safety stock, incorporating the order line size distribution. We have derived these heuristics ourselves.The first two heuristics are based on the insight that Slimstock currently does not explicitly take the undershoot into account in the configuration of our research, the (R,s,nQ) model. Slimstock argues that this is not necessary since enough batches of IOQ (and the initial minimum order quantity) are ordered to exceed the order level, whereas undershoot generally is a problem if only an amount of Q can be ordered (Slim4-developer, 2016a). Yet, it is still very well possible that the stock on hand level drops below the reorder point with an arbitrary amount, which especially is the case with multi-modal order line size SKUs. This undershoot could result in the problematic 'useless inventory' we aim to abolish. Hence, if we incorporate this undershoot, we would expect better results for multi-modal order line size SKUs. Next to the literature-based method from Section 3.4.1, we introduce two more undershoot-related heuristics, and one target fill rate method.

4.5.1 "Overwrite"-method

In this method we overwrite the safety stock level with the expected undershoot. In the standard (R,s,nQ) model undershoot is not included at all, so we suspect safety stocks to be overestimated when directly adding the expected undershoot, as in the "Add undershoot" method. Therefore we expect this method to perform reasonable, although we are not sure whether this performs well for unimodal order line size SKUs. Since this type of SKUs in general is fast-moving, the safety stocks are underestimated

4.5.2 "Max"-method

The second heuristic takes the maximum of the safety stock as calculated by the standard methods of the (R,s,nQ) model - thus without incorporating the undershoot - and the expected undershoot. We do this in order to construct a hybrid between the standard method, which we presume to work fine for unimodal order line size SKUs, and the undershoot method which we presume to work fine for multi-modal order line size SKUs. In order to prevent the earlier mentioned underestimation of the safety stock, we take the maximum between both. Formally, we define Equation 4.1.

$$\text{Safety stock} = max[\text{standard method}, E(Under_{total})] \tag{4.1}$$

4.5.3 "First moment β^*"-method

The final heuristic performs a multiplication between the average order line size and the volume fill rate, resembling the in-service order line size. Thus, we obtain Equation 4.2. This "First moment β^*"-method has the advantage that it works with target volume fill rate constraints, and as such it generates lower order advices for SKUs with lower target fill rates. Moreover, in our heuristic we are fine with not serving 1-β^* of the demand, by constructing a strictly lower 'expected' order line size than in a regular average order line size. Finally, in order to be able to serve the expected number of customers during the cover period, we multiply with λ_{L+R}. We differentiate between instances in which we aggregate (by taking the average) on shop-level in case of 1 or less orders within the initialization period.

$$\text{Safety stock} = \lambda_{L+R} \cdot \beta^* \cdot \left[\sum_{k=1}^{\infty} k \cdot p(SOL = k\right] \tag{4.2}$$

The "First moment β^*" for the aggregated example in Section 4.3 would come down to:

$$\text{Safety stock} = 1.10 \cdot 96.86\% \cdot \left[\sum_{k=1}^{\infty} k \cdot p(SOL = k\right] \tag{4.3}$$

Thus, $1.10 \cdot 96.86\% \cdot [1 \cdot 29.13\% + 2 \cdot 41.73\% + 3 \cdot 3.15\% + 4 \cdot 11.02\% + 5 \cdot 2.36\% + 6 \cdot 8.66\% + 7 \cdot 0.00\% + 8 \cdot 2.36\% + 10 \cdot 0.79\% + 24 \cdot 0.79\% = 2.76] = 1.10 \cdot 0.9686 \cdot 2.76 = 2.94$.

4.6 Performance measurement

We need to measure performance of our inventory model. In accordance with the 'conditions for implementations' as stated in Section 2.4, we do not incorporate measures like computation time or programming effort. We focus on quality-related measures, and we exclude the results from the initialization period, as these data were used for setting the inventory parameters. We list our main KPIs in Table 4.8.

Description	**Symbol** for each tuple (i,j)	**Formula**
Observed stock on hand	$\overline{SoH}$	$\frac{1}{M-t^-}\sum_{t=t^-+1}^{M} SoH_{i,j,t}$
Observed volume fill rate	$\overline{\beta}$	$1-\frac{\text{Unmet demand}}{\text{Total demand}}$
Observed order line fill rate	$\overline{\gamma}$	$1-\frac{\text{Unmet customers}}{\text{Total \# customers}}$

Table 4.8: *Key Performance Indicators and their notation*

Recall that M equals the maximum number of review periods as stated in Section 4.2.1. As we evaluate a service-level inventory model, the observed volume fill rate should be greater than or equal to the target fill rate. If this condition is met, we strive to find the lowest average stock on hand, as this is a main objective for inventory control in general. Next, if there is no difference on the first two KPIs we employ the observed order line fill rate, a final important KPI for retailers.

4.7 Validation and verification

As our model is an abstraction of reality, it of utmost importance to validate the model with reality and to verify that the model indeed functions as described, and programming is done correctly. We exploit the book of Averill Law (2015) as starting point for this matter. *Validation* continuously took place, for frequent brief meetings took place with developers and the consultant for company X, all of whom can be regarded to as *subject-matter experts* due to long experience in both inventory modelling, both from a practical and theoretical point of view. Furthermore, mainly in the early stages of modelling, our inventory model was constantly checked with a compact dataset of Slim4 that contained no exceptions, trends, seasonality or other complex details.

Furthermore about validation, the author's own experience in university courses and intuition was exploited. Next to this *collection of high quality information* (Law, 2015), we kept a written assumption document, and several *structured walkthroughs* with the supervisor and other colleagues at Slimstock took place.

Regarding verification of our model, we have started compact with a snapshot-like inventory model, and kept adding features. Effort was put in testing whether or not earlier created features were harmed by new ones. In addition, the model was tested for all kinds of settings, with for example reorder points set hard at zero, or incredibly high initial inventories. This was a test for robustness, and adaptations were made to make our model capable of handling a large parameter space.

Next, we introduce some kind of *warm-up period.* The sales data from the initialization period were used for setting the inventory parameters, which were used for calculating the inventory variables for the remaining months. Performance was only measured for these remaining months. Furthermore, we only used SKUs that showed a stationary weekly demand pattern. Whether or not the weeks within the initialization period were suitable (i.e. was likely to be from the same population or not), was tested through a independent samples t-Test on the weekly demand of the first eight weeks versus the weeks used for performance measurement. Significance level α was as usual set on 5%, and of we used the methods for unequal sample sizes (8 weeks versus 18 weeks) and unequal variances (we do not know the variance beforehand, so we need to incorporate the possibility they differ). It turned out that of our initial dataset, containing all kinds of both types of SKUs, we could erase 15.65% of our transactions from 5.82% of our SKU-shop combinations, due to being non-stationary.

4.8 Conclusions

We finalize this chapter by first answering the sub questions, and we conclude with our research question for our inventory model. Implications for next the chapter are given.

a. *What are the different possibilities, and which alternative performs best?*
Next to the standard and literature-based methods, we can incorporate the empirical order line size distribution by regarding to the demand process as compound Poisson process, via own-derived equations. Three heuristics for the safety stock are discussed, and we include a different reorder point heuristic in Appendix D that works with the customer arrival rate and the order line size distribution. Demand during the cover period can be modelled as normal, or compound Poisson, implying to use methods from literature. We cannot determine which method performs best a priori. We provided the Key Performance Indicators through which we can differentiate.

b. *How can the inventory model be validated and verified?*
Several ways for both validating and verifying our model were used. Most methods come down to either discussing with my external supervisor, my own thoughts, comparing with numerical results from Slim4 itself, or by explicitly checking on assumptions we have made. Most took place in a natural way, as a result of several years of practising modelling during university courses.

c. *Which SKUs are classified and how is this done?*
We have derived our own method of classifying, since using an expert panel as suggested in appendix A would be too time-expensive and potentially inconsistent. Furthermore, in literature there were no classification methods found that focussed on the order line size distribution, see Section 3.3. We based our classification method on (and validated with) earlier methods with an expert panel. This leads to the classification of mainly SKU groups like cutlery and dining room equipment as they show a multi-modal order line size distribution.

We conclude with our answer to this chapter's research question:
How can the inventory management at best incorporate multi-modal order line demand?
We have obtained three heuristics: "Overwrite"-, "Max"-, and the "First moment β^*"-method, for directly incorporating the empirical order line size distribution, alongside for example a compound Poisson distribution from literature (see Section 3.4.1), that indirectly comprises the empirical order line size distribution. Next, we measure performance mainly in terms of volume fill rate and average stock on hand, followed by the observed order line fill rate. In order to determine which of the aforementioned methods works best, in Chapter 5 we should construct a experiment with historical sales data.

Experiment

Now that we have developed an inventory model, we can construct our experiment. We start by gathering the right data, thereby selecting the right SKUs for testing. Next, we provide the experimental design, addressing all relevant experimental parameters. Finally, we discuss some final aspects and set-up a preliminary experiment in order to filter out some bad-performing heuristics. In Section 5.7 we analyse the experiment's results. Through this experiment, we want to find the best-performing methods and heuristics for generating ordering policies. In this chapter we answer the research questions:

4. Which heuristic performs best on both the products classified and not classified as 'multi-modal order line size SKUs'?

    a. *Which data are needed and how do we gather them?*

    b. *Which of the methods and heuristics for generating ordering policies performs best?*

First, we address some aspects related to our experiment like the variables at hand, data exclusion, the SKUs we select, and our programming efforts. Finally, we select the most promising methods through a small experiment we run before our final experiment.

5.1 Selection and classification of SKUs

For each SKU one has the choice whether to hold inventory or not, many reasons exist for either choice. Non-stocked items are ordered at the supplier only in case demand occurs (analogue to Make-to-Order environments), and no inventory is held. As this is of no interest for our inventory model, we restrict ourselves to stocked items (analogue to Make-to-Stock). Second, we needed stationary demand, so we performed t-Tests on all SKU-shop combinations, as described in Section 4.7. And third, we exclude SKU-shop combinations with more than 15 customers per day, as with such a high customer arrival rate, it is expected to have no use to look at order line size distributions. Furthermore, this last limitation provided much programming convenience. For an elaboration on the data gathering, we refer to Appendix E.

Our specific SKUs coming from company X are selected based on their likeliness to have a potential for improvement. This particular attribute, of having a multi-modal order line size distribution, led to the following conclusions: (a) the order line size distribution is modelled with its empirical distribution, as it showed very differing and hardly-tractable patterns, and (b) the entire ordering process can be regarded to as compound Poisson process. Finally, we decided to also model the arrival of customers empirically, by loading the sales data in our model, as if it were future demand. See Table C1 in appendix C for the selected SKUs and some SKU information. Finally, in general the cumulative probabilities of order line sizes of up to 1 or 2, are low for multi-modal order line size SKUs, and higher for unimodal order line size SKUs.

5.2 Experimental design

In order to find the best-performing method for dealing with multi-modal order line size SKUs and with unimodal order line size SKUs, we conduct an experiment, in which we construct three variables and differing number of levels per variable. To be precise, each scenario comprises one level of each of the

[I.] Out-of-stock	**[II.] Safety stock**	**[III.] Reorder point**
A. Full back-ordering sect.4.1	A. Standard (β^*) Eq. 2.4, sect. 2.1	Standard Eq.2.5, sect. 2.1
B. Conditional back-ordering sect.4.1	B. *Cycle Service Level* Eq. 3.3, sect. 3.4.1	
C. Lost sales sect.4.1	C. *Tijms approximation* Eqs.3.8 & 3.9, sect.3.4.1	
	D. *Compound Poisson* β^* Eq.3.14, sect.3.4.1	
	E. *Power Approximations* Sect.3.4.1	
	F. *"Add undershoot"-method* Eqs. 3.22 & 3.23 Sect.3.4.1	
	G. *"Overwrite"-method* Sect. 4.5.1	
	H. *"Max"-method* Eq. 4.1, sect.4.5.2	
	I. *"First moment β^*"-method* Eq. 4.2, Sect.4.5.3	

Table 5.1: *All scenarios for running experiments*

three columns (variables) from Table 5.1. Note: setting III.B means the in-service order line size times the expected number of customers during the cover period. For example scenario 1 consists of I.A, II.A, and III, i.e.: Full back-ordering with standard (R,s,nQ) model like settings for calculation of the safety stock and the reorder point. Henceforth, we construct 3 (Variable I) · 9 (Variable II) · 1 (Scenario III.A) = 27 scenarios. Next, we need to initialize parameters. First one is the *BO* factor. We test with the value 3, as this value yielded highly realistic values of for example 12 mugs or 18 cake forks. We do not vary this factor, as it would double our experiment size.
For our model, we chose to have equal review periods and lead times, both lasting one week. Placing replenishment orders takes place on Monday morning before receiving new SKUs or selling anything, and replenishment arrives on the next Monday morning before opening the shop, and henceforth is available on Monday. This setting is chosen as it makes programming more convenient, and it is within reasonable scales for retail environments. In reality review periods and lead times are slightly shorter, most often each day or twice a week (Slim4-developer, 2016a).

The six months of planning horizon start with the 'initialization period', of which performance is not taken into account. The initial stock on hand, i.e.: at the beginning of our experiment, per SKU on shop-level, is set equal to the average stock: $\textit{safety stock} + \frac{Qr}{2}$, with Qr defined in Table 4.2, in accordance with Silver, Pyke, Thomas (2017). Next, the initial stock on order, the amount that could have been ordered at the supplier before our inventory model would even start, is initialized as zero.

5.3 Data exclusion

Finally, we excluded the transactions from Table 5.2, as they are regarded extreme outliers (all fall outside 95%-confidence intervals around the average order line size), which seriously disturb performance. These extreme big order lines would normally be back-ordered, complete separate from the 'regular' demand process.

SKU	**Order line size (# items)**
MM SKU 1	21, 30, 40, 100, 204
MM SKU 4	36
MM SKU 5	30
MM SKU 9	24, 30
MM SKU 13	50
MM SKU 14	24
MM SKU 19	24, 30, 40
UM SKU 23	25, 26, 30, 35, 40, 50
UM SKU 25	60
UM SKU 26	25, 29, 40
UM SKU 27	24, 29
UM SKU 28	24, 30, 39
UM SKU 29	24
UM SKU 32	24
UM SKU 33	25
UM SKU 34	24
UM SKU 36	26, 30

Table 5.2: *Transactions containing these order line sizes were excluded from the experiment*

5.4 Programming

We have written a *Visual Basic for Applications*(VBA) code in *Microsoft Excel*, included in appendix B, through which we were able to calculate performance on experiment levels as follows: SKU-shop combination, back-ordering versus conditional back-ordering versus lost sales, heuristics for safety stock. These results were analysed based on their score on the first three KPIs, namely the observed volume fill rate (as high as possible), average stock on hand (as low as possible), and the observed order line fill rate (as high as possible). Their relative difference compared to the base line of both safety stock and reorder point calculations from the standard approach, were calculated per out-of-stock setting, and averaged on SKU-level. Finally, in order to determine the robustness, we calculated standard deviations between performances on SKU-shop level.

5.5 Selecting most promising methods

In order to select the best methods for our final experiments, we now start by discussing the assumptions of Tijms & Groenevelt's approximation. All assumptions are evaluated per SKU-shop combination, and when one of the assumptions was not attained, the method was not performed. The first assumption that requires demand to be smooth, is covered as we focus on stationary SKUs. The second assumption - of setting a lower bound to the target fill rate of 90% - results in this heuristic not being calculated for SKU-shop combinations for which the target volume fill rate is below 90%, which occurred in 0.20% of the SKU-shop combinations. The third assumption, which prescribes that the review period should not be too small compared to the lead time - in general holds, as we set the review period equal to the lead time in our experiment. Finally, the fourth assumption, involving the coefficient of variation ($CV_{D_{L+R}}$), which should not exceed 0.5, depends per SKU-shop combination. In 51.39% of the SKU-shop combinations, the normality assumption did not hold. In Appendix F, we have made an attempt to transform the k determination of the Ju(k) part from Equation 3.10, into an approximation of the Gamma distribution, thereby avoiding the evaluation of incomplete gamma integrals (Tijms & Groenevelt, 1984). We did not succeed in this, hence, we abolished this method from further evaluation.

Next, we decided to abolish the literature method involving the cycle service level. This method is not capable of incorporating an order line size, hence an improvement is not expected, and it would not solve the problem of not incorporating order line sizes. On the contrary for example, the compound Poisson approach directly incorporates the order line size distribution.

Finally, we have performed a first test involving a limited amount of both multi-modal and unimodal order line size SKUs. This led to the conclusion that Nenes' approximations perform considerably worse than all other methods, so we do not include this method in our final experiment.

5.6 Differences between out-of-stock configurations

As we currently have three configurations for out-of-stock occurrences (full lost sales, conditional back-ordering, and full back-ordering), we first want to know the differences between the impacts on performance each out-of-stock configuration has. Data can be found in appendix C in Table C2. Moreover, in case the conditional back-ordering does not clearly differs from either lost sales or full back-ordering, it would only slow our experiment. We have tested the performance of the three out-of-stock configurations for the 38 SKUs of company X, for the standard method. To this extent we calculated the absolute differences between the three out-of-stock configurations for each of the three KPIs (stock on hand, volume fill rate, order line fill rate). In case one of the three KPIs differed among the two compared out-of-stock configurations, this SKU-shop combination was marked as *Difference between out-of-stock configuration 1 and 2*. It turned out that in 68.66% of the SKU-shop combinations, the full lost sales and full back-ordering setting did not differ (again, for standard methods, which do incorporate a different equation for each setting, see Section 2.1). As these equations (see Equation 3.4 and 2.2) resemble for large fill rates $\left[\frac{(1-\beta^*)}{\beta^*} \xrightarrow{\beta^* \to 1} (1-\beta^*)\right]$, and in our dataset target fill rates are generally high (on average 97.83%), it makes perfectly sense that both settings do not differ at all in many cases.

In respectively 81.84% and 71.35% of the SKU-shop combinations, there was no difference between full lost sales and conditional back-ordering, and full back-ordering and conditional back-ordering. As confirmation of the assumption made in Section 4.1, we expected conditional back-ordering to perform more like lost sales than like full back-ordering, since formulae of lost sales were applied for the widest range of order line sizes. This is confirmed by these results, conditional back-ordering yielded rather similar results. As a result of this similarity, we think it is not an interesting configuration for further research within our scope. Hence, in the remainder of our research, we do not include the conditional back-ordering configuration. Concluding, we run our experiment with the design as denoted in Table 5.3.

