MSc Thesis

A time-based order fill rate model for spare parts

Erik Raesen BSc. – 0021822

Graduation project

Master of Science in Applied Mathematics and
Master of Science in Industrial Engineering and Management

University of Twente
Enschede - The Netherlands

VanDerLande Industries

Dr. Matthieu van der Heijden
Prof.Dr. Johann Hurink

Drs. Katja Kleinveld
Ir. Joost Herman

Abstract

The service central warehouse at Vanderlande holds a number of spare parts on stock to fill customer orders for spare parts. In this research the supply chain of spare parts at Vanderlande is analyzed and subsequently an inventory model for the service central warehouse is developed that optimizes the percentage of customer orders that is completely filled within a given timeframe, with respect to the total holding costs. Real life data has been entered into the model and resulted in significant improvements in both holding costs and order fill rate in comparison with the current stock levels.

Preface

This thesis is the result of a graduation project conducted at Vanderlande Industries BV in Veghel, in order to finalize two master studies: Applied Mathematics and Industrial Engineering & Management, both at the University of Twente in Enschede.

Beside the research-related goals, I have tried to achieve some personal goals during my research project at Vanderlande.

- **Autonomously conduct a large research:** Since the research size equals a study load of 55 ECTS we may define this research as a large project. Although I have experience in managing and conducting small research projects of at most 20 ECTS like my Bachelor Thesis and a traineeship, it is somewhat different to manage and conduct a research of this size. Through this research I have experienced how to do this.
- **Deliver a complete and useable end-product to Vanderlande:** Besides designing a theoretical model that optimizes the problem stated by Vanderlande, it is my goal to deliver a tool of some kind that applies the model to the actual situation of Vanderlande and can be used at all times. It is more satisfying to me if the model I design is actually used in practice, but to achieve this Vanderlande must have the ability to apply the model.
- **Experience working in a large company:** All my previous relevant work experiences in companies, either in a traineeship or in (holiday) jobs, were in relatively small companies. With over 1,800 employees, out of which nearly 1,000 people work in Veghel, Vanderlande could be called a large company. During my research project I hope to experience e.g. the hierarchy in such a company. This can help me in determining in what kind of company I would like to work in the future.

I would like to thank my supervisors of the university, Matthieu van der Heijden and Johann Hurink, for their support and feedback throughout the research. The research took longer than expected and planned and sometimes structure was hard to find. Their advice helped me to keep in control of the research and not to get fixated on details too long.

Of course a big “Thank you” to my daily supervisors at Vanderlande, Katja Kleinveld and Joost Herman, is in place here. They gave me the necessary insights in the complicated world of spare parts supply at Vanderlande. I have enjoyed working with them and all the other colleagues of the Supply Chain Management Services department. Also I would like to thank Harold Bol for supervising me in the first part of my research.

Finally, I would like to thank everyone around me for making the eight years of study at the university the best eight years of my life so far. In particular I want to thank my parents for their support; not only financially, but also by allowing me to take the maximum out of my time at the university. I also want to thank my girlfriend Ineke for her support; it still amazes me that something this good was already this close to me the whole time.

Erik Raesen
Veghel, October 2009

Management summary

This report is the result of a graduation project executed at the service department of Vanderlande Industries BV in Veghel, the Netherlands. The service department is amongst other activities responsible for the supply of spare parts to customers of Vanderlande around the world. In this supply the department makes use of the service central warehouse.

Service is becoming increasingly important as a business unit of Vanderlande. Therefore a roadmap has been defined by the service department to improve themselves in a number of areas. The goal of this research is twofold. First we analyze the spare parts supply chain from the point of view of the service central warehouse. Next we develop an inventory model for this warehouse. This approach is formulated in the research problem:

> *Develop an inventory policy for spare parts at the service central warehouse of Vanderlande Industries, taking into account the structure of the supply chain. Analyze how the stock levels resulting from the model influence the performance towards customers regarding on-time-delivery and customer order lead time.*

The analysis of the spare parts supply chain has led to a division of the customer demand for spare parts into five demand streams.

- Emergency orders
- Preventive maintenance orders
- Replenishment orders
- Spare part packages
- Revisions, modifications and retrofits (RMR)

Only preventive maintenance orders and replenishment orders are of interest for the service central warehouse. The other demand streams only scarcely occur or do not influence the stock levels since they are generally not picked from stock. The two relevant demand streams are grouped together in one stream, service orders, since a separation between these streams is not made currently and we expect the demand streams to be much alike.

We have developed an inventory model that optimizes the time-based order fill rate against minimal holding costs. The model makes use of a greedy heuristic to find the optimal reorder points for each spare part; the heuristic terminates at either the target order fill rate or the target holding costs as indicated by the user.