5.7 Results

After cutting the methods down, we have run the experiment using company X's dataset containing 38 SKUs, of which 22 are classified multi-modal order line size SKUs. We run the experiment for both lost sales and full back-ordering, for the standard methods, compound Poisson methods for target fill rate, and our heuristics, resulting in $2 \cdot 6 = 12$ scenarios per SKU-shop combination. For each SKU-shop combination,

[I.] Out-of-stock	**[II.] Safety stock**	**[III.] Reorder point**
A. Full back-ordering sect.4.1	A. Standard (β^*) Eq. 2.4, sect. 2.1	Standard Eq.2.5, sect. 2.1
B. Lost sales sect.4.1	B. *Compound Poisson* β^* Eq.3.14, sect.3.4.1	
	C. *"Add undershoot"-method* Eqs. 3.22 & 3.23 Sect.3.4.1	
	D. *"Overwrite"-method* Sect. 4.5.1	
	E. *"Max"-method* Eq. 4.1, sect.4.5.2	
	F. *"First moment* β^* *"-method* Eq. 4.2, Sect.4.5.3	

Table 5.3: *All scenarios for the final experiment*

we calculate each method's relative difference with the standard (target fill rate) approach. Stock on hand differences are calculated as relative difference compared to the baseline, whereas for observed fill rates absolute differences- so, %-points - are calculated, since relative difference between two percentages is not very intuitive. Next, these SKU-shop combinations are averaged on SKU-level, resulting in relative performance per heuristics per SKU, compared to their standard (R,s,nQ) model-baseline. Ideally, we look for heuristics with the highest increase on average observed volume fill rate and the smallest relative increase of average observed stock on hand. We start with the literature-based methods, followed by our own-derived heuristics, and conclude by making a comparison between the methods.

We discuss the results in this hierarchy: 1. Heuristics, 2. Type of SKU, 3. Out-of-stock configuration. Hence, we first discuss the first method's multi-modal SKUs' performance given lost sales, continue with its back-ordering case performance, and - after discussing all intermediate results - conclude with the last method's unimodal order line size SKUs' performance given back-ordering.

5.7.1 Means for comparing methods

First, we found that the order line fill rate hardly differed from the corresponding volume fill rate. Therefore, we focus on the volume fill rate, as this is a more frequently used measure. In accordance with the objective of Slimstock, we strive to reach the volume fill rate, against the smallest increase of stock as possible. The straightforward performance ratio of *relative difference in stock on hand* divided by *absolute difference in volume fill rate* provides a reasonable insight per heuristic, for both directions of increase and decrease of stock on hand and fill rate. For example, a value of '3' means that for 1%-point increase of volume fill rate, the stock on hand rises with 3.00%. We want this measure to be as small as possible, maybe even to be negative (meaning both less stock on hand and a higher volume fill rate). In case the volume fill rate decreased or did not change at all, we excluded this performance ratio, as we it would result in errors (0 in denominator), or in false positive results. Since a decrease of both stock on hand and volume fill rate would result in a division of two negative numbers, this ratio becomes positive, yet it reflects something no stock-keeping company would want. Therefore, we set SKUs with a lower volume fill rate than the standard approach arbitrarily large and grey out results for both out-of-stock configurations (in order to keep equal comparisons) in the result tables. SKUs with equal volume fill rates are set 'n/a', yet they are incorporated in the calculation of the average performance ratio. Furthermore, we are aware of the non-linear relationship between target fill rate and stock levels for normal demand, as depicted in Figure 5.1. As we have no direct insights on the nature of the relationship between observed fill rate and stock on hand levels for our heuristics, we just assume linearity, thereby justifying our performance ratio. Moreover, the relationship generally shows good resemblance with linearity for lower (target) fill rates, which is relevant since many heuristics have to improve a volume fill rate that is in general within the

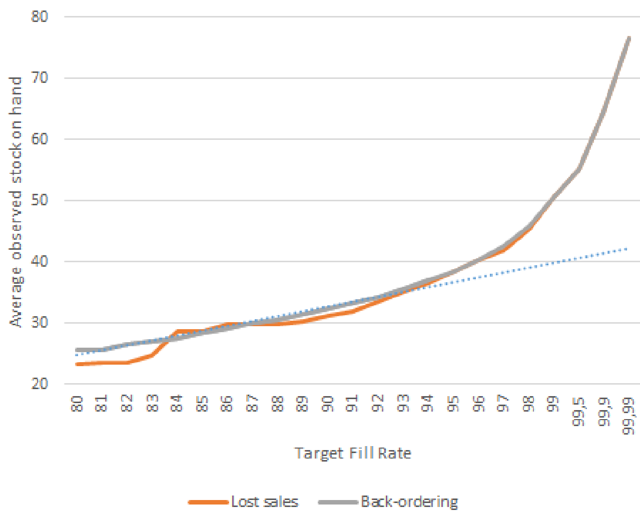

Figure 5.1: *Non-linear relationship between target fill rate and stock level, with a blue-dotted line showing the linear part of the figure*

linear part of the figure, that ranges approximately up to 94% (see the dashed line in Figure 5.1). So, for these lower fill rates the performance ratio predicts a linear increasing/decreasing stock on hand (yet a smaller increase than the exact performance ratio, since this also includes the non-linear, high fill rates). But for the higher fill rates stock on hand would increase by far more than the performance ratio times the increase of target fill rate. The difference between the lost sales and back-ordering setting is due to the slight adjustment for lost sales, as addressed in Section 3.4.1. Note: the slight bump for the lost sales line around 84%-87%, is most likely due to the batch size of IOQ, which is directly involved in the target fill rate calculations.

Finally, next to this efficiency indicator, we measure the effectiveness. To this extent, we compute the percentage of SKU-shop combinations for which the observed volume fill rate meets the target fill rate.

5.7.2 Performance literature methods

We now discuss the results of the literature methods of our final experiments, namely the "Add undershoot"-method for undershoot, and the target fill rate compound Poisson methods. Data are to be found in Tables C4 and C8 in Appendix C.

"Add undershoot"-method

Averaged on SKU-level, we have calculated relative differences of the stock on hand levels, absolute differences on observed volume fill rates, and their corresponding performance ratios, compared to the standard baseline performance, as denoted in Table 2.1 in Section 2.3. These differences are denoted in Table 5.4.

<u>Multi-modal order line size SKUs</u>
In the lost sales case, the "Add undershoot"-method performs rather bad compared to the standard methods: on average the observed volume fill rate increases with 5.06%-point, against a price of 21.39% more

stock on hand, yielding a performance ratio of 4.22. In case of back-ordering, this method performs similar, with 5.02%-point more observed volume fill rate and 21.19% more stock on hand, yielding a ratio of 4.22.

<u>Unimodal order line size SKUs</u>

For unimodal order line size SKUs and lost sales on the other hand, on average, it results in 1.98%-point more volume fill rate, against 13.16% more stock on hand, yielding a ratio of 6.66. When full back-ordering is assumed, observed volume fill rate increases with 1.76%-point, and stock on hand rises with 13.08%-point, resulting in a ratio of 7.42.

	"Add undershoot"-method						Compound Poisson method					
	Lost sales			*Back-ordering*			*Lost sales*			*Back-ordering*		
SKU	**SoH**	**β**	**Ratio**	**SoH**	**β**	**Ratio**	**SoH**	**β**	**Ratio**	**SoH**	**β**	**Ratio**
MM SKU 1	20.30%	2.71%	7.50	21.41%	2.23%	9.58	5.47%	2.74%	1.99	6.10%	2.64%	2.31
MM SKU 2	15.63%	1.16%	13.46	15.39%	1.16%	13.25	18.38%	1.10%	16.66	17.85%	1.10%	16.18
MM SKU 3	26.14%	3.48%	7.51	26.12%	3.48%	7.51	61.81%	5.88%	10.51	61.85%	5.88%	10.52
MM SKU 4	21.67%	6.49%	3.34	19.26%	5.89%	3.27	16.20%	7.38%	2.20	13.72%	6.43%	2.13
MM SKU 5	13.72%	3.72%	3.69	15.10%	4.18%	3.62	36.64%	9.80%	3.74	38.94%	10.43%	3.73
MM SKU 6	13.71%	1.42%	9.63	11.89%	1.42%	8.35	37.42%	1.85%	20.22	35.71%	2.87%	12.46
MM SKU 7	23.88%	3.66%	6.52	20.43%	4.50%	4.54	45.01%	6.73%	6.68	41.04%	7.15%	5.74
MM SKU 8	21.36%	5.57%	3.84	22.93%	5.33%	4.30	63.41%	11.31%	5.60	65.05%	10.91%	5.96
MM SKU 9	2.90%	0.80%	3.61	3.00%	0.45%	6.65	4.87%	0.86%	5.65	4.87%	0.41%	11.94
MM SKU 10	21.85%	1.51%	14.50	22.24%	2.59%	8.59	53.12%	5.19%	10.24	53.62%	6.80%	7.89
MM SKU 11	25.31%	4.66%	5.43	25.91%	4.19%	6.18	50.07%	9.31%	5.38	50.18%	7.69%	6.52
MM SKU 12	33.99%	6.69%	5.08	32.52%	5.70%	5.71	59.75%	10.87%	5.50	59.12%	10.54%	5.61
MM SKU 13	26.12%	10.58%	2.47	25.37%	9.87%	2.57	58.80%	16.78%	3.50	58.35%	16.54%	3.53
MM SKU 14	17.71%	5.06%	3.50	19.27%	5.04%	3.82	114.54%	8.81%	13.00	120.57%	9.10%	13.26
MM SKU 15	27.70%	6.74%	4.11	27.77%	7.58%	3.67	62.16%	12.63%	4.92	63.65%	15.19%	4.19
MM SKU 16	26.91%	7.79%	3.46	27.00%	7.55%	3.58	42.33%	12.34%	3.43	41.97%	12.59%	3.33
MM SKU 17	22.53%	9.06%	2.49	20.26%	7.55%	2.68	25.47%	9.11%	2.80	23.20%	7.34%	3.16
MM SKU 18	29.10%	8.97%	3.24	30.24%	10.58%	2.86	75.45%	15.83%	4.77	77.55%	17.88%	4.34
MM SKU 19	23.06%	4.14%	5.57	23.12%	4.77%	4.85	10.06%	2.70%	3.72	10.11%	3.41%	2.97
MM SKU 20	16.10%	7.51%	2.14	14.39%	6.09%	2.36	53.03%	12.10%	4.38	52.54%	10.58%	4.97
MM SKU 21	21.99%	6.04%	3.64	22.71%	5.58%	4.07	45.02%	7.15%	6.30	45.56%	6.41%	7.11
MM SKU 22	18.86%	3.66%	5.16	19.73%	4.64%	4.25	97.61%	12.30%	7.94	105.78%	13.37%	7.91
Avg. MM	21.39%	5.06%	4.22	21.19%	5.02%	4.22	47.12%	8.31%	5.67	47.61%	8.42%	5.65
UM SKU 23	15.59%	3.50%	4.45	15.06%	1.96%	7.67	6.11%	3.80%	1.61	6.06%	2.26%	2.68
UM SKU 24	22.58%	5.31%	4.25	21.07%	4.52%	4.66	2.08%	2.82%	0.74	0.64%	2.06%	0.31
UM SKU 25	10.84%	1.28%	8.49	11.59%	2.38%	4.87	4.19%	3.04%	1.38	3.84%	2.76%	1.40
UM SKU 26	14.38%	0.38%	37.84	14.57%	0.43%	33.76	-27.44%	-0.46%	59.21	-27.52%	-0.77%	35.84
UM SKU 27	15.62%	1.33%	11.78	15.94%	1.47%	10.85	8.82%	3.24%	2.72	8.85%	2.43%	3.64
UM SKU 28	12.23%	1.51%	8.11	11.93%	1.21%	9.88	22.74%	4.66%	4.88	23.55%	4.79%	4.92
UM SKU 29	7.88%	0.42%	18.98	8.18%	0.38%	21.78	8.98%	1.73%	5.18	9.23%	1.75%	5.28
UM SKU 30	11.46%	3.20%	3.58	11.36%	3.28%	3.46	-19.96%	-5.96%	3.35	-23.30%	-7.53%	3.09
UM SKU 31	12.90%	2.49%	5.19	14.25%	2.53%	5.64	-12.50%	-0.09%	145.49	-8.13%	-0.83%	9.79
UM SKU 32	11.67%	1.11%	10.53	10.74%	0.83%	13.01	9.58%	2.69%	3.56	9.27%	2.60%	3.57
UM SKU 33	9.96%	1.38%	7.24	9.57%	1.12%	8.52	-57.12%	-13.79%	4.14	-56.73%	-14.41%	3.94
UM SKU 34	10.30%	2.25%	4.58	10.43%	2.48%	4.20	-52.19%	-12.03%	4.34	-53.00%	-12.57%	4.22
UM SKU 35	15.72%	3.68%	4.27	15.57%	2.29%	6.79	25.26%	4.78%	5.28	24.70%	3.99%	6.19
UM SKU 36	16.36%	1.70%	9.60	16.62%	1.66%	10.03	5.81%	2.01%	2.90	5.71%	1.80%	3.16
UM SKU 37	16.86%	1.98%	8.51	15.99%	1.41%	11.31	8.87%	3.02%	2.94	8.28%	2.71%	3.06
UM SKU 38	6.23%	0.13%	47.81	6.33%	0.23%	28.02	6.53%	0.48%	13.67	6.79%	0.57%	11.85
Avg. UM	13.16%	1.98%	6.66	13.08%	1.76%	7.42	9.91%	2.93%	4.08	9.72%	2.52%	4.19
Difference	62.51%	156.18%	-55.34%	62.02%	184.84%	-54.16%	375.68%	183.17%	66.26%	389.77%	234.24%	58.25%
Avg. total	17.27%	3.52%	4.91	17.13%	3.39%	5.06	28.51%	5.62%	5.07	28.66%	5.47%	5.24

Table 5.4: *Relative differences of stock levels and absolute differences of observed volume fill rates through the "Add undershoot"- and "Compound Poisson"-method, compared with the standard method. Results that yield a decrease of volume fill rate are greyed out and not included in the averaged results.*

Target fill rate compound Poisson method

Multi-modal order line size SKUs
In the lost sales case, the compound Poisson method performs very well, but rather inefficient, compared to the standard methods: on average the observed volume fill rate increases with 8.31%-point, against a high price of 47.12% more stock on hand, yielding a performance ratio of 5.67. In case of back-ordering, this method performs similar with 8.42%-point more observed volume fill rate and 47.61% more stock on hand, yielding a ratio of 5.65.

Unimodal order line size SKUs
For unimodal order line size SKUs and lost sales on the other hand, on average, this method results in 2.93%-point more volume fill rate, against 9.91% more stock on hand, yielding a ratio of 4.08. When full back-ordering is assumed, observed volume fill rate increases with 2.52%-point, and stock on hand rises with 9.72%-point, resulting in a ratio of 4.19. However, SKUs 26, 30, 31, 33, and 34 were not incorporated, since they perform worse than the standard situation. Three out of five (UM SKU 30, 33, and 34) perform even considerable worse. Hence, following from Section 2.4, condition [a.] is met, but [b.] is unmet. Thus, this method is not recommendable for unimodal order line size SKUs. An explanation that the target fill rate compound Poisson method (using the empirical distribution for the order line size distribution) works well for multi-modal order line size SKUs, but performs unstable and worse for unimodal order line size SKUs, is that the first category roughly has only few order lines per period of time, compared to the latter. So, in general (no clear quantified boundaries could be found) holds: the bigger the number of orders per period of time, the bigger the resemblance with the standard method. Since the standard method assumes normal demand, this similarity is logical, since a compound Poisson distribution can be well approximated by a normal distribution for large demand (Axsäter, 2006).

5.7.3 Performance own heuristics

Finally, we analyse the performance of our own derived “Overwrite”, “Max” and “First moment β^*”-methods, and we proceed in this order. Data are to be found in tables C5, C6, C7 in appendix C. Again, averaged on SKU-level, we have calculated relative differences of the stock on hand levels, absolute differences on observed volume fill rates, and their corresponding performance ratios, compared to the standard baseline performance, as denoted in Table 2.1 in Section 2.3. These differences are denoted in Table 5.5.

“Overwrite”-method

Multi-modal order line size SKUs
The “Overwrite”-method performs very unstable, for only 36.36% of the SKUs provided a positive result on the difference of the volume fill rate compared to the standard methods. However, on average, SKUs 3, 6, 7, 8, 14, 18, 20, and 21 yield an average of 3.26%-point more volume fill rate, against on average 10.16% more stock on hand for lost sales, yielding a ratio of 3.74. In the back-ordering setting, these values amount 3.11%-point, 10.07% and 5.81 respectively. Yet, since this occurs in only 36.36% of the multi-modal SKUs, this is a very unstable method for this type of SKU.

Unimodal order line size SKUs

This method acts even more unstable for unimodal order line size SKUs, with only UM SKU 23 having an increase in volume fill rate compared to the standard situation. Thus, this method is completely non-applicable for this type of SKU. Since only 1 SKU is considered in the average, the average differences and performance ratios are arbitrary.

	"Overwrite"-method						"Max"-method						"First moment β^*"-method					
	Lost sales			*Back-ordering*			*Lost sales*			*Back-ordering*			*Lost sales*			*Back-ordering*		
SKU	**SoH**	β	**Ratio**	**SoH**	β	**Ratio**	**SoH**	β	**Ratio**	**SoH**	β	**Ratio**	**SoH**	β	**Ratio**	**SoH**	β	**Ratio**
MM SKU 1	-12.97%	-3.99%	3.25	-13.39%	-3.77%	3.55	1.11%	-0.33%	-3.33	1.75%	0.00%	n/a	28.47%	3.55%	8.01	28.46%	3.80%	7.48
MM SKU 2	11.35%	-2.03%	-5.58	10.83%	-2.56%	-4.24	12.50%	-0.46%	-26.90	12.28%	-0.46%	-26.44	42.83%	0.06%	n/a	42.53%	0.06%	n/a
MM SKU 3	17.58%	3.48%	5.05	16.99%	2.95%	5.76	18.14%	3.48%	5.21	18.23%	2.95%	6.18	73.23%	4.95%	14.78	73.23%	4.95%	14.78
MM SKU 4	-12.70%	-1.20%	10.54	-15.40%	-4.13%	3.73	1.31%	0.66%	1.97	0.62%	0.66%	0.94	16.72%	5.10%	3.28	13.26%	2.74%	4.84
MM SKU 5	-9.15%	0.69%	-13.33	-8.77%	0.01%	n/a	2.53%	1.85%	1.37	3.14%	1.72%	1.83	0.01%	3.75%	0.00	0.55%	3.07%	0.18
MM SKU 6	6.45%	-0.21%	-30.54	5.27%	0.31%	17.29	9.88%	1.42%	6.94	8.85%	1.42%	6.21	47.68%	2.33%	20.46	46.53%	3.09%	15.07
MM SKU 7	-4.96%	-0.59%	8.42	-8.64%	-0.68%	12.76	4.73%	-0.40%	-11.70	1.02%	0.55%	1.86	31.79%	4.57%	6.96	28.33%	4.55%	6.23
MM SKU 8	-1.83%	-0.30%	6.07	-1.70%	0.42%	-4.03	5.25%	1.61%	3.27	6.12%	1.85%	3.32	35.17%	8.85%	3.98	36.19%	7.46%	4.85
MM SKU 9	2.51%	0.43%	5.80	2.41%	0.15%	15.69	2.54%	0.43%	5.87	2.47%	0.15%	16.13	12.26%	1.56%	7.86	12.07%	0.84%	14.39
MM SKU 10	-2.74%	-1.00%	2.73	-2.73%	-0.76%	3.58	4.65%	1.10%	4.22	4.65%	1.10%	4.22	22.51%	1.47%	15.29	22.12%	3.36%	6.58
MM SKU 11	2.36%	0.22%	10.55	2.21%	-0.04%	-55.87	9.16%	2.15%	4.26	8.94%	1.65%	5.41	59.58%	6.08%	9.80	59.45%	5.24%	11.34
MM SKU 12	-13.22%	-2.91%	4.54	-13.49%	-3.63%	3.72	3.02%	2.26%	1.33	3.23%	1.89%	1.70	41.72%	8.30%	5.03	41.20%	8.87%	4.65
MM SKU 13	15.05%	7.86%	1.91	13.83%	6.88%	2.01	19.41%	8.77%	2.21	18.29%	8.14%	2.25	38.20%	13.11%	2.91	36.45%	13.00%	2.80
MM SKU 14	-4.65%	-1.51%	3.08	-3.79%	-1.41%	2.68	0.24%	0.90%	0.27	1.28%	0.95%	1.35	22.95%	4.64%	4.95	26.26%	4.89%	5.38
MM SKU 15	11.49%	0.62%	18.48	13.96%	3.13%	4.46	17.47%	3.57%	4.89	18.29%	4.11%	4.45	43.29%	8.34%	5.19	44.52%	10.48%	4.25
MM SKU 16	-6.78%	-2.34%	2.89	-7.22%	-1.46%	4.96	3.43%	0.88%	3.88	3.04%	1.55%	1.96	12.64%	6.67%	1.90	12.81%	6.99%	1.83
MM SKU 17	-3.93%	0.76%	-5.17	-6.16%	-0.85%	7.22	4.34%	2.92%	1.49	3.07%	2.22%	1.38	39.00%	8.89%	4.39	37.69%	6.29%	5.99
MM SKU 18	-4.58%	0.13%	-35.14	-5.25%	-0.09%	56.49	3.72%	0.97%	3.82	4.22%	0.75%	5.62	60.79%	13.62%	4.46	63.23%	15.67%	4.03
MM SKU 19	-28.00%	-11.25%	2.49	-28.16%	-11.38%	2.47	0.11%	0.05%	2.41	-0.11%	0.51%	-0.21	-12.30%	-4.14%	2.97	-12.76%	-4.12%	3.10
MM SKU 20	10.82%	6.60%	1.64	8.81%	5.18%	1.70	13.00%	6.92%	1.88	11.18%	5.50%	2.03	63.10%	14.03%	4.50	61.20%	12.74%	4.80
MM SKU 21	14.81%	2.93%	5.06	14.34%	2.11%	6.79	16.73%	4.61%	3.63	16.88%	3.79%	4.45	41.80%	6.87%	6.08	42.33%	6.13%	6.90
MM SKU 22	18.18%	3.45%	5.27	19.00%	4.43%	4.28	18.50%	3.66%	5.06	19.35%	4.64%	4.17	79.57%	11.99%	6.64	85.58%	13.15%	6.51
Avg. MM	10.16%	3.26%	3.74	10.07%	3.11%	5.81	8.08%	2.54%	3.37	7.99%	2.40%	3.86	38.73%	6.60%	6.82	38.76%	6.54%	6.64
UM SKU 23	-6.61%	2.66%	-2.49	-7.85%	1.03%	-7.62	3.53%	3.09%	1.14	1.72%	1.51%	1.14	-8.82%	1.87%	-4.70	-10.46%	0.24%	-43.02
UM SKU 24	-26.52%	-7.10%	3.74	-27.86%	-7.31%	3.81	0.00%	0.00%	n/a	0.00%	0.00%	n/a	10.77%	3.35%	3.21	9.90%	1.82%	5.44
UM SKU 25	-10.56%	-1.76%	6.01	-12.15%	-1.40%	8.70	-0.24%	-0.65%	0.37	-0.30%	-0.22%	1.37	-7.60%	-1.09%	7.00	-9.21%	-1.25%	7.36
UM SKU 26	-37.95%	-2.57%	14.79	-37.53%	-2.63%	14.27	0.00%	0.00%	n/a	0.00%	0.00%	n/a	-27.19%	-1.18%	23.04	-26.85%	-1.20%	22.39
UM SKU 27	-10.07%	-0.50%	20.18	-10.63%	-0.79%	13.50	0.16%	0.02%	6.69	0.29%	0.02%	11.76	-6.12%	-0.19%	33.06	-6.38%	-0.22%	29.05
UM SKU 28	-17.09%	-3.00%	5.69	-17.93%	-3.19%	5.63	0.63%	0.11%	5.96	0.71%	0.00%	n/a	-9.79%	-1.84%	5.32	-11.07%	-2.17%	5.10
UM SKU 29	2.92%	-0.34%	-8.49	3.12%	-0.47%	-6.69	5.14%	0.33%	15.53	5.25%	0.26%	20.15	4.65%	-0.25%	-18.51	5.04%	-0.34%	-14.81
UM SKU 30	-19.00%	-5.55%	3.43	-22.48%	-7.73%	2.91	0.96%	0.53%	1.81	0.42%	0.46%	0.90	-9.08%	-2.88%	3.15	-11.15%	-4.40%	2.54
UM SKU 31	-28.40%	-7.69%	3.69	-26.22%	-7.96%	3.29	0.00%	0.00%	n/a	0.00%	0.00%	n/a	-17.77%	-5.44%	3.27	-16.61%	-5.92%	2.81
UM SKU 32	-14.95%	-2.77%	5.39	-16.05%	-2.89%	5.55	0.13%	0.00%	n/a	0.26%	0.00%	n/a	-10.50%	-1.87%	5.63	-12.16%	-1.91%	6.38
UM SKU 33	-60.04%	-16.43%	3.65	-59.20%	-16.26%	3.64	0.00%	0.00%	n/a	0.00%	0.00%	n/a	-53.42%	-12.94%	4.13	-53.73%	-13.14%	4.09
UM SKU 34	-58.27%	-16.07%	3.63	-57.47%	-15.89%	3.62	0.00%	0.00%	n/a	0.00%	0.00%	n/a	-52.22%	-13.07%	4.00	-51.92%	-12.59%	4.12
UM SKU 35	6.54%	0.79%	8.27	6.01%	-0.26%	-23.22	7.78%	1.50%	5.17	7.42%	0.46%	16.30	50.18%	6.65%	7.54	49.97%	5.79%	8.63
UM SKU 36	-20.91%	-2.98%	7.01	-20.68%	-3.15%	6.57	0.28%	0.03%	7.97	0.23%	0.03%	7.76	-11.08%	-1.15%	9.66	-10.56%	-1.10%	9.62
UM SKU 37	-23.43%	-2.76%	8.49	-24.38%	-2.36%	10.33	0.00%	0.00%	n/a	0.00%	0.00%	n/a	-0.09%	0.28%	-0.33	-0.83%	0.43%	-1.95
UM SKU 38	-0.18%	-0.06%	2.89	-0.33%	-0.11%	3.08	1.99%	0.01%	294.56	2.01%	-0.04%	-53.58	0.44%	-0.53%	-0.83	0.13%	-0.49%	-0.26
Avg. UM	-0.03%	1.72%	2.89	-0.92%	0.39%	-15.42	1.33%	0.40%	6.32	1.16%	0.20%	9.67	13.01%	3.04%	1.43	12.14%	2.07%	-7.72
Difference	-253.84%	22.59%	-250.29%	-228.28%	201.55%	-176,29%	507.35%	531.79%	-46.75%	586.20%	1122.33%	-60.03%	197.67%	117.36%	376.43%	219.18%	216.00%	-186.02%
Avg. total	1.78%	2.96%	0.60	1.11%	2.07%	0.54	4.70%	1.47%	3.20	4.58%	1.30%	3.53	25.87%	4.82%	5.37	25.45%	4.31%	5.91

Table 5.5: *Relative differences of stock levels and absolute differences of observed volume fill rates through the "Overwrite"-, "Max"- and "First moment β^*"-method, compared with the standard method. Results that yield a decrease of volume fill rate are greyed out and not included in the averaged results.*

"Max"-method

Multi-modal order line size SKUs

Next heuristic we have measured is the "Max"-method. In the lost sales case, the "Max"-method performs very similar compared to the standard methods. This was expected, since in this method we take the maximum of the safety stock from the standard approach, and the undershoot. To be precise, 89.39% of the multi-modal order line size SKUs and 86.21% of the unimodal order line size SKUs provided equal safety stocks and thereby equal results as the standard methods. This fractions are obtained from the lost sales case, for back-ordering these amounts are 89.50% and 86.40%. Only MM SKU 1 and MM SKU 2 are skipped due to under-performance of the volume fill rate compared to the standard method. For lost sales, on average the observed volume fill rate increases with 2.54%-point, against the small price of 8.08% more stock on hand, yielding a performance ratio of 3.37. In the back-ordering setting, these values amount 2.40%-point, 7.99% and 3.86 respectively. All in all, this method is stable but very conservative.

It does not provide enough increase of volume fill rate.

Unimodal order line size SKUs
For unimodal order line size SKUs on the other hand, on average and provided lost sales, this method yields 0.40%-point more volume fill rate, against 1.33% more stock on hand, yielding a ratio of 6.32. For the back-ordering setting, these values amount 0.20%-point, 1.16% and 9.67.

"First moment β^*"-method

Multi-modal order line size SKUs
Finally, we have measured the "First moment β^*"-method. In the lost sales case, the "First moment β^*"-method performs rather 'amplified' compared to the standard methods. By this we mean that in general this method yields rather much more volume fill rate, but also extremely much more stock on hand. This amplification is stronger for multi-modal order line size SKUs - having rather high median order line sizes - than for unimodal order line size SKUs. For lost sales, on average the observed volume fill rate increases with 6.60%-point, against the high price of 38.73% more stock on hand, yielding a performance ratio of 6.82. In the back-ordering setting, these values amount 6.54%-point, 38.76% and 6.64 respectively. 9.10% of the SKUs have a lower volume fill rate than the standard method, which seems reasonable well.

Unimodal order line size SKUs
For unimodal order line size SKUs on the other hand, on average and provided lost sales, this method yields 3.04%-point more volume fill rate, against 13.01% more stock on hand, yielding a ratio of 1.43. For the back-ordering setting, these values amount 2.07%-point, 12.14% and -7.72. Next, 75% of the SKUs have a lower volume fill rate than the standard method, which is very unstable. Thus, this method is completely non-applicable for this type of SKU.

5.7.4 Comparison result

Now that we have analysed the performance of all heuristics, we can compare and clarify them. First of all, we make a listing based on the effectiveness of the method, compared to the baseline. To this extent we make an ordered listing, based on the earlier introduced performance ratio. We take into account that we rather want higher service levels against higher stock on hands than vice versa. Furthermore, methods that were found to be too unstable (we took 10% as maximum for the percentage of SKUs that resulted in a lower volume fill rate than in the standard approach), are greyed out again. See Table 5.6. We were surprised by the fact that compound Poisson appears to be rather inefficient, since it ought be an appropriate model for slow-moving, non-unit order line size SKUs. Next, it is clear that the unimodal order line size SKUs hardly have potential for improvement through a different heuristic. This can also be a result of the fact that this type of SKU is not the main topic of this research, thus we did not attempt to improve this type in the first place. Next, we expected "Max"-method to perform rather efficient, as this acts as hybrid with the standard model, hence this would only increase the stock on hand slightly.

	Lost sales		**Back-ordering**	
Rank	*MM (ratio)*	*UM (ratio)*	*MM (ratio)*	*UM (ratio)*
1.	"Max" (3.37)	"Max" (6.32)	"Max" (3.86)	"Max" (9.67)
2.	"Add undershoot" (5.45)	"Add undershoot" (12.20)	"Add undershoot"" (5.29)	"Add undershoot" (11.53)
3.	Compound Poisson (6.78)	"First moment β^*" (1.43)	Compound Poisson (6.62)	Compound Poisson (4.19)
4.	"First moment β^*" (6.82)	Compound Poisson (4.08)	"First moment β^*" (6.64)	"First moment β^*" (-7.72)
5.	"Overwrite" (3.74)	"Overwrite" (n/a)	"Overwrite" (5.81)	"Overwrite" (n/a)

Table 5.6: *Ranking of heuristics per out-of-stock setting. Methods that yield a decrease of volume fill rate for a significant number of SKUs, are greyed out and not included in the averaged results. Overwrite is not applicable for unimodal order line size SKUs, since only 1 SKU yields an improvement, hence basing the average on one observation.*

Next to this efficiency ratios, we want to regard the percentage of SKU-shop combinations that for which the input target fill rate was met. This gives the overall effectiveness performance of a heuristic, see Table 5.7.

		Standard	**"Add undershoot"**	**Compound Poisson**	**"Overwrite"**	**"Max"**	**"First moment β^*"**
MM SKUs	*Lost sales*	67.61%	78.47%	<u>84.79%</u>	66.93%	71.48%	80.12%
	Back-ordering	67.61%	78.51%	<u>84.90%</u>	66.81%	71.67%	80.07%
UM SKUs	*Lost sales*	64.15%	<u>71.97%</u>	64.95%	50.46%	65.07%	55.38%
	Back-ordering	64.52%	<u>72.24%</u>	64.67%	50.69%	65.07%	55.75%

Table 5.7: *Comparison of different methods on the percentage of SKU-shop combinations per type of SKU for which target fill rate is attained*

We underlined the highest value per type of SKU and out-of-stock configuration. Thus, when implementing the compound Poisson approach for multi-modal order line size SKUs for lost sales, the observed volume fill rates would be met in $\frac{84.79\% - 67.61\%}{67.61} = 25.41\%$ more of the SKU-shop combinations. For back-ordering, this would be 25.57%. When implementing the "Add undershoot"-method for unimodal order line size SKUs for lost sales, the observed volume fill rates would be met in 12.19% more of the SKU-shop combinations, and given back-ordering this would be 11.97%. We were astonished by the poor performance in the standard model, even more underlining the power a more sophisticated model such as Slim4 - which would show better results. Next, the "Overwrite"-method performs poorly as well. This was to be expected, since it makes a rather arbitrary application of undershoot in the safety stock calculations. As expected, this approach leads to an underestimation of the safety stock for unimodal order line size SKUs (see Table 5.9), thereby leading to poor performance.

Next, we take a look at the implications for average stock on hand and reorder points. First, see Table 5.8 for the average stock on hand levels per type of SKU and out-of-stock configuration.

		Standard	**"Add undershoot"**	**Compound Poisson**	**"Overwrite"**	**"Max"**	**"First moment β^*"**
MM SKUs	*Lost sales*	<u>9.15</u>	10.89	12.66	9.17	9.79	12.22
	Back-ordering	<u>9.15</u>	10.89	12.68	9.13	9.78	12.22
UM SKUs	*Lost sales*	22.31	25.07	18.72	<u>16.03</u>	22.55	17.94
	Back-ordering	22.38	25.13	18.72	<u>16.02</u>	22.59	17.90

Table 5.8: *Comparison of different methods on the average stock on hand per type of SKU*

In Table 5.8, we see that the stock on hand levels would increase when implementing the compound

Poisson and "Add undershoot"-method. The latter yields a rather small increase of stock on hand. This is logical, since undershoot generally is only small. Furthermore, for unimodal order line size SKUs, some having near-pure Poisson demand (i.e. larger probabilities of order line sizes of 1), the compound Poisson method yields less stock on hand. As we see in Table 5.7, its performance is also low for this type of SKUs. This makes sense, since compounding the order line size distribution for near-pure Poisson leads to a normal Poisson distribution. This does not perform well for slow-moving SKUs, as we also see in our experiment, which thereby follows general inventory theory (Silver et al., 2017). Next, we compare the average safety stocks, see Table 5.9.

		Standard	**"Add undershoot"**	**Compound Poisson**	**"Overwrite"**	**"Max"**	**"First moment β^*"**
MM SKUs	*Lost sales*	1.87	3.66	5.61	1.79	2.49	5.04
	Back-ordering	1.89	3.68	""	""	2.50	""
UM SKUs	*Lost sales*	9.60	12.48	5.81	2.87	9.85	5.02
	Back-ordering	9.65	12.52	""	""	9.89	""

Table 5.9: *Comparison of different methods on the average safety stock per type of SKU*

In Table 5.9 we see that safety stocks from methods *compound Poisson*, "Overwrite", and "First moment β^*" are indeed equal for both out-of-stock configurations. This functions as verification, for we did not differ for both settings in these methods. It is clear to see that safety stocks are underestimated in the standard model, for both types of SKUs.

5.7.5 Sensitivity analysis standard model

Other methods generally yield a higher service level, against higher stock on hand. A different method to increase the fill rate would therefore be to keep increasing the target service level, until the 'true' target fill rate is reached. In this way there would be no reason to alter the inventory model. Therefore, we conducted a sensitivity analysis on five target service levels using the standard model. The five levels are 80%, 90%, 95%, 98%, and 99%. Through this brief analysis we are able to compare current performance from the standard model with alternative approaches. The results can be found in Table 5.10. Note: we restricted ourselves to the lost sales case, since there is no significant difference in case of full back-ordering.

We see in Table 5.10 that how much one would increase the target service level for multi-modal order line size SKUs in the standard model, observed volume fill rates such as in alternative methods like the compound Poisson method, are never achieved. Moreover, the highest average observed volume fill rate for multi-modal order line size SKUs amounts 91.76%, whereas compound Poisson yields an average of 95.29%. For unimodal on the other hand, we see that when the target fill rates are all set on 99%, an average observed volume fill rate of 94.94% is achieved. This greatly resembles the average observed volume fill rate of the "Add undershoot" - given lost sales - of 94.36%, which requires an average stock of 25.07. When simply setting the target fill rate equal to 99%, one would achieve very similar volume fill rates as when incorporating the undershoot effect, yet against less stock on hand of 22.78. Thus, although the "Add undershoot" method outperforms the current standard model, this standard model might be tricked into yielding better result. We stick to the aforementioned method as 'best' methods, but we include this analysis in our conclusions and recommendations.