The current performance with respect to service orders is a time-based order fill rate of 60.2% and holding costs of €37,000 for a timeframe of 15 workdays. The inventory model results in a decrease of holding costs of nearly 20% at a time-based order fill rate of at least 60%. In practice this order fill rate will be higher since in some cases spare parts can be delivered from a warehouse nearby, hence reducing the assumed replenishment lead time. Spare parts with low replenishment lead times are not held on stock anymore since these spare parts can be delivered within the timeframe by the supplier. Spare parts with a high demand rate, low holding costs or underperforming suppliers (with respect to on time delivery of replenishment orders) are put on stock.

Analysis of different scenarios shows that the order fill rate increases slowly if the holding costs exceed €70,000. This means that very high holding costs only lead to slightly better order fill rates. Different timeframes can be chosen in the execution of the model; the higher the timeframe, the lower the number of spare parts with low replenishment lead times that is held on stock. This affects the order fill rate for other timeframes; a high order fill rate for a timeframe of 20 days can result in a very low order fill rate from stock.

The model has been implemented at Vanderlande through manuals and training. Two interfaces are created; one to convert the data from the database to input files for the model and one to apply the model for the real life data. In this interface the user is able to change a number of parameters such as the timeframe and target time-based order fill rate and/or holding costs. Beside applying the model, it is possible for the user to calculate the time-based order fill rate and holding costs by manually determining desired performance for categories of spare parts.

To improve the performance of the inventory model, the quality of the input data needs to be improved. In our model, the holding costs are only an estimate of the true holding costs. Analysis has shown that the holding costs largely contributes to the decision to put spare parts on stock. A better model to forecast spare parts demand is needed as well. Currently the demand is forecasted based on historical data alone; using e.g. information on the installed base can improve the forecasting model.

Interesting areas of further research are a division between the desired performance for preventive maintenance and replenishment orders, and including system locations into a two-echelon inventory model. A division between the two demand streams can only be carried out if detailed information of the characteristics of each of these demand streams is available; therefore the type of demand streams needs to be recorded in the future. Extending the inventory model to a two-echelon model can be carried out using the research of (van Sommeren, 2007) – concerning optimal stock levels at system locations – and this research as a basis.

Index

1. Introduction

In this chapter a brief description is given of the company and the department in which this research is conducted. In section 1.2 the research questions and goals are formulated.

1.1. Company description

Vanderlande Industries was founded in 1949 as “Machinefabriek E. van der Lande” in Veghel, the Netherlands. At that time it was a general machinery and construction company which manufactured hoists, cranes and conveyor equipment for bulk material handling and oil drums. Through joint ventures and establishments the company has grown to a multinational company. Today, Vanderlande Industries has over 1800 employees in four continents. Its headquarters and factory where parts are being produced are still located in Veghel.

Vanderlande Industries’ activities can be divided into four categories or business areas:

- *Baggage handling:* Vanderlande Industries designs, builds and services baggage handling systems for airports of all sizes. These systems are located at several major international airports such as Schiphol Airport, London Heathrow and Hong Kong Int. Airport.
- *Distribution:* Vanderlande Industries provides automated handling systems for order selection and sorting in distribution centers. These systems are installed at amongst others Daimler, PLUS retail and Amazon.de.
- *Parcel & Post:* Vanderlande Industries offers a wide range of technologies for parcel handling and documents. The system size varies from some thousands parcels per day up to 100,000 parcels per hour. Customers with parcel & post systems are amongst others UPS, TNT and DHL.
- *Services:* The services Vanderlande Industries provides vary from preventive and corrective maintenance to complete spare part deliveries from a ‘one-stop-shop’ and operation and process support through e.g. training.

Due to increasing worldwide competition and shrinking profit margins, high-technology-product manufacturers are forced to find new ways to differentiate themselves from their competitors. Providing a fast, high-quality after-sales service to customers can help to achieve this. Therefore, service is becoming an increasingly important activity for Vanderlande. Besides, as opposed to the other business areas where turnover is generated through projects, service generates a more continuous cash flow through service contracts. This improves the continuity of the company, which is important especially in the current global economic situation.

1.1.1. Organization

Vanderlande Industries is split up in customer centers. A customer center can be seen as a separate company; it has its own project, service, and financial organization. A customer of Vanderlande is connected to one of the customer centers, based on the type of system and/or geographical location. If a customer needs a new system, an update, or service – either maintenance or spare parts – he turns to his ‘own’ customer center. Figure 1 shows an overview of the customer centers of Vanderlande Industries.