	β: 80%*		*β*: 90%*		*β*: 95%*		*β*: 98%*		*β*: 99%*	
SKU	**SoH**	**VFR**	**SoH**	**VFR**	**SoH**	**VFR**	**SoH**	**VFR**	**SoH**	**VFR**
MM SKU 1	10.06	83.46%	10.64	85.91%	12.69	88.97%	14.05	91.18%	15.19	93.34%
MM SKU 2	12.45	97.44%	12.55	97.97%	12.84	97.97%	13.27	99.48%	14.21	99.54%
MM SKU 3	12.89	94.33%	13.11	94.33%	13.18	95.46%	13.06	93.61%	14.10	94.34%
MM SKU 4	4.79	77.80%	5.01	79.90%	5.85	85.58%	6.72	88.11%	7.37	92.00%
MM SKU 5	4.83	89.48%	5.10	89.71%	5.46	91.30%	6.35	92.99%	6.90	95.61%
MM SKU 6	5.39	92.93%	5.45	92.63%	5.64	94.08%	6.66	97.50%	6.79	97.76%
MM SKU 7	5.16	83.47%	5.42	83.88%	5.98	88.87%	6.95	92.44%	7.76	94.18%
MM SKU 8	5.39	82.06%	5.64	82.63%	6.09	85.99%	6.88	89.48%	7.25	90.07%
MM SKU 9	33.85	96.48%	33.85	96.48%	33.85	96.48%	33.90	96.66%	34.18	97.22%
MM SKU 10	5.00	78.58%	5.32	81.99%	6.40	87.46%	7.54	93.86%	7.96	95.20%
MM SKU 11	9.91	83.05%	10.30	84.71%	10.99	86.59%	12.72	89.74%	13.37	91.74%
MM SKU 12	3.81	66.95%	4.73	76.01%	5.28	81.33%	6.26	88.38%	7.00	89.19%
MM SKU 13	5.75	78.56%	5.92	79.34%	6.14	79.63%	7.23	84.40%	7.24	85.41%
MM SKU 14	4.72	80.59%	4.95	82.46%	6.14	87.61%	6.94	92.57%	7.46	92.92%
MM SKU 15	5.12	80.52%	5.31	80.52%	5.45	80.52%	6.49	88.74%	6.65	89.65%
MM SKU 16	5.09	64.07%	5.37	68.13%	6.14	75.93%	7.23	82.53%	7.75	84.30%
MM SKU 17	4.92	74.31%	5.13	76.12%	5.49	79.01%	6.55	85.61%	7.08	89.21%
MM SKU 18	5.06	63.49%	5.43	65.63%	6.62	73.55%	7.60	80.82%	8.14	83.93%
MM SKU 19	5.28	69.92%	6.58	78.64%	7.60	83.74%	9.09	89.87%	9.89	92.12%
MM SKU 20	5.54	82.98%	5.59	83.34%	5.77	84.15%	5.88	85.73%	6.55	89.38%
MM SKU 21	5.35	81.19%	5.41	81.19%	5.58	81.33%	6.62	91.11%	6.67	91.11%
MM SKU 22	4.24	86.95%	4.24	86.95%	4.26	86.95%	4.26	86.95%	4.99	90.58%
Avg. MM	7.48	81.30%	7.78	83.11%	8.34	86.02%	9.19	90.08%	9.75	91.76%
UM SKU 23	21.39	95.02%	21.46	95.24%	23.07	95.60%	26.15	96.03%	28.18	96.28%
UM SKU 24	7.57	69.84%	9.13	72.63%	10.72	75.39%	12.69	79.28%	13.84	81.74%
UM SKU 25	7.77	92.35%	7.94	92.42%	8.44	93.27%	9.68	95.47%	10.27	96.32%
UM SKU 26	18.42	94.29%	22.60	95.33%	26.40	96.65%	30.59	97.92%	33.45	98.18%
UM SKU 27	18.63	96.60%	18.97	96.69%	20.23	97.24%	22.77	97.91%	24.41	98.12%
UM SKU 28	7.50	90.51%	7.75	90.97%	8.55	92.57%	9.73	95.27%	10.47	96.34%
UM SKU 29	26.33	97.80%	26.44	97.80%	26.62	97.84%	27.20	98.14%	28.32	98.34%
UM SKU 30	17.69	79.82%	18.62	81.55%	20.59	84.82%	23.59	88.99%	25.50	90.49%
UM SKU 31	8.36	75.15%	9.09	77.04%	10.42	80.20%	12.15	83.13%	13.46	85.99%
UM SKU 32	7.34	91.18%	7.45	91.24%	8.08	93.23%	9.27	95.96%	9.95	96.64%
UM SKU 33	30.26	83.78%	38.29	90.56%	45.16	93.29%	52.96	95.04%	58.18	96.00%
UM SKU 34	29.60	84.53%	37.70	90.16%	44.35	92.90%	51.77	96.33%	57.27	97.47%
UM SKU 35	9.27	91.15%	9.38	91.38%	9.58	91.70%	10.15	94.79%	10.92	94.81%
UM SKU 36	8.91	88.68%	9.64	90.11%	10.78	92.58%	12.45	94.73%	13.70	96.10%
UM SKU 37	8.28	90.26%	9.09	91.32%	10.45	93.50%	12.15	95.86%	13.32	96.86%
UM SKU 38	18.30	96.97%	18.30	96.97%	18.47	97.10%	19.31	98.88%	20.00	99.33%
Avg. UM	15.35	88.62%	17.32	90.09%	18.87	91.74%	21.41	93.98%	22.78	94.94%
Difference	-51.26%	-8.26%	-55.11%	-7.74%	-55.81%	-6.24%	-57.07%	-4.15%	-57.21%	-3.34%
Avg. total	11.42	84.96%	13.44	86.60%	13.60	88.88%	15.30	92.03%	16.63	93.35%

Table 5.10: Stock levels and volume fill rate per β^*, through the standard model, for lost sales .

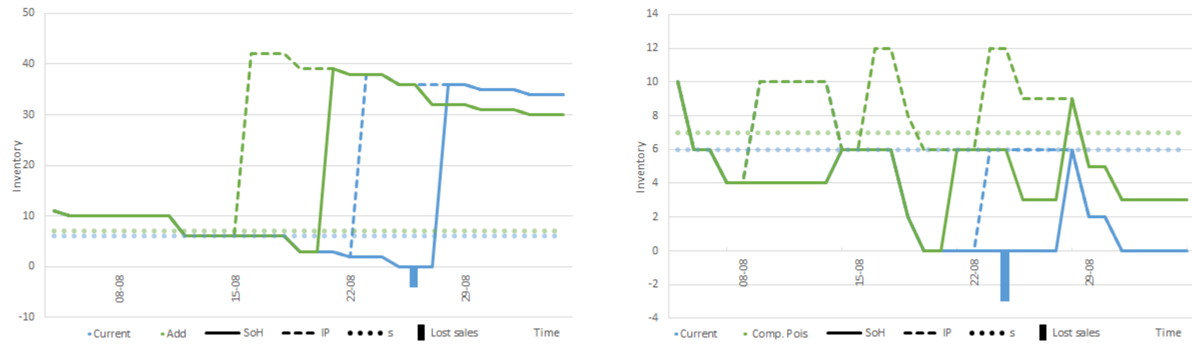

(a) *A toy SKU, showing better performance for the "Add undershoot"-method due to a more accurate estimation of the safety stock*

(b) *A tea glass SKU, showing better performance for the compound Poisson method due to a more accurate estimation of the safety stock*

Figure 5.2: Inventory positions of an arbitrary multi-modal and unimodal order line size SKU for four weeks, given lost sales. The bars depict sales that are lost due to insufficient inventory

5.7.6 Inferences best methods

Now that we have selected the two best performing methods, namely "Add undershoot" for unimodal order line size SKUs and target fill rate compound Poisson for multi-modal order line size SKUs, we provide some inferences using these methods. First, in Figures 5.2a and 5.2b we see the differences between the standard method, and the "Add undershoot"-method and compound Poisson method respectively.

In Figure 5.3 we depicted the transitions of relative performance per SKU, when implementing their new heuristic. We depict a weighted (using the total demand of the SKU) centre of gravity per SKU as cloud, since this is a more visually attractive measure than depicting an average. For the multi-modal order line size SKUs it is clear to see that many SKUs currently under-perform, while they would perform very well under the compound Poisson approach. Even so, one could argue that they over-perform, although we would only suggest this if a significant number of SKUs would on average reach 100% observed volume fill rate. Next, for the unimodal order line size SKUs, we see a more conservative improvement, and most SKUs perform around their target under the "Add undershoot"-method. We also added the results from the 99% target fill rate from the sensitivity analysis, as comparison. We see that this 'tweaked' standard method performs worse that the compound Poisson for multi-modal order line sizes SKUs. In addition, for unimodal order line size SKUs, performance in terms of observed volume fill rate is very similar.

Next, when we look at figure Figure 5.4 - where we depicted the stock on hand level of the new method compared to the stock on hand level of the standard method - we see that the compound Poisson method's stock on hand levels for multi-modal order line size SKUs is not similar to the ones using the standard method. Differences are very big and varying, especially for low stock on hand levels. For unimodal order line size SKUs on the other hand, the "Add undershoot"-method (as the name already suggests) *adds* - a reasonable amount of - stock compared to the standard situation. This effect gets larger for larger baseline stock on hand levels.

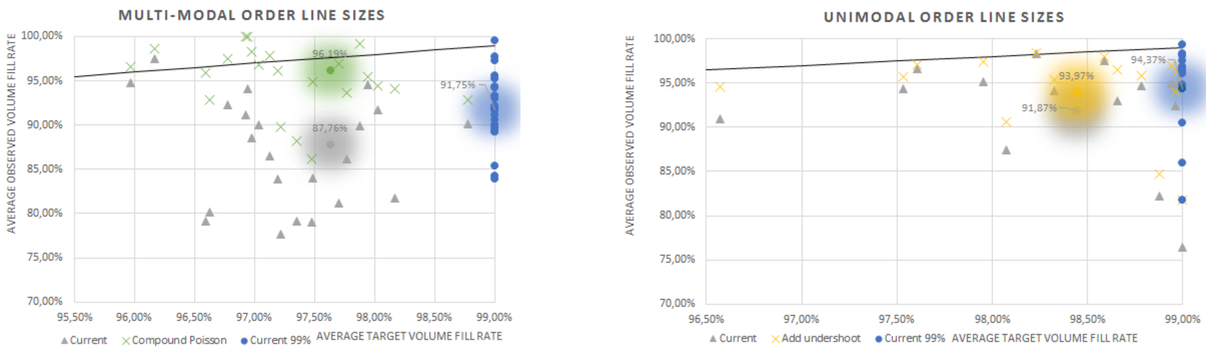

Figure 5.3: *The observed volume fill rates for the compound Poisson and "Add undershoot"-method set out against the SKUs' target volume fill rates*

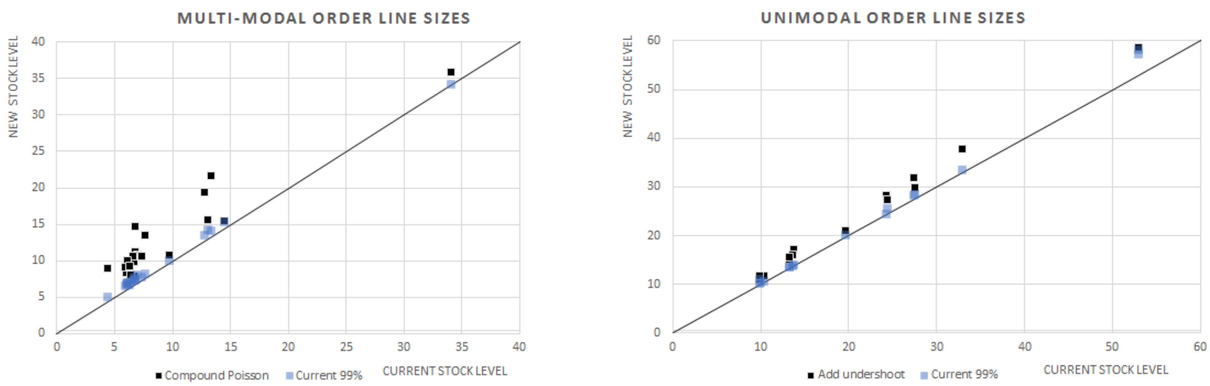

Figure 5.4: *The observed stock on hand levels for the compound Poisson and "Add undershoot"-method set out against the SKUs' baseline stock on hand levels*

Finally, we refer to Figure 5.5, where we clearly see the difference in performance between the three methods at hand. Multi-modal order line sizes SKUs perform best on volume fill rate when modelling them with the compound Poisson distribution, with slightly more stock on hand. On the other hand, the "Add undershoot"-method has a relatively smaller increase of volume fill rate, and needs more stock on hand. For unimodal order line size SKUs, we see that the compound Poisson method needs less stock hand but also gets somewhat more unstable (many SKUs' volume fill rates improve, yet the average is lower, compared to the standard method). The "Add undershoot"-method provides more performance with regards to volume fill rate, but the stock on hand also increases considerably.

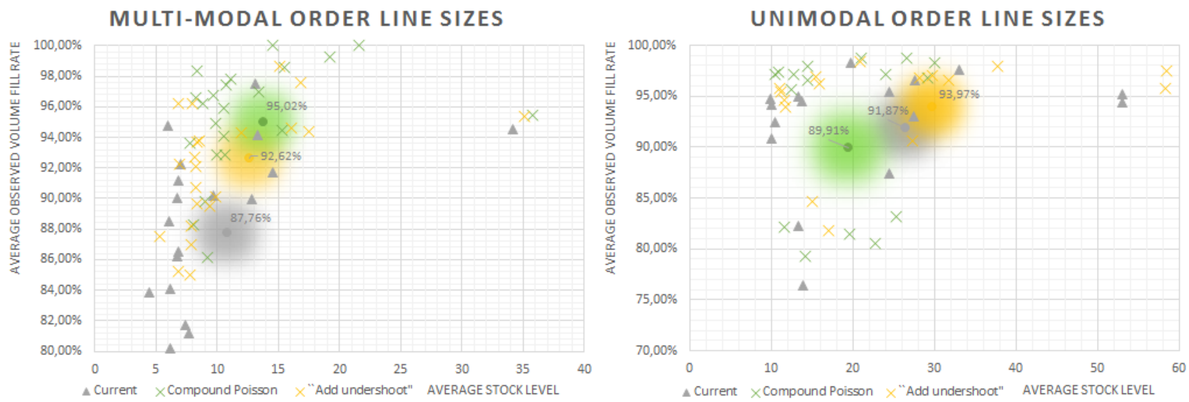

Figure 5.5: *The 22 multi-modal and 16 unimodal order line size SKUs and their performance for the two best-performing methods and the standard approach. The clouds represent the centre of gravity of the performance of each method.*

5.8 Conclusions

We now answer the sub research questions, and the main research question for our experiment:

a. *Which data are needed and how do we gather them?*
We need the following information about SKU-shop combinations: SKU code, shop identifier, replenishment minimum order quantity, replenishment incremental order quantity, target volume fill rate. Next, we need point-of-sales data, also taking place on SKU-shop level, which is gathered through digital shopping receipts: SKU code, shop identifier, sold quantity, and transaction date. Data is gathered through connections with the *SQL-database* of the companies of research.

b. *Which of the methods and heuristics for generating ordering policies performs best?*
We found clear distinctions between the two types of SKUs for most heuristics. Multi-modal order line size SKUs' inventories can at best be managed by using a target fill rate compound Poisson distribution. When doing so, the order line size distribution is incorporated as compounding distribution. This method provides an average observed volume fill rate that is 8.31%-point higher than currently, with 47.12% more stock on hand. For the other type of SKUs, the unimodal order line size SKUs, we found the "Add undershoot"-method to perform best. Hence, by incorporating the expected undershoot due to the order line size distribution and the review frequency, we can increase the observed volume fill rate by 1.98%-point, with 13.16% more stock on hand. Yet, since we can also adjust the standard approach to have extremely high target fill rates of 99%, which yields very similar results. So, we would not recommend to implement this method yet, since it yields no significant improvement compared to "stretching" the current way of working.

Finally, we conclude this chapter by answering the main research question of the experiment:

Which heuristic performs best on both the products classified and not classified as 'multi-modal order line size SKUs'?

As discussed when answering this chapter's research question [c.], there is no heuristic that performs best on both classifications. But as it is very well possible to clearly classify SKUs through our classification method, there is no need any more for a “one size fits all’ method for both types of SKUs. The SKUs of which the order line size distribution is classified as unimodal, ought to stay at the standard methods using a target fill rate constraint, and in addition they should incorporate the undershoot effect, as pilot for further testing. For the other class, the multi-modal order line size SKUs, the compound Poisson distribution should be used for modelling demand and handling inventory.

Conclusions and recommendations

We conclude our research by briefly summing up the conclusions on the research questions. Next, we formulate recommendations on implementing our solution.

Conclusions

1. How can the inventory management at best incorporate multi-modal order line demand?
Slim4 works with demand classes that take into account both the total demand during a period of time, and the customer arrival rate, no information on order line size is included. Yet, a large number of SKUs have such a specific order line distribution, hence it would be interesting to include this element into the inventory model. Currently, 67.61% of the multi-modal order line size SKUs from our dataset reach their target fill rate (for lost sales, for back-ordering this amounts 67.61% as well). The other type of SKU performs slightly worse on average (though there are three SKUs responsible for a big drop of average volume fill rate), with 64.15% and 64.52% of the SKUs meeting their target volume fill rate.

Hence, breaking the standard thinking of demand as total demand per period of time into a customer arrival rate and an order line size component seems interesting. By incorporating new inventory policies we attempt to serve more demand, as the stock on hand should never drop to a useless level: the amount of items that no customer wants to buy (for example three mugs). We found that the most issues with standard methods often arise in SKUs that are ordered at the supplier in small batches (for company X in at most packs of six), have high median order line sizes (for company X on average approximately 2.5), are slow-moving (for company X with not more than 30 customer orders per year), have small probabilities of selling less than 1 or 2 item(s) per order line, and in general have slightly lower target fill rates.

2. What can we learn from literature about optimization of multi-modal order line size SKUs inventory?
Sales can be transformed into true demand by enriching data with theoretical distributions, or through an empirical distribution. However, this is not preferable, as we do not have access to daily inventory level data, nor is it valid in our setting to only accommodate for larger order line sizes and not on other unobserved lost sales. So, we assume demand to equal sales. The demand process can at best be modelled via a compound Poisson process, therein focussing on the order line size distribution. We prefer to develop our own method for classifying SKUs based on their empirical distribution.

To the best of our knowledge there is hardly any research done on identifying modes within order line size distribution, except for mixture models. As we prefer the empirical distribution over a mixture of theoretical distributions, and for mixing the empirical distribution would make no sense, we feel the urge to develop our own method for classifying SKUs on their order line size distribution. Thus, we answer research question 2a by stating that we prefer to develop our own method for classifying SKUs based on their empirical distribution. The most important heuristics from literature for incorporating the 'multi-modal order line size effect' from SKUs classified as such, focus on incorporating the order line size distribution into safety stock and reorder points. Safety stock determinations are divided into either a direct service level method, an approximation, Poisson-based method, including undershoot, and a base stock level approximative method. The reorder point calculation methods include a distribution-free method.

3. How can the inventory management at best incorporate multi-modal order line demand?
We have obtained three heuristics: the "Overwrite"-, "Max"-, and "First moment β^*"-method, for directly incorporating the empirical order line size distribution. These heuristics are based on the expected undershoot, or a target fill rate altered first moment. Second, we incorporate for example a compound Poisson method from literature, that indirectly comprises the empirical order line size distribution. Next, we measure performance mainly in terms of volume fill rate and average stock on hand, followed by the observed order line fill rate. Yet, we cannot determine which method performs best a priori, so we need to test the methods.

4. Which heuristic performs best on both the products classified and not classified as 'multi-modal order line size SKUs'?
We have derived our own method of classifying, which we based on (and validated with) earlier methods with an expert panel. This leads to the classification of mainly SKU groups like cutlery and dining room equipment as they show a multi-modal order line size distribution. We have abolished our heuristic for setting the reorder point, and we have narrowed down to heuristics for safety stocks. Distinctions between the two types of SKUs are clear for most heuristics. Multi-modal order line size SKUs' inventories can at best be managed by using a target fill rate compound Poisson distribution. When doing so, the order line size distribution is incorporated as compounding distribution. This method provides an average observed volume fill rate that is 8.31%-point higher than currently, with 47.12% more stock on hand. When back-ordering is assumed, these values are 8.42%-point and 47.61%. For the other type of SKUs, the unimodal order line size SKUs, we found the "Add undershoot"-method to perform best. Hence, by incorporating the expected undershoot due to the order line size distribution and the review frequency, we can increase the observed volume fill rate by 1.98%-point, with 13.16% more stock on hand. Again, for back-ordering, these values amount 1.76%-point and 13.08%. However, if we set the target fill rates equal to 99% in the standard model, we obtain similar observed fill rates, making it uncertain whether this "Add undershoot" truly is better than the standard approach. There is no heuristic that performs best on both classifications. But as we clearly classify the SKUs, there is no need any more for a 'one size fits all' method for both types of SKUs.

Recommendations

We propose recommendations for implementation of our solution for Slimstock. Where possible we provide practical and managerial implications.

First, the classification scheme as we have developed in Section 5.1 should be implemented as such. The standard exact configuration of looking at the five most occurring order line sizes could be extended for SKUs being sold in many distinct order line sizes. In addition, the threshold should be re-established when including more modes. Henceforth, the first results for other retail clients' data should be useful for further verification and fine-tuning of the scheme. As two classes are comprises we recommended to add one boolean attribute to each SKU, with 0 for unimodal order line size SKUs, and 1 for multi-modal order line size SKUs. Classification itself is already implemented in Microsoft Excel, but it can rather easily be done in the SQL-databases as well.

Next, SKUs with multi-modal demand should be modelled through a compound Poisson distribution. This distribution incorporates the specific order lines, and works well for SKUs that do not sell often. Consequently, inventory should be held based on a target fill rate constraint, just like the standard situation. Currently, this method involves a cumbersome calculation implemented in Microsoft Excel. We propose to stick to these calculations, and implement them in Slim4, as transferring from and to Excel is no option. No extra data is required, since we remain at target fill rates and data on order line sizes is already present.

Furthermore, we do not propose to implement the “Add undershoot”-method for unimodal order line size SKUs yet. Although in this way the undershoot is correctly considered, thereby leading to less occurrences of ‘useless inventory’, and increasing the volume fill rate, the difference with the standard model is not significant. By simply increasing the target fill rates to 99% we obtain similar results, which diminishes the added value of a new, to be implemented, heuristic. Since our focus was not on this type of SKU, more practical research should be conducted on improving this method, since we acknowledge its theoretical added value. We ran our experiment with constant review periods, which impacts the undershoot review frequency component. Its sensitivity for other lengths of review periods should be found, before implementing this component of the undershoot. Next, it would be interesting to start with adding two attributes, namely the expected undershoot due to the order line size distribution and its variance. In the meantime, we advice not to change the safety stock determination for unimodal order line size SKUs.

Finally, our recommendations touch upon the core of an inventory model, namely demand modelling and replenishment order advice generation. Therefore, it is of utmost importance that a long-term pilot is executed. During this period the impact of aspects such as trends, seasonality, and promotion, on the “Add undershoot”- and compound Poisson method should be researched. We propose to execute a pilot among both existing and new clients, and to run it parallel to the standard Slim4 implementation. The volume fill rate and average stock level for multi-modal order line size SKUs (assuming sales are lost in case of insufficient stock on hand), should on average be around 8%-point higher and 45-50% higher for our solution. For unimodal order line size SKUs this range should be 2%-point and 13%.

Discussion and future research

Now that we have completed our research it is time to critically reflect on the project and to provide directions for future research. We present two main topics for discussion and eight directions for future research.

Discussion

First of all we have tested the different methods within our inventory model with empiric data. Data was limited, as we had six months of transactions. However, more history would simply mean more accuracy in results. This could also have been accomplished by implementing a theoretical probability distribution for the demand process, as long-term simulation is then a possibility. Another possibility was to abolish the SKU-shop combinations having less than a threshold transactions within the six months. Namely, small numbers of transactions led to trivial performances. For example, if two out of four transactions coincidentally have rarely-occurring sizes and insufficient stock levels, fill rates are only 50%. On the other hand, if only one transaction takes place, which is successful, the fill rate equals 100%. Yet, performance is only based on one successful instance.

Next, we had some difficulties in programming our Excel spreadsheet in the right way to quickly run through the experiment. One could imagine that loading 64,393 transactions for 2,888 SKU-shop combinations, running through 12 configurations for controlling inventory, and measuring performance, gets quite cumbersome in Excel. Yet, due the limited period of time, and the lack of a good alternative, programming in Excel was the best option. Nevertheless, we are confident that Slimstock can follow-up our recommendations in their software, which is not based on Excel nor VBA-code.

Future research

Next, there have been some topics we stumbled upon during our project that are worth researching in the future. First of all, it would be interesting to gather data on the true demand distribution of the multi-modal order line size SKUs. It is a possibility to consult customers either during the point of sale (e.g.: "Was this the right amount of items, or did you want to buy more items of this SKU or any other SKU?"), or when customers leave the shop empty-handed (e.g.: "Did you observe insufficient inventory of the SKU of your interest?"). This would greatly contribute to insights on a 'truly expected' order line size, and how much inventory should be minimally kept (the safety stock).

Furthermore, in accordance with the aforementioned direction for future research, it would also be possible to transform historic sales into demand in a more theoretic way. If correct censor rates could be obtained, the censored data (i.e.: lost sales are unobservable in the transaction data as they are not registered) could be directly used if transformed correctly. We suggest that both small order line sizes (customers only wanting one single mug tend to leave the shops without wanting to back-order) and big order line sizes are censored.

A retailer should think about the balance between stock level and (volume) fill rate. As we have seen, multi-modal order line size SKUs can show higher volume fill rates through different heuristics. Roughly speaking, these heuristics yield different service levels against different increased stock levels. The degree to which a retailer would want to increase the SKU's stock level, should be further researched, as this would effect the definition of 'best' heuristic.

We chose to leave the mixture modelling topic for modelling customer demand. We still believe it would be an interesting way for being able to find a theoretical distribution for demand. Yet, methods should be developed to find a method to adequately fit each SKU to a mixture distributions automatically. Furthermore, implications for the inventory model when using for example a mixture of negative binomial distributions for the order line size demand distribution should be researched. Hence, which assumptions for calculating safety stock are still valid, or perhaps new methods for calculating safety stock should be derived.

Next, in our research we assumed that the order line size distribution would not be impacted by trends or seasons. More research however would be needed to find out if this order line size distribution indeed does not differ significantly. Furthermore, the (long-term) effects of promotion on the order line size distribution is an interesting extension for our findings. In this same line of reasoning, we hypothesize that attributes like the physical location (in the city centre or outside the city with big parking lots) influences the order line size distribution. Therefore, this could be taken into account when constructing an inventory policy for this location.

Furthermore, we took one of the aspects of an inventory model very static: the moment of ordering and replenishment. Namely, both the review period and lead time were set deterministically, and both are equal to one week. This however influences the inventory model, and it would be interesting to research what the influence of releasing this assumption would mean for our heuristics for incorporating the order line size distribution.

It could be interesting to evaluate other values for the BO-factor as well, e.g. values 1.1 (everything right above the average order line size is back-ordered), and 7 (in this setting the BO threshold already reach 'business-to-business order lines' of for example 21 plates). This range of states for parameter BO would provide interesting insights on the impact of this conditional back-ordering setting. This contributes to the effectiveness of the conditional back-ordering setting, which is currently not very effective.

Finally, more research can be done in developing a decision support system for environments in which customers cannot pick their items from shelf but are helped by the shop instead. In such settings, it would be interesting to help decisions like serving a big order line, given a certain order line size distribution, the standard stock on hand, and the remaining expected demand until replenishment.

References

Agrawal, N., & Smith, S. A. (1996). Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Research Logistics (NRL)*, *43*(6), 839-861.

Akaike, H. (1973). Information Theory and an Extension of the Maximum Likelihood Principle. , 199–213.

Axsäter, S. (2006). *Inventory Control* (2nd ed.). Springer-Verlag New York Inc.

Bijvank, M., & Vis, I. (2011). Lost-sales inventory theory: A review. *European Journal of Operational Research*, *215*(1), 1 - 13.

Bisi, A., Dada, M., & Tokdar, S. (2011). A Censored-Data Multiperiod Inventory Problem with Newsvendor Demand Distributions. *Manufacturing & Service Operations Management*, *13*(4), 525-533.

Bratley, P., Fox, B. L., & Schrage, L. E. (1986). *A Guide to Simulation (2nd ed.).* New York, USA: Springer-Verlag New York, Inc.

Breuer, L. (2002). An EM Algorithm for Batch Markovian Arrival Processes and its Comparison to a Simpler Estimation Procedure. *Annals of Operations Research*, *112*(1), 123–138.

Breuer, L., & Kume, A. (2010). An EM Algorithm for Markovian Arrival Processes Observed at Discrete Times. In B. Muller-Clostermann, K. Echtle, & E. P. Rathgeb (Eds.), *Measurement, Modelling, and Evaluation of Computing Systems and Dependability and Fault Tolerance: 15th International GI/ITG Conference, MMB and DFT 2010, Essen, Germany, March 15-17, 2010. Proceedings* (pp. 242–258). Berlin, Heidelberg: Springer Berlin Heidelberg.

Chen, J., & Kalbfleisch, J. (1996). Penalized minimum-distance estimates in finite mixture models. *Canadian Journal of Statistics*, *24*(2), 167–175.

Chen, J., & Khalili, A. (2009). Order Selection in Finite Mixture Models With a Nonsmooth Penalty. *Journal of the American Statistical Association*, *104*(485), 187-196.

Cook, J. (2017). *Notes on the Negative Binomial Distribution.* Retrieved from `https://www.johndcook.com`

Cordeiro, J. D., & Kharoufeh, J. P. (2010). Batch Markovian Arrival Processes (BMAP). In *Wiley Encyclopedia of Operations Research and Management Science.* John Wiley and Sons, Inc.

Croston, J. D. (1972). Forecasting and Stock Control for Intermittent Demands. *Operational Research Quarterly (1970-1977)*, *23*(3), 289-303.

Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, *39*(1), 1–38.

Ehrhardt, R. (1979). The Power Approximation for Computing (s, S) Inventory Policies. *Management Science*, *25*(8), 777-786.

Fan, J., & Li, R. (2001). Variable Selection via Nonconcave Penalized Likelihood and Its Oracle Properties. *Journal of the American Statistical Association*, *96*(456), 1348-1360.

Figueiredo, M. A. T., & Jain, A. K. (2002). Unsupervised learning of finite mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *24*(3), 381-396.

Friend, J. K. (1960). Stock Control with Random Opportunities for Replenishment. *Journal of the Operational Research Society*, *11*(3), 130–136.

Hadley, G., & Whitin, T. (1963). *Analysis of inventory systems*. Prentice-Hall.

Hill, R. M. (1988). Stock control and the undershoot of the re-order level. *The Journal of the Operational Research Society*, *39*(2), 173-181.

Hu, M., Pavlicova, M., & Nunes, E. V. (2011). Zero-inflated and Hurdle Models of Count Data with Extra Zeros: Examples from an HIV-Risk Reduction Intervention Trial. *The American Journal of Drug and Alcohol Abuse*, *37*(5), 367 - 375.

Johnson, N., Kemp, A., & Kotz, S. (2005). *Univariate Discrete Distributions*. Wiley.

Johnston, F. R., Boylan, J. E., & Shale, E. A. (2003). An examination of the size of orders from customers, their characterisation and the implications for inventory control of slow moving items. *Journal of the Operational Research Society*, *54*(8), 833–837.

Larsen, C., & Thorstenson, A. (2008). A Comparison between the Order and the Volume Fill Rate for a Base-Stock Inventory Control System under a Compound Renewal Demand Process. *The Journal of the Operational Research Society*, *59*(6), 798-804.

Lau, H.-S., & Lau, A. H.-L. (1996). Estimating the demand distributions of single-period items having frequent stockouts. *European Journal of Operational Research*, *92*(2), 254 - 265.

Law, A. M. (2015). *Simulation Modeling and Analysis* (5th ed.). McGraw-Hill Higher Education.

Lindsay, B. (1995). *Mixture Models: Theory, Geometry, and Applications*. Institute of Mathematical Statistics.

Lucantoni, D. M. (1991). New results on the single server queue with a Batch Markovian Arrival Process. *Communications in Statistics. Stochastic Models*, *7*(1), 1-46.

Mukhopadhyay, S. (2017). Large-scale mode identification and data-driven sciences. *Electronic Journal of Statistics*, *11*(1), 215–240.

Nahmias, S. (1993). *Production and Operations Analysis*. Irwin.

Nenes, G., Panagiotidou, S., & Tagaras, G. (2010). Inventory management of multiple items with irregular demand: A case study. *European Journal of Operational Research*, *205*(2), 313 - 324.

Richardson, S., & Green, P. J. (1997). On Bayesian Analysis of Mixtures with an Unknown Number of Components (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *59*(4), 731–792.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*, 461–464.

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2017). *Inventory and Production Management in Supply Chains* (4th ed.). CRC Press.

Slim4-developer. (2016a). Internal meeting.

Slim4-developer. (2016b). Internal meeting.

Slim4-developer. (2016c). Internal meeting.

Tijms, H., & Groenevelt, H. (1984). Simple approximations for the reorder point in periodic and continuous review (s, S) inventory systems with service level constraints. *European Journal of Operational*

Research, *17*(2), 175 - 190.

Tsybakov, A. (2008). *Introduction to Nonparametric Estimation.* Springer New York.

van Donselaar, K. H., & Broekmeulen, R. A. (2013). Determination of safety stocks in a lost sales inventory system with periodic review, positive lead-time, lot-sizing and a target fill rate. *International Journal of Production Economics*, *143*(2), 440 - 448.

Vereecke, A., & Verstraeten, P. (1994). An inventory management model for an inventory consisting of lumpy items, slow movers and fast movers. *International Journal of Production Economics*, *35*(1), 379 - 389.

Willemain, T. R., Smart, C. N., Shockor, J. H., & DeSautels, P. A. (1994). Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston's method. *International Journal of Forecasting*, *10*(4), 529 - 538.

Williams, B. D., & Tokar, T. (2008). A review of inventory management research in major logistics journals: Themes and future directions. *The International Journal of Logistics Management*, *19*(2), 212-232.

Winston, W. L. (2003). *Operations Research: Applications and Algorithms* (4th ed.). Cengage Learning, Inc.

Xu, C., & Chen, J. (2015). A Thresholding Algorithm for Order Selection in Finite Mixture Models. *Communications in Statistics - Simulation and Computation*, *44*(2), 433-453.

Appendices

A Analysis problem size

First we investigate the slow-movers, by postulating that if SKUs have a multi-modal order demand pattern, then their most frequently sold order line sizes should have a reasonable big discrepancy between their median and mean. To this end we investigated the six most frequently sold order line sizes per SKU, in order to get a representative view on the top-selling order line sizes. As metric for dispersion we needed a ratio variable without unit, in order to compare the SKUs of varying order line sizes. We took the ratio of median divided by mean as measure, since a value close to 1 would suggest a smoothly distributed order pattern. Henceforth, SKUs having a median/mean-ratio substantially lower or higher than 1, were classified as 'interesting for further research'. All SKUs were visually checked through figures like Figure A1, acting as confirmation of our classification.

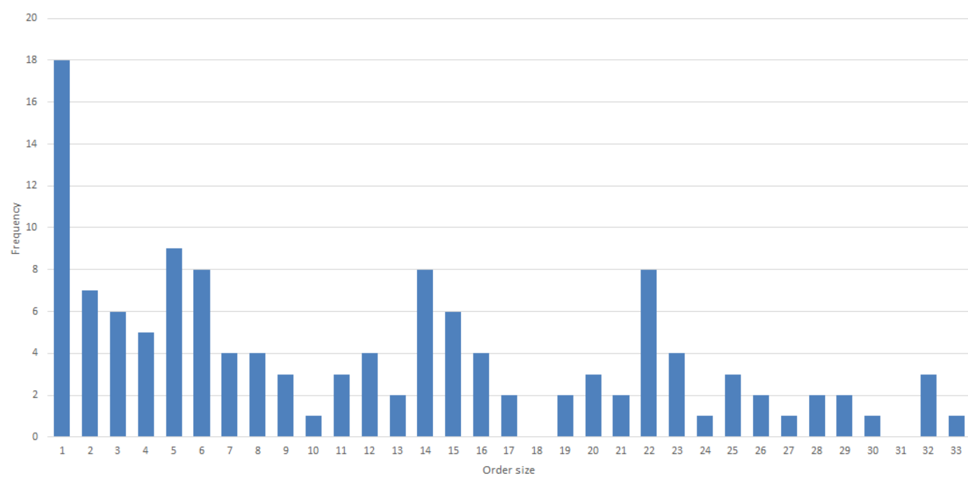

Figure A1: *Order line size figure of slow-moving SKU with non-smoothly distributed order line sizes, clearly multi-modal*

Secondly, we investigated the fast-movers. When looking at the figures of the fast-movers, roughly two subcategories were seen by a panel of experts (my external supervisor and his colleague), namely SKUs whose order line sizes resembled a (combination of) Poisson-distributed pattern(s) (see Figure A2 for a presumably uni-modal Poisson-mixture example) and patterns that resembled an approximative straight line (see Figure A3 for presumably multi-modal example due to big discrepancies with a linear trend). This led to a categorization based on the skewness of all the sales data of the SKUs (not just its top-6 order line sizes). In agreement with this panel of experts, we found that all SKUs skewed 1,75 or less -

so, hardly skewed - roughly resemble a straight line, and on the other hand all SKUs skewed more than 1,75 roughly resemble a Poisson distribution. This value of 1,75 was empirically found by looking at the figures, and all SKUs with a skewness close to 1,75 were manually classified. Finally, all SKUs from these two subcategories were manually reviewed by the panel of experts, and when either a big discrepancy with the trend line (in case of the 'straight line'-category) or a graphically well-visible second (and consequently third, or even fourth) maximum (in case of the Poisson distribution) was seen, we classified the SKU as 'interesting for further research'.

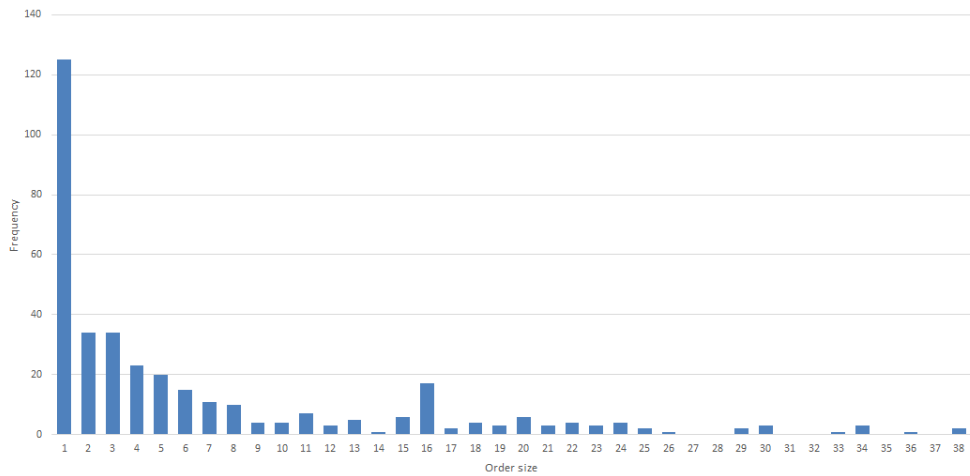

Figure A2: *Order line size figure of fast-moving SKU, presumably a simple mixture distribution*

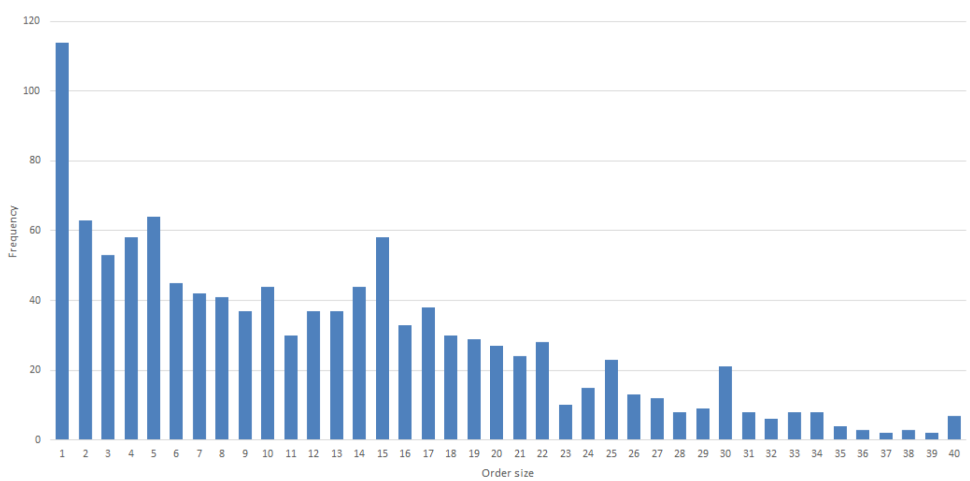

Figure A3: *Order line size figure of fast-moving SKU, presumably approximately monotonically decreasing*

B VBA code for experiment

VBA code as implemented in Chapter 5 is found in Figure B1.

```
Sub Processing()
'definition
Dim cs As Worksheet
Dim res As Worksheet
Dim art As Worksheet
Set cs = Sheets("Controlsheet")
Set res = Sheets("Result")
Set art = Sheets("ArticleInfo_LOOKUP")
Dim ac As Range
Dim artcode As Range
Dim oos As Range
Dim outofstock As Range
Dim hr As Range
Dim heuristic As Range
Dim rc As Integer
'scenarios
Set artcode = art.Range("B2:B2888")
Set outofstock = cs.Range("A42:A44")
Set heuristic = cs.Range("A50:A56")
'disable visuals, to increase speed
Application.ScreenUpdating = False
'clean result sheet
res.Range("A2:AB1000000").ClearContents
'initialization of counter
rc = 2
'enter three for-loops, in sequence of the scenarios
For Each ac In artcode
cs.Range("B37") = ac
    For Each oos In outofstock
    cs.Range("B38") = oos
        For Each hr In heuristic
        cs.Range("B39") = hr
        'manually calculate, to improve speed
        Application.Calculate
        If Not Application.CalculationState = xlDone Then
        DoEvents
        End If

        'write down results
        res.Range("A" & rc) = ac.Value
        res.Range("B" & rc) = bl.Value
        res.Range("C" & rc) = ss.Value
        res.Range("D" & rc) = cs.Range("B27").Value
        'etc. etc.
        rc = rc + 1
        Next
    Next
Next
Application.ScreenUpdating = True
MsgBox ("Experiment finished")
End Sub
```

Figure B1: *VBA code as implemented in our experiment*

C Numerical results

SKU	**Category**	**MOQ/IOQ**	**Median X**	**#order lines**	**Total demand (items)**	β^*	$\mathbf{P}(X \leq 1)$	$\mathbf{P}(X \leq 2)$
MM SKU 1	Kitchen equipment	12	4	6.73	3,062	98.03%	12.40%	39.53%
MM SKU 2	Cutlery	12	3	2.68	506	96.17%	24.75%	56.44%
MM SKU 3	Cutlery	12	5	1.78	707	96.94%	19.63%	34.58%
MM SKU 4	Kitchen equipment	6	2	7.06	1,938	97.77%	18.85%	44.26%
MM SKU 5	Kitchen equipment	6	1	5.79	1,185	96.98%	24.14%	50.00%
MM SKU 6	Kitchen equipment	6	2	3.63	713	95.96%	29.13%	70.87%
MM SKU 7	Kitchen equipment	6	2	4.83	1,259	97.04%	27.06%	64.71%
MM SKU 8	Kitchen equipment	6	2	3.62	1,066	97.12%	18.10%	48.57%
MM SKU 9	Kitchen equipment	36	2	9.99	2,703	97.94%	27.38%	60.71%
MM SKU 10	Kitchen equipment	6	2	3.58	1,064	96.78%	20.00%	40.00%
MM SKU 11	Kitchen equipment	12	5	3.61	2,011	97.87%	16.21%	53.12%
MM SKU 12	Kitchen equipment	4	3	3.60	1,480	97.48%	28.21%	48.72%
MM SKU 13	Kitchen equipment	6	2	2.90	855	96.59%	30.19%	79.25%
MM SKU 14	Kitchen equipment	6	2	3.74	1,127	96.93%	10.00%	50.00%
MM SKU 15	Kitchen equipment	6	2	3.40	1,095	96.62%	22.22%	66.67%
MM SKU 16	Kitchen equipment	6	2	5.75	1,808	98.17%	12.50%	25.00%
MM SKU 17	Kitchen equipment	6	2	4.95	1,241	97.35%	34.43%	69.81%
MM SKU 18	Kitchen equipment	6	4	3.56	1,450	97.70%	55.61%	86.22%
MM SKU 19	Kitchen equipment	6	2	12.97	4,843	98.78%	48.28%	85.06%
MM SKU 20	Kitchen equipment	6	2	3.70	1,036	97.22%	30.33%	71.31%
MM SKU 21	Kitchen equipment	6	2	3.15	529	97.48%	28.28%	60.61%
MM SKU 22	Kitchen equipment	4	2	2.21	432	97.19%	40.00%	64.00%
Average multi-modal order line size SKUs		8.27	2.50	4.69	1,460	97.28%	26.26%	57.70%
UM SKU 23	Others	36	1	17.26	2,399	98.66%	48.76%	75.48%
UM SKU 24	Others	12	2	75.30	9,238	99.00%	46.45%	79.04%
UM SKU 25	Others	12	1	10.87	2,478	98.33%	52.07%	80.18%
UM SKU 26	Others	12	1	53.32	8,618	98.59%	58.07%	83.26%
UM UM SKU 27	Others	30	1	40.62	7,212	98.98%	56.05%	83.77%
UM SKU 28	Others	12	1	35.48	4,791	98.96%	87.24%	95.81%
UM SKU 29	Others	24	1	11.80	2,466	97.61%	44.92%	73.31%
UM SKU 30	Others	24	1	58.96	11,329	98.07%	52.01%	79.93%
UM SKU 31	Others	12	1	64.72	2,714	98.88%	80.08%	93.68%
UM SKU 32	Others	12	1	22.68	3,622	98.78%	70.19%	89.72%
UM SKU 33	Others	12	1	42.44	10,243	97.54%	52.11%	79.83%
UM SKU 34	Others	12	1	40.98	9,121	97.95%	49.15%	76.06%
UM SKU 35	Others	10	3	3.30	976	96.57%	27.78%	52.78%
UM SKU 36	Others	12	1	50.73	7,911	98.97%	73.64%	92.99%
UM SKU 37	Others	12	2	24.34	6,085	98.94%	41.58%	82.46%
UM SKU 38	Others	30	1	13.33	2,623	98.23%	61.18%	84.65%
Average unimodal order line size SKUs		17.13	1.25	35.38	5,739	98.38%	56.33%	81.43%
Differences between averages		-51.69%	45.45%	-86.74%	-1.12%	-74.57%	41.29%	69.57%
Average - total		12.70	2.45	20.04	3,599	97.83%	-53.38%	-29.14%

Table C1: *SKUs of research with their static data*

	Lost sales			*Conditional back-ordering*			*Full back-ordering*		
SKU	**SoH**	$\boldsymbol{\beta}$	$\boldsymbol{\gamma}$	**SoH**	$\boldsymbol{\beta}$	$\boldsymbol{\gamma}$	**SoH**	$\boldsymbol{\beta}$	$\boldsymbol{\gamma}$
MM SKU 1	14.53	91.71%	94.71%	14.85	95.54%	96.13%	14.52	91.80%	94.87%
MM SKU 2	13.10	97.51%	98.88%	13.11	97.51%	98.88%	13.13	97.51%	98.88%
MM SKU 3	13.35	94.12%	95.38%	13.33	94.12%	95.38%	13.39	94.12%	95.38%
MM SKU 4	6.73	86.22%	90.13%	6.71	87.49%	90.08%	6.86	86.83%	91.47%
MM SKU 5	6.11	88.52%	92.81%	6.09	88.39%	92.47%	6.02	87.88%	92.56%
MM SKU 6	6.03	94.77%	95.47%	6.05	94.85%	95.47%	6.08	94.53%	95.35%
MM SKU 7	6.72	90.04%	93.18%	6.70	91.33%	93.19%	6.94	89.96%	92.96%
MM SKU 8	6.83	86.52%	92.60%	6.75	86.76%	92.60%	6.75	87.40%	93.06%
MM SKU 9	34.15	94.57%	95.98%	34.29	95.94%	96.46%	34.33	95.63%	97.00%
MM SKU 10	7.01	92.26%	94.58%	7.01	93.24%	94.58%	6.98	90.65%	93.76%
MM SKU 11	12.83	89.93%	91.46%	12.68	90.52%	90.93%	12.81	90.76%	92.48%
MM SKU 12	6.19	84.05%	88.24%	6.19	85.34%	87.85%	6.22	84.55%	89.31%
MM SKU 13	6.67	79.10%	85.51%	6.68	80.29%	85.69%	6.73	79.21%	86.14%
MM SKU 14	6.80	91.19%	93.02%	6.77	91.39%	93.02%	6.61	90.90%	92.94%
MM SKU 15	6.17	80.21%	87.83%	6.19	82.71%	88.15%	6.17	78.64%	86.77%
MM SKU 16	7.42	81.72%	88.95%	7.41	81.75%	89.21%	7.44	81.18%	89.01%
MM SKU 17	6.41	79.14%	84.59%	6.38	80.93%	84.30%	6.50	80.77%	86.35%
MM SKU 18	7.67	81.16%	87.17%	7.66	81.16%	87.17%	7.54	79.11%	85.65%
MM SKU 19	9.72	90.18%	93.07%	9.71	92.02%	93.19%	9.72	89.66%	92.67%
MM SKU 20	5.92	77.68%	82.92%	5.90	81.76%	82.25%	5.97	78.68%	84.04%
MM SKU 21	6.36	78.99%	85.97%	6.31	80.18%	85.35%	6.34	79.73%	86.29%
MM SKU 22	4.48	83.89%	87.79%	4.37	83.55%	87.30%	4.30	82.82%	87.90%
Avg. MM	9.15	86.98%	90.92%	9.14	88.04%	90.89%	9.15	86.92%	91.13%
UM SKU 23	27.48	93.02%	96.47%	26.81	96.18%	96.89%	27.86	94.56%	97.05%
UM SKU 24	13.84	76.45%	81.52%	13.89	89.55%	82.04%	14.12	77.42%	81.90%
UM SKU 25	10.02	94.11%	95.68%	10.00	95.93%	95.22%	10.01	93.87%	95.78%
UM SKU 26	33.03	97.55%	98.20%	32.98	98.69%	98.22%	33.07	97.63%	98.24%
UM SKU 27	24.39	95.47%	97.62%	24.37	96.35%	97.66%	24.39	96.09%	97.79%
UM SKU 28	10.42	92.45%	95.39%	10.42	93.59%	95.40%	10.41	92.42%	95.27%
UM SKU 29	27.57	96.60%	97.49%	27.49	97.66%	97.63%	27.41	96.49%	97.55%
UM SKU 30	24.46	87.42%	89.21%	24.44	93.51%	89.73%	24.51	87.52%	89.37%
UM SKU 31	13.27	82.22%	84.61%	13.14	91.40%	84.81%	13.20	82.71%	85.45%
UM SKU 32	9.89	94.68%	96.78%	9.88	95.96%	96.57%	9.90	94.64%	96.89%
UM SKU 33	53.04	94.34%	94.71%	53.12	97.54%	94.95%	53.54	94.87%	95.26%
UM SKU 34	52.98	95.19%	96.54%	53.10	97.41%	96.62%	53.10	95.08%	96.63%
UM SKU 35	9.97	90.89%	93.35%	9.97	91.84%	93.35%	9.97	91.75%	94.06%
UM SKU 36	13.65	94.54%	95.86%	13.62	97.31%	95.70%	13.65	94.75%	95.96%
UM SKU 37	13.24	94.92%	96.90%	13.19	96.47%	97.04%	13.29	95.20%	97.13%
UM SKU 38	19.69	98.29%	98.79%	19.66	98.44%	98.90%	19.61	98.19%	98.79%
Avg. UM	22.31	92.38%	94.32%	22.25	95.49%	94.42%	22.38	92.70%	94.57%
Difference	-59.01%	-5.85%	-3.60%	-58.92%	-7.81%	-3.74%	-59.10%	-6.23%	-3.64%
Avg. total	15.73	89.68%	92.62%	15.70	91.76%	92.66%	15.76	89.81%	92.85%

Table C2: *Comparison performance standard method for all three out-of-stock configurations*

	Lost sales					*Full back-ordering*				
SKU	**SoH**	$\boldsymbol{\beta}$	$\boldsymbol{\gamma}$	**ss**	**s**	**SoH**	$\boldsymbol{\beta}$	$\boldsymbol{\gamma}$	**ss**	**s**
MM SKU 1	14.53	91.71%	94.71%	5.21	9.50	14.52	91.80%	94.87%	5.18	9.41
MM SKU 2	13.10	97.51%	98.88%	0.78	2.27	13.13	97.51%	98.88%	0.79	2.29
MM SKU 3	13.35	94.12%	95.38%	0.75	1.83	13.39	94.12%	95.38%	0.76	1.83
MM SKU 4	6.73	86.22%	90.13%	2.43	4.53	6.86	86.83%	91.47%	2.46	4.86
MM SKU 5	6.11	88.52%	92.81%	1.49	3.36	6.02	87.88%	92.56%	1.51	3.40
MM SKU 6	6.03	94.77%	95.47%	0.53	1.81	6.08	94.53%	95.35%	0.56	1.86
MM SKU 7	6.72	90.04%	93.18%	1.90	3.81	6.94	89.96%	92.96%	1.93	3.83
MM SKU 8	6.83	86.52%	92.60%	1.58	3.28	6.75	87.40%	93.06%	1.61	3.28
MM SKU 9	34.15	94.57%	95.98%	0.22	1.41	34.33	95.63%	97.00%	0.22	1.42
MM SKU 10	7.01	92.26%	94.58%	2.07	3.85	6.98	90.65%	93.76%	2.11	3.86
MM SKU 11	12.83	89.93%	91.46%	3.06	5.34	12.81	90.76%	92.48%	3.09	5.36
MM SKU 12	6.19	84.05%	88.24%	3.07	5.08	6.22	84.55%	89.31%	3.10	5.21
MM SKU 13	6.67	79.10%	85.51%	1.11	2.41	6.73	79.21%	86.14%	1.14	2.44
MM SKU 14	6.80	91.19%	93.02%	1.65	3.44	6.61	90.90%	92.94%	1.69	3.46
MM SKU 15	6.17	80.21%	87.83%	1.16	2.34	6.17	78.64%	86.77%	1.18	2.38
MM SKU 16	7.42	81.72%	88.95%	2.55	4.58	7.44	81.18%	89.01%	2.56	4.59
MM SKU 17	6.41	79.14%	84.59%	1.52	2.91	6.50	80.77%	86.35%	1.55	3.20
MM SKU 18	7.67	81.16%	87.17%	2.70	4.67	7.54	79.11%	85.65%	2.72	4.69
MM SKU 19	9.72	90.18%	93.07%	5.59	9.97	9.72	89.66%	92.67%	5.61	10.00
MM SKU 20	5.92	77.68%	82.92%	0.58	1.47	5.97	78.68%	84.04%	0.59	1.51
MM SKU 21	6.36	78.99%	85.97%	1.01	2.18	6.34	79.73%	86.29%	1.04	2.18
MM SKU 22	4.48	83.89%	87.79%	0.26	1.03	4.30	82.82%	87.90%	0.27	1.03
Avg. MM	9.15	86.98%	90.92%	1.87	3.68	9.15	86.92%	91.13%	1.89	3.73
UM SKU 23	27.48	93.02%	96.47%	6.35	13.03	27.86	94.56%	97.05%	6.37	13.05
UM SKU 24	13.84	76.45%	81.52%	8.31	18.47	14.12	77.42%	81.90%	8.33	18.53
UM SKU 25	10.02	94.11%	95.68%	2.44	5.23	10.01	93.87%	95.78%	2.45	5.28
UM SKU 26	33.03	97.55%	98.20%	17.38	33.93	33.07	97.63%	98.24%	17.47	34.04
UM SKU 27	24.39	95.47%	97.62%	6.53	16.27	24.39	96.09%	97.79%	6.55	16.32
UM SKU 28	10.42	92.45%	95.39%	3.34	7.96	10.41	92.42%	95.27%	3.35	7.98
UM SKU 29	27.57	96.60%	97.49%	1.48	4.56	27.41	96.49%	97.55%	1.50	4.56
UM SKU 30	24.46	87.42%	89.21%	9.35	18.87	24.51	87.52%	89.37%	9.42	18.94
UM SKU 31	13.27	82.22%	84.61%	6.08	12.84	13.20	82.71%	85.45%	6.10	12.84
UM SKU 32	9.89	94.68%	96.78%	2.86	6.62	9.90	94.64%	96.89%	2.87	6.63
UM SKU 33	53.04	94.34%	94.71%	38.33	55.01	53.54	94.87%	95.26%	38.56	55.30
UM SKU 34	52.98	95.19%	96.54%	38.14	55.28	53.10	95.08%	96.63%	38.31	55.44
UM SKU 35	9.97	90.89%	93.35%	0.76	2.23	9.97	91.75%	94.06%	0.77	2.24
UM SKU 36	13.65	94.54%	95.86%	5.48	13.60	13.65	94.75%	95.96%	5.49	13.61
UM SKU 37	13.24	94.92%	96.90%	5.39	11.28	13.29	95.20%	97.13%	5.41	11.28
UM SKU 38	19.69	98.29%	98.79%	1.40	4.33	19.61	98.19%	98.79%	1.41	4.36
Avg. UM	22.31	92.38%	94.32%	9.60	17.47	22.38	92.70%	94.57%	9.65	17.52
Difference	-59.01%	-5.85%	-3.60%	-80.47%	-78.91%	-59.10%	-6.23%	-3.64%	-80.36%	-78.71%
Avg. total	15.73	89.68%	92.62%	5.74	10.58	15.76	89.81%	92.85%	5.77	10.63

Table C3: *Performance of standard method*

	Lost sales					*Full back-ordering*				
SKU	**SoH**	**β**	**γ**	**ss**	**s**	**SoH**	**β**	**γ**	**ss**	**s**
MM SKU 1	17.48	94.42%	96.54%	8.21	12.35	17.62	94.04%	96.26%	8.25	12.37
MM SKU 2	15.15	98.67%	99.09%	2.54	4.12	15.15	98.67%	99.09%	2.56	4.12
MM SKU 3	16.84	97.60%	98.55%	4.22	5.44	16.89	97.60%	98.55%	4.23	5.44
MM SKU 4	8.19	92.71%	95.29%	3.85	6.15	8.18	92.73%	95.34%	3.87	6.17
MM SKU 5	6.95	92.24%	95.00%	2.40	4.20	6.93	92.06%	94.94%	2.42	4.23
MM SKU 6	6.85	96.19%	97.18%	1.34	2.59	6.80	95.95%	97.06%	1.37	2.59
MM SKU 7	8.32	93.71%	95.25%	3.32	5.07	8.35	94.46%	95.93%	3.35	5.08
MM SKU 8	8.29	92.09%	95.62%	3.15	4.80	8.29	92.72%	96.08%	3.18	4.80
MM SKU 9	35.14	95.37%	96.90%	1.50	2.48	35.36	96.08%	97.56%	1.51	2.48
MM SKU 10	8.54	93.77%	96.63%	3.81	5.70	8.53	93.24%	96.19%	3.85	5.74
MM SKU 11	16.08	94.59%	96.39%	6.29	8.83	16.13	94.96%	96.30%	6.31	8.83
MM SKU 12	8.30	90.74%	93.68%	5.18	7.32	8.24	90.25%	93.44%	5.21	7.32
MM SKU 13	8.41	89.69%	93.26%	2.83	4.19	8.43	89.08%	92.85%	2.86	4.19
MM SKU 14	8.00	96.24%	96.64%	3.16	4.74	7.89	95.95%	96.49%	3.20	4.76
MM SKU 15	7.88	86.95%	92.18%	2.73	4.30	7.88	86.22%	91.79%	2.75	4.32
MM SKU 16	9.42	89.51%	93.68%	4.66	6.67	9.45	88.73%	93.29%	4.68	6.67
MM SKU 17	7.85	88.21%	92.47%	2.86	4.35	7.81	88.32%	92.53%	2.90	4.40
MM SKU 18	9.90	90.14%	93.02%	5.12	6.99	9.83	89.69%	92.61%	5.14	6.99
MM SKU 19	11.96	94.33%	95.97%	8.08	12.36	11.96	94.42%	96.26%	8.10	12.36
MM SKU 20	6.87	85.19%	89.44%	1.53	2.52	6.83	84.77%	89.19%	1.54	2.52
MM SKU 21	7.76	85.03%	90.76%	2.53	3.68	7.77	85.30%	91.48%	2.55	3.68
MM SKU 22	5.32	87.55%	90.98%	1.26	2.03	5.15	87.46%	91.91%	1.26	2.03
Avg. MM	10.89	92.04%	94.75%	3.66	5.49	10.89	91.94%	94.78%	3.68	5.50
UM SKU 23	31.76	96.52%	98.31%	10.43	17.11	32.05	96.52%	98.31%	10.45	17.16
UM SKU 24	16.97	81.76%	85.47%	11.60	21.87	17.09	81.94%	85.47%	11.62	21.87
UM SKU 25	11.11	95.38%	96.82%	3.70	6.60	11.18	96.25%	97.32%	3.72	6.60
UM SKU 26	37.77	97.93%	98.39%	22.42	38.96	37.89	98.06%	98.44%	22.51	39.07
UM SKU 27	28.20	96.80%	98.47%	10.14	19.93	28.27	97.55%	98.73%	10.17	19.95
UM SKU 28	11.69	93.96%	96.53%	4.73	9.26	11.65	93.63%	96.21%	4.75	9.26
UM SKU 29	29.75	97.01%	98.04%	3.91	6.78	29.65	96.86%	97.99%	3.92	6.80
UM SKU 30	27.26	90.62%	92.03%	12.70	22.13	27.30	90.80%	92.12%	12.77	22.23
UM SKU 31	14.98	84.71%	86.68%	8.11	14.92	15.08	85.24%	87.27%	8.13	14.92
UM SKU 32	11.05	95.79%	97.51%	4.16	7.84	10.96	95.47%	97.36%	4.17	7.85
UM SKU 33	58.33	95.71%	95.98%	43.55	60.28	58.66	95.99%	96.23%	43.78	60.52
UM SKU 34	58.44	97.43%	98.13%	43.70	60.79	58.64	97.56%	98.33%	43.87	61.08
UM SKU 35	11.54	94.57%	96.32%	2.25	3.46	11.53	94.05%	96.06%	2.26	3.47
UM SKU 36	15.88	96.24%	96.99%	7.90	15.96	15.91	96.41%	97.26%	7.91	15.99
UM SKU 37	15.48	96.90%	98.26%	7.53	13.40	15.41	96.61%	98.05%	7.55	13.43
UM SKU 38	20.92	98.42%	98.98%	2.80	5.67	20.85	98.42%	98.98%	2.81	5.68
Avg. UM	25.07	94.36%	95.81%	12.48	20.31	25.13	94.46%	95.88%	12.52	20.37
Difference	-56.57%	-2.46%	-1.10%	-70.65%	-72.95%	-56.69%	-2.67%	-1.15%	-70.58%	-72.97%
Avg. total	17.98	93.20%	95.28%	8.07	12.90	18.01	93.20%	95.33%	8.10	12.94

Table C4: *Performance of "Add undershoot"-method*

	Lost sales					*Full back-ordering*				
SKU	**SoH**	$\boldsymbol{\beta}$	$\boldsymbol{\gamma}$	**ss**	**s**	**SoH**	$\boldsymbol{\beta}$	$\boldsymbol{\gamma}$	**ss**	**s**
MM SKU 1	12.65	87.72%	92.21%	3.07	7.08	12.57	88.04%	92.43%	3.07	7.08
MM SKU 2	14.59	95.48%	97.19%	1.76	3.46	14.55	94.96%	96.92%	1.76	3.46
MM SKU 3	15.69	97.60%	98.55%	3.45	4.37	15.67	97.07%	98.09%	3.45	4.37
MM SKU 4	5.88	85.02%	89.38%	1.41	3.60	5.80	82.70%	88.32%	1.41	3.60
MM SKU 5	5.55	89.20%	92.26%	0.91	2.56	5.49	87.90%	91.41%	0.91	2.56
MM SKU 6	6.42	94.56%	96.18%	0.81	2.17	6.40	94.83%	96.23%	0.81	2.17
MM SKU 7	6.39	89.46%	92.76%	1.42	3.46	6.34	89.29%	93.09%	1.42	3.46
MM SKU 8	6.70	86.22%	91.34%	1.57	3.38	6.63	87.82%	92.54%	1.57	3.38
MM SKU 9	35.00	95.00%	96.47%	1.28	2.28	35.16	95.78%	97.25%	1.28	2.28
MM SKU 10	6.81	91.26%	93.65%	1.74	3.53	6.79	89.89%	92.97%	1.74	3.53
MM SKU 11	13.13	90.15%	92.81%	3.22	5.75	13.10	90.73%	93.44%	3.22	5.75
MM SKU 12	5.37	81.13%	86.89%	2.11	4.34	5.38	80.92%	86.97%	2.11	4.35
MM SKU 13	7.67	86.96%	91.56%	1.72	3.37	7.66	86.09%	91.32%	1.72	3.37
MM SKU 14	6.48	89.67%	92.12%	1.51	3.30	6.36	89.49%	91.96%	1.51	3.30
MM SKU 15	6.88	80.83%	87.44%	1.57	3.24	7.03	81.77%	88.47%	1.57	3.24
MM SKU 16	6.92	79.38%	86.72%	2.11	3.81	6.91	79.73%	87.12%	2.11	3.81
MM SKU 17	6.16	79.90%	85.62%	1.33	2.54	6.10	79.92%	85.93%	1.33	2.55
MM SKU 18	7.32	81.29%	87.44%	2.41	4.57	7.15	79.02%	85.60%	2.41	4.57
MM SKU 19	7.00	78.93%	84.43%	2.48	6.75	6.98	78.28%	84.48%	2.48	6.75
MM SKU 20	6.56	84.28%	88.36%	0.95	2.16	6.49	83.86%	88.11%	0.95	2.16
MM SKU 21	7.30	81.92%	89.64%	1.52	3.10	7.24	81.84%	89.73%	1.52	3.10
MM SKU 22	5.29	87.34%	90.66%	0.99	2.00	5.12	87.25%	91.58%	0.99	2.00
Avg. MM	9.17	86.97%	91.08%	1.79	3.67	9.13	86.69%	91.09%	1.79	3.67
UM SKU 23	25.66	95.68%	96.84%	4.08	10.76	25.67	95.59%	96.60%	4.08	10.76
UM SKU 24	10.17	69.36%	74.79%	3.28	13.57	10.19	70.11%	76.25%	3.28	13.57
UM SKU 25	8.96	92.35%	93.92%	1.26	4.16	8.80	92.48%	94.40%	1.26	4.16
UM SKU 26	20.49	94.99%	96.32%	5.03	21.59	20.66	95.00%	96.49%	5.03	21.59
UM SKU 27	21.93	94.97%	97.20%	3.61	13.42	21.80	95.30%	97.52%	3.61	13.42
UM SKU 28	8.64	89.45%	92.36%	1.39	5.99	8.54	89.24%	92.16%	1.39	5.99
UM SKU 29	28.38	96.26%	97.66%	2.43	5.35	28.27	96.02%	97.58%	2.43	5.35
UM SKU 30	19.81	81.87%	84.15%	3.35	12.89	19.00	79.79%	82.96%	3.35	12.89
UM SKU 31	9.50	74.54%	77.45%	2.03	8.76	9.74	74.75%	77.43%	2.03	8.76
UM SKU 32	8.41	91.90%	94.34%	1.31	5.05	8.31	91.75%	94.58%	1.31	5.05
UM SKU 33	21.20	77.91%	79.37%	5.15	21.63	21.85	78.61%	80.97%	5.15	21.63
UM SKU 34	22.11	79.12%	82.62%	5.57	22.77	22.59	79.19%	82.83%	5.57	22.77
UM SKU 35	10.62	91.68%	94.35%	1.49	2.50	10.57	91.49%	94.32%	1.49	2.50
UM SKU 36	10.80	91.56%	92.93%	2.42	10.60	10.82	91.60%	93.14%	2.42	10.60
UM SKU 37	10.14	92.16%	94.73%	2.14	8.02	10.05	92.84%	95.03%	2.14	8.02
UM SKU 38	19.65	98.22%	98.82%	1.40	4.39	19.55	98.08%	98.71%	1.40	4.39
Avg. UM	16.03	88.25%	90.49%	2.87	10.72	16.02	88.24%	90.68%	2.87	10.72
Difference	-42.79%	-1.45%	0.65%	-37.70%	-65.72%	-43.01%	-1.76%	0.45%	-37.70%	-65.71%
Avg. total	12.60	87.61%	90.78%	2.33	7.20	12.58	87.46%	90.89%	2.33	7.20

Table C5: *Performance of "Overwrite"-method*

	Lost sales					*Full back-ordering*				
SKU	**SoH**	$\boldsymbol{\beta}$	γ	**ss**	**s**	**SoH**	$\boldsymbol{\beta}$	γ	**ss**	**s**
MM SKU 1	14.69	91.38%	94.45%	5.43	9.61	14.77	91.80%	94.87%	5.45	9.64
MM SKU 2	14.74	97.05%	98.27%	1.98	3.68	14.74	97.05%	98.27%	1.99	3.68
MM SKU 3	15.77	97.60%	98.55%	3.64	4.57	15.83	97.07%	98.09%	3.64	4.57
MM SKU 4	6.82	86.89%	90.52%	2.49	4.57	6.90	87.50%	91.85%	2.51	4.88
MM SKU 5	6.27	90.37%	93.62%	1.59	3.37	6.21	89.60%	93.46%	1.61	3.40
MM SKU 6	6.62	96.19%	97.18%	1.01	2.37	6.62	95.95%	97.06%	1.01	2.40
MM SKU 7	7.04	89.64%	92.91%	2.13	4.13	7.01	90.52%	93.67%	2.15	4.16
MM SKU 8	7.19	88.13%	93.34%	2.12	3.91	7.16	89.24%	93.98%	2.13	3.91
MM SKU 9	35.01	95.00%	96.47%	1.31	2.29	35.18	95.78%	97.25%	1.31	2.29
MM SKU 10	7.33	93.36%	95.47%	2.35	4.18	7.30	91.75%	94.65%	2.38	4.20
MM SKU 11	14.01	92.08%	93.91%	4.12	6.64	13.96	92.42%	94.40%	4.13	6.66
MM SKU 12	6.38	86.31%	90.11%	3.20	5.29	6.42	86.44%	90.79%	3.23	5.43
MM SKU 13	7.96	87.87%	92.02%	2.08	3.72	7.96	87.35%	92.07%	2.09	3.74
MM SKU 14	6.82	92.09%	93.56%	1.91	3.69	6.70	91.85%	93.48%	1.92	3.71
MM SKU 15	7.25	83.78%	89.39%	1.87	3.53	7.30	82.75%	88.78%	1.88	3.53
MM SKU 16	7.67	82.60%	89.48%	3.03	4.70	7.67	82.73%	89.69%	3.03	4.72
MM SKU 17	6.69	82.06%	87.27%	1.88	3.18	6.70	83.00%	87.99%	1.90	3.22
MM SKU 18	7.95	82.14%	87.96%	3.22	5.34	7.86	79.86%	86.13%	3.23	5.36
MM SKU 19	9.73	90.23%	93.17%	5.60	9.98	9.71	90.16%	92.98%	5.62	10.01
MM SKU 20	6.69	84.60%	88.77%	1.13	2.33	6.64	84.18%	88.52%	1.13	2.35
MM SKU 21	7.42	83.60%	90.14%	1.69	3.28	7.40	83.52%	90.23%	1.70	3.28
MM SKU 22	5.31	87.55%	90.98%	1.01	2.02	5.13	87.46%	91.91%	1.01	2.02
Avg. MM	9.79	89.11%	92.62%	2.49	4.38	9.78	89.00%	92.73%	2.50	4.42
UM SKU 23	28.45	96.11%	97.64%	6.80	13.50	28.34	96.07%	97.57%	6.82	13.53
UM SKU 24	13.84	76.45%	81.52%	8.31	18.47	14.12	77.42%	81.90%	8.33	18.53
UM SKU 25	10.00	93.45%	95.23%	2.49	5.28	9.98	93.65%	95.73%	2.50	5.33
UM SKU 26	33.03	97.55%	98.20%	17.38	33.93	33.07	97.63%	98.24%	17.47	34.04
UM SKU 27	24.43	95.50%	97.62%	6.63	16.33	24.46	96.11%	97.79%	6.65	16.38
UM SKU 28	10.48	92.56%	95.53%	3.42	8.04	10.49	92.42%	95.27%	3.43	8.05
UM SKU 29	28.99	96.93%	97.92%	3.05	5.99	28.85	96.75%	97.85%	3.06	5.99
UM SKU 30	24.70	87.95%	89.79%	9.57	19.11	24.62	87.98%	89.86%	9.60	19.12
UM SKU 31	13.27	82.22%	84.61%	6.08	12.84	13.20	82.71%	85.45%	6.10	12.84
UM SKU 32	9.90	94.68%	96.78%	2.88	6.65	9.92	94.64%	96.89%	2.89	6.66
UM SKU 33	53.04	94.34%	94.71%	38.34	55.01	53.54	94.87%	95.26%	38.56	55.30
UM SKU 34	52.98	95.19%	96.54%	38.14	55.28	53.10	95.08%	96.63%	38.31	55.44
UM SKU 35	10.75	92.40%	94.71%	1.70	2.71	10.71	92.21%	94.68%	1.71	2.73
UM SKU 36	13.69	94.57%	95.90%	5.52	13.64	13.68	94.78%	95.98%	5.54	13.65
UM SKU 37	13.24	94.92%	96.90%	5.39	11.28	13.29	95.20%	97.13%	5.41	11.28
UM SKU 38	20.08	98.29%	99.01%	1.86	4.84	20.00	98.15%	98.90%	1.86	4.84
Avg. UM	22.55	92.70%	94.54%	9.85	17.68	22.59	92.85%	94.70%	9.89	17.73
Difference	-56.60%	-3.86%	-2.03%	-74.70%	-75.22%	-56.70%	-4.15%	-2.07%	-74.71%	-75.09%
Avg. total	16.17	90.90%	93.58%	6.17	11.03	16.18	90.93%	93.71%	6.20	11.07

Table C6: *Performance of "Max"-method*

	Lost sales					*Full back-ordering*				
SKU	**SoH**	β	γ	**ss**	**s**	**SoH**	β	γ	**ss**	**s**
MM SKU 1	18.67	95.26%	97.20%	9.38	13.33	18.65	95.61%	97.28%	9.38	13.33
MM SKU 2	18.72	97.57%	98.61%	6.33	7.66	18.72	97.57%	98.61%	6.33	7.66
MM SKU 3	23.12	99.07%	99.38%	10.87	11.83	23.20	99.07%	99.38%	10.87	11.83
MM SKU 4	7.86	91.32%	94.46%	3.36	5.65	7.77	89.57%	93.18%	3.36	5.65
MM SKU 5	6.12	92.26%	94.71%	1.64	3.43	6.05	90.96%	93.66%	1.64	3.43
MM SKU 6	8.90	97.10%	98.18%	3.88	4.69	8.91	97.62%	98.35%	3.88	4.69
MM SKU 7	8.86	94.61%	96.65%	3.82	5.61	8.90	94.51%	96.56%	3.82	5.61
MM SKU 8	9.23	95.37%	97.15%	3.94	5.51	9.19	94.86%	96.95%	3.94	5.51
MM SKU 9	38.33	96.13%	97.86%	4.09	4.82	38.48	96.47%	97.99%	4.09	4.82
MM SKU 10	8.58	93.73%	95.85%	3.75	5.39	8.52	94.01%	96.09%	3.75	5.39
MM SKU 11	20.47	96.01%	97.03%	10.76	12.87	20.43	96.01%	97.03%	10.76	12.87
MM SKU 12	8.77	92.34%	95.23%	5.60	7.49	8.78	93.41%	95.66%	5.60	7.49
MM SKU 13	9.21	92.21%	95.95%	3.79	5.01	9.18	92.21%	95.95%	3.79	5.01
MM SKU 14	8.36	95.82%	96.76%	3.52	5.24	8.35	95.79%	96.81%	3.52	5.24
MM SKU 15	8.84	88.55%	93.43%	3.87	5.59	8.92	89.12%	93.76%	3.87	5.59
MM SKU 16	8.36	88.39%	92.03%	3.76	5.75	8.40	88.17%	92.16%	3.76	5.75
MM SKU 17	8.91	88.03%	91.95%	3.92	5.50	8.94	87.07%	91.51%	3.92	5.50
MM SKU 18	12.33	94.78%	96.57%	7.84	9.59	12.31	94.78%	96.57%	7.84	9.59
MM SKU 19	8.52	86.05%	90.49%	4.42	8.60	8.48	85.54%	90.07%	4.42	8.60
MM SKU 20	9.65	91.71%	93.76%	4.13	4.77	9.62	91.43%	93.51%	4.13	4.77
MM SKU 21	9.02	85.86%	91.84%	4.14	4.88	9.02	85.86%	91.84%	4.14	4.88
MM SKU 22	8.04	95.88%	96.80%	3.97	4.97	7.98	95.97%	96.85%	3.97	4.97
Avg. MM	12.22	93.09%	95.54%	5.04	6.74	12.22	92.98%	95.44%	5.04	6.74
UM SKU 23	25.05	94.90%	97.07%	3.38	9.89	24.94	94.80%	96.82%	3.38	9.89
UM SKU 24	15.33	79.80%	83.86%	10.19	20.17	15.52	79.25%	83.50%	10.19	20.17
UM SKU 25	9.26	93.02%	94.79%	1.54	4.29	9.09	92.62%	95.04%	1.54	4.29
UM SKU 26	24.05	96.37%	97.36%	8.70	25.21	24.19	96.43%	97.36%	8.70	25.21
UM SKU 27	22.90	95.29%	97.73%	4.94	14.67	22.83	95.87%	97.92%	4.94	14.67
UM SKU 28	9.40	90.62%	93.97%	2.19	6.78	9.26	90.25%	93.53%	2.19	6.78
UM SKU 29	28.86	96.35%	97.88%	3.00	5.86	28.79	96.15%	97.83%	3.00	5.86
UM SKU 30	22.24	84.54%	86.81%	6.14	15.47	21.78	83.12%	85.71%	6.14	15.47
UM SKU 31	10.91	76.79%	79.46%	3.62	10.28	11.00	76.79%	79.79%	3.62	10.28
UM SKU 32	8.85	92.81%	95.06%	1.79	5.44	8.69	92.74%	95.20%	1.79	5.44
UM SKU 33	24.70	81.40%	82.66%	8.61	25.03	24.77	81.72%	83.09%	8.61	25.03
UM SKU 34	25.31	82.12%	84.78%	9.13	26.28	25.53	82.49%	84.87%	9.13	26.28
UM SKU 35	14.97	97.54%	99.15%	5.92	7.64	14.96	97.54%	99.15%	5.92	7.64
UM SKU 36	12.14	93.39%	94.58%	4.00	12.06	12.21	93.65%	95.04%	4.00	12.06
UM SKU 37	13.23	95.20%	97.26%	5.60	11.22	13.18	95.62%	97.53%	5.60	11.22
UM SKU 38	19.78	97.75%	98.32%	1.63	4.46	19.64	97.70%	98.28%	1.63	4.46
Avg. UM	17.94	90.49%	92.55%	5.02	12.80	17.90	90.42%	92.54%	5.02	12.80
Difference	-31.86%	2.87%	3.24%	0.23%	-47.36%	-31.74%	2.83%	3.14%	0.23%	-47.36%
Avg. total	15.08	91.79%	94.04%	5.03	9.77	15.06	91.70%	93.99%	5.03	9.77

Table C7: *Performance of "First moment β^*"-method*

	Lost sales					*Full back-ordering*				
SKU	**SoH**	$\boldsymbol{\beta}$	γ	**ss**	**s**	**SoH**	$\boldsymbol{\beta}$	γ	**ss**	**s**
MM SKU 1	15.33	94.45%	96.33%	6.21	10.15	15.40	94.44%	96.38%	6.21	10.15
MM SKU 2	15.51	98.62%	99.42%	3.33	4.46	15.48	98.62%	99.42%	3.33	4.46
MM SKU 3	21.60	100.00%	100.00%	9.43	10.28	21.67	100.00%	100.00%	9.43	10.28
MM SKU 4	7.82	93.60%	96.05%	3.67	5.70	7.80	93.26%	95.55%	3.67	5.70
MM SKU 5	8.35	98.32%	98.82%	4.00	5.52	8.36	98.32%	98.82%	4.00	5.52
MM SKU 6	8.28	96.62%	98.08%	2.51	3.34	8.25	97.40%	98.32%	2.51	3.34
MM SKU 7	9.74	96.78%	98.32%	5.08	6.59	9.78	97.11%	98.61%	5.08	6.59
MM SKU 8	11.16	97.84%	98.81%	6.26	7.54	11.13	98.31%	98.99%	6.26	7.54
MM SKU 9	35.81	95.43%	97.18%	1.56	2.33	36.00	96.03%	97.70%	1.56	2.33
MM SKU 10	10.73	97.45%	98.41%	6.04	7.56	10.72	97.45%	98.41%	6.04	7.56
MM SKU 11	19.25	99.24%	99.14%	9.67	11.78	19.24	98.46%	98.47%	9.67	11.78
MM SKU 12	9.89	94.92%	97.03%	6.83	8.64	9.89	95.09%	97.00%	6.83	8.64
MM SKU 13	10.59	95.89%	97.76%	5.38	6.39	10.65	95.75%	97.76%	5.38	6.39
MM SKU 14	14.59	100.00%	100.00%	10.11	11.52	14.59	100.00%	100.00%	10.11	11.52
MM SKU 15	10.01	92.84%	96.22%	5.33	6.43	10.10	93.84%	96.87%	5.33	6.43
MM SKU 16	10.56	94.06%	96.73%	6.05	7.74	10.57	93.78%	96.64%	6.05	7.74
MM SKU 17	8.04	88.25%	92.65%	3.28	4.51	8.00	88.12%	92.26%	3.28	4.51
MM SKU 18	13.46	96.99%	98.05%	9.08	10.60	13.40	96.99%	98.05%	9.08	10.60
MM SKU 19	10.69	92.89%	95.22%	6.76	11.05	10.70	93.07%	95.49%	6.76	11.05
MM SKU 20	9.06	89.78%	92.93%	3.65	4.22	9.11	89.26%	92.75%	3.65	4.22
MM SKU 21	9.22	86.14%	91.84%	4.46	5.18	9.22	86.14%	91.84%	4.46	5.18
MM SKU 22	8.85	96.19%	97.39%	4.75	5.02	8.85	96.19%	97.39%	4.75	5.02
Avg. MM	12.66	95.29%	97.11%	5.61	7.12	12.68	95.35%	97.12%	5.61	7.12
UM SKU 23	29.16	96.82%	98.52%	7.83	14.47	29.54	96.82%	98.52%	7.83	14.47
UM SKU 24	14.13	79.27%	82.91%	8.29	18.51	14.21	79.48%	83.26%	8.29	18.51
UM SKU 25	10.44	97.15%	97.67%	2.82	5.69	10.40	96.63%	97.30%	2.82	5.69
UM SKU 26	23.96	97.09%	97.92%	8.36	24.89	23.97	96.86%	97.75%	8.36	24.89
UM SKU 27	26.54	98.71%	99.20%	8.74	18.38	26.54	98.52%	99.16%	8.74	18.38
UM SKU 28	12.79	97.12%	98.45%	5.94	10.48	12.86	97.21%	98.47%	5.94	10.48
UM SKU 29	30.05	98.33%	98.62%	4.42	7.20	29.94	98.24%	98.61%	4.42	7.20
UM SKU 30	19.58	81.46%	83.83%	3.15	12.42	18.80	79.99%	82.98%	3.15	12.42
UM SKU 31	11.61	82.14%	84.11%	5.24	11.80	12.12	81.88%	84.62%	5.24	11.80
UM SKU 32	10.84	97.37%	98.49%	3.91	7.61	10.81	97.24%	98.37%	3.91	7.61
UM SKU 33	22.74	80.55%	82.01%	6.75	23.09	23.17	80.45%	82.27%	6.75	23.09
UM SKU 34	25.33	83.15%	86.09%	8.33	25.54	24.95	82.50%	85.31%	8.33	25.54
UM SKU 35	12.49	95.67%	97.68%	3.38	4.36	12.44	95.74%	97.74%	3.38	4.36
UM SKU 36	14.44	96.54%	97.49%	6.38	14.49	14.43	96.55%	97.48%	6.38	14.49
UM SKU 37	14.42	97.94%	99.13%	6.59	12.50	14.39	97.90%	99.09%	6.59	12.50
UM SKU 38	20.97	98.76%	99.36%	2.82	5.77	20.94	98.76%	99.36%	2.82	5.77
Avg. UM	18.72	92.38%	93.84%	5.81	13.58	18.72	92.17%	93.77%	5.81	13.58
Difference	-32.36%	3.14%	3.48%	-3.44%	-47.59%	-32.28%	3.44%	3.58%	-3.44%	-47.59%
Avg. total	15.69	93.83%	95.48%	5.71	10.35	15.70	93.76%	95.45%	5.71	10.35

Table C8: *Performance of target fill rate compound Poisson method*

D Method for reorder point

We focus on the determination of expected demand during the cover period, and we simply add the safety stock to establish a reorder point. Instead of averaging all demand and customers as is standard practice within a (R,s,nQ) model in order to determine the expected demand during the cover period, we can incorporate our in-service order line size. If multiplied by the rounded customer arrival rate during the cover period, we establish our measure for expected demand during the cover period:

$$s = \text{safety stock} + \left\lceil \lambda_{L+R} \cdot E(X) \right\rceil$$

Let λ_{L+R} be the arrival rate for customers during the cover period. Since safety stocks calculated through literature-based methods of Tijms & Groenevelt, Nenes, and the compound Poisson method, are a direct result of subtracting the corresponding expected demand during the cover period from the reorder point, we only combine our own heuristics for safety stock with this heuristic for the reorder point.

It appeared that our heuristics for setting the reorder point performs extremely worse on average stock on hand. Relative to the standard methods, even the 'best' III.B. method *ssA* (an obsolete heuristic) resulted in an average stock on hand that was at least 82% higher. Hence, we see no point in further investigating this *second reorder point determination*. This method was already arguable in itself, since reorder point by its definition solely consists of the safety stock and the expected demand during the lead time (and review period).

E Data gathering

In accordance with the qualitative description as provided in Section 2.1 and our analysis in Chapter 4, we develop a framework for measuring performance of the standard (R,s,nQ) methods, literature-based methods, and our own heuristics. Data are gathered from two main datasets of company X, the *ArticleInfo* table and *Transactions* table. The *ArticleInfo* table contains the SKUs sold by company X, through attributes like among others the SKU code, shop identifier, replenishment minimum order quantity, replenishment incremental order quantity, the target volume fill rate. The *Transactions* table on the other hand contains all the sales transactions made by customers of company X, one could say this is the digital equivalence of a shopping receipt. This table includes among others the transaction number, SKU code, sold quantity, and the transaction date.

The data from the *ArticleInfo* table functions as input parameters for inventory calculations. The sales data from the *Transactions* table are loaded into the framework, and used to evaluate the distinct methods, see Section 5.2. Evaluation of performance is done on each SKU code, hence on each SKU-shop, as stock-keeping and sales take place on shop level.

For the setting of inventory control parameters we divide the half year of data into two parts. We work with 38 SKUs, sold in 94 shops, and accounting for 64,393 transactions. The first eights weeks, from now on to be called *initialization period* are used for calculating the *sample mean* $\bar{X}$ and *sample standard deviation* s of the weekly demand, the customer arrival rate, and the in-service order line size. This initialization period is sufficient, as we focus on SKUs that show a stationary demand pattern - tested trough a t-Test - which we discuss in more detail in Section 5.1. In case of one or no orders at all of the specific SKU-shop combination took place during this initialization period - which would result in a meaningless mean and standard deviation - the average and standard deviation of the weekly demand for the SKU is averaged on all shops selling this SKU. In this case, we exploit the population standard deviation, as we observe the full population.

F Gamma approximation Tijms & Groenevelt

In order to transform the k determination of the Ju(k) part from Equation 3.10, into an approximation of the Gamma distribution, we employed Equation *Appendices*.1 for $0.5 < CV_{D_{L+R}} \leq 1/\sqrt{2}$, based on the mixture of two Erlang distributions E_{k-1} and E_k with $k = \lceil \frac{1}{CV_{D_{L+R}}} \rceil$ and $k \geq 2$ (Tijms & Groenevelt, 1984).

$$k(x) = p \cdot b^{k-1} \cdot \frac{x^{k-2}}{(k-2)!} \cdot e^{-b \cdot x} + (1-p) \cdot b^k \cdot \frac{x^{k-1}}{(k-1)!} \cdot e^{-b \cdot x}, \qquad x \geq 0 \qquad (Appendices.1)$$

where

$$p = \frac{1}{1 + CV^2_{D_{L+R}}} \cdot \left[k \cdot CV^2_{D_{L+R}} - \sqrt{k \cdot (1 + CV^2_{D_{L+R}}) - k^2 \cdot CV^2_{D_{L+R}}} \right] \text{ and } b = \frac{k-p}{E(D_{L+R})} \quad (Appendices.2)$$

Next, we employed Equation *Appendices*.3 for $CV_{D_{L+R}} > 1/\sqrt{2}$ (Tijms & Groenevelt, 1984), using the hyper-exponential distribution.

$$k(x) = p \cdot b_1 \cdot e^{-b_1 \cdot x} + (1-p) \cdot b_2 \cdot e^{-b_2 \cdot x}, \qquad x \geq 0 \qquad (Appendices.3)$$

where

$$b_{1,2} = \frac{2}{E(D_{L+R})} \pm \frac{2}{E(D_{L+R})} \cdot \sqrt{\frac{CV^2_{D_{L+R}} - \frac{1}{2}}{1 + CV^2_{D_{L+R}}}} \text{ and } p = \frac{b_1 \cdot (b_2 \cdot E(D_{L+R}) - 1)}{b_2 - b_1} \qquad (Appendices.4)$$

Unfortunately, Equations *Appendices*.1 and *Appendices*.3 led to bad results around the threshold for the coefficient of variation of $1/\sqrt{2}$. For $1/\sqrt{2} < CV_{D_{L+R}} \leq 1$ the p of Equation *Appendices*.3 got negative (as also stated by the authors), making no sense in any more.

Next, we did not manage to get the method for $CV_{D_{L+R}} \leq 1/\sqrt{2}$ to yield results (k-factors were found to be negative, which makes no sense). As 18.51% of the SKU-shop combinations of research show a $CV_{D_{L+R}} \leq 1$, for which we did not find a working method using Gamma distribution, nor may we apply the normal distribution - which is only valid for $CV_{D_{L+R}} < 0.5$ or even $CV_{D_{L+R}} < 0.33$ (Silver et al., 2017). Hence, a rather large bit (18.51%) of the SKU-shop combinations cannot be calculated using the normal or Gamma distribution. Thus, we exclude this approach - the only one left assuming normal or Gamma demand - from further exploration in the final experiment.