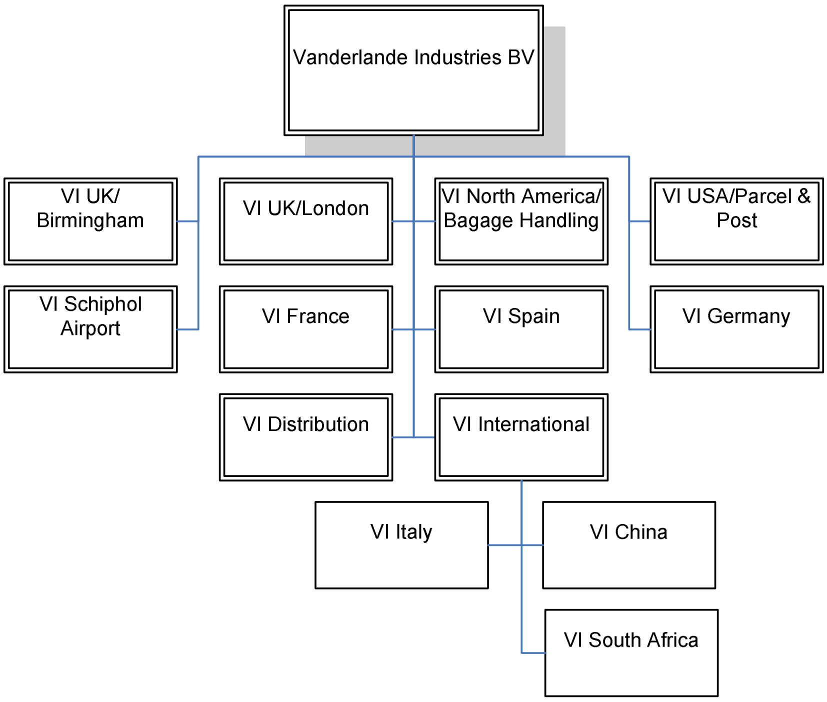

Figure 1 Vanderlande Industries and its customer centers

The customer centers Italy, China, and South Africa do not have a complete organization and are assisted by the customer center International in some operations, e.g. the processing of service orders. At the headquarters in Veghel, the general management of Vanderlande Industries, the customer center Distribution and the main part of customer center International are located.

1.1.2. Service department

Next to the customer centers, the corporate Vanderlande ("Vanderlande Industries BV") has a service department as well, organized in the business unit services. Figure 2 shows how the business unit is built up.

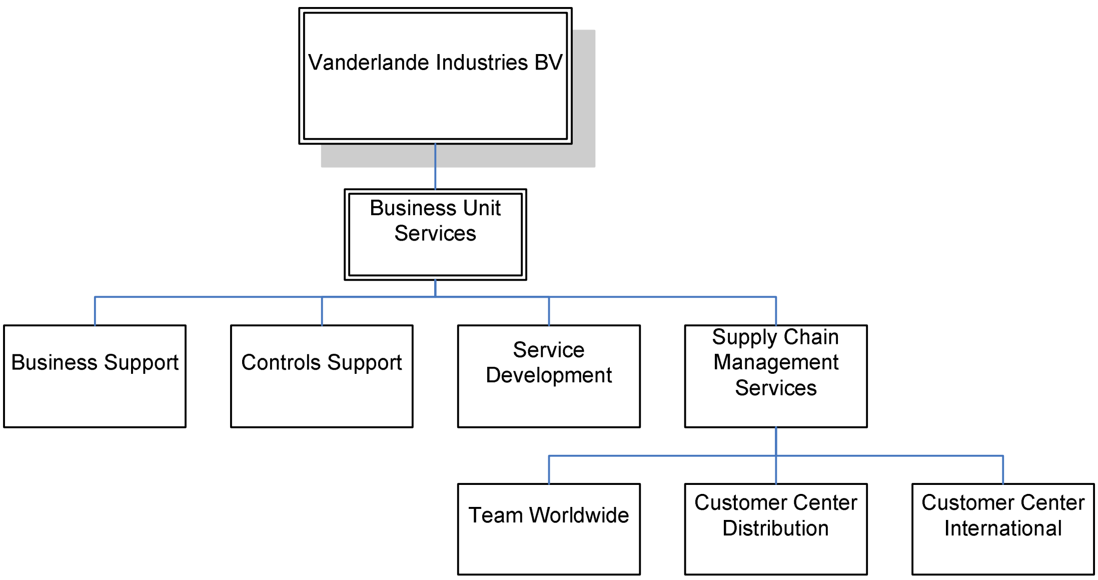

Figure 2 Organizational chart of service department

This research is conducted within the Supply Chain Management Services (SCMS) department and is mainly focused on the activities of the team Worldwide. This team is responsible for the global supply of spare parts. As stated in the previous subsection, customers turn to their customer center in case they need spare parts. The corresponding customer center then orders these spare parts at the team Worldwide.

1.1.3. Service central warehouse

In supplying spare parts the service department makes use of the service central warehouse (SCW), located in Veghel, where a number of spare parts is held on stock permanently. The amount of spare parts that need to be kept on stock has been determined at the start of the SCW in January 2007. Before the start of the SCW, a dedicated stock of spare parts was available at the warehouse of Vanderlande itself. Stock heights were determined through some simple calculation rules based on the number of orders per year and average yearly demand. Since then these amounts have not been reconsidered. The stock height is the lowest amount to have on stock at the warehouse; in case the stock drops below this amount, a replenishment order is placed at the supplier. In general a base-stock policy is applied, which means that after each customer order that is picked from stock, the stock is replenishment up to the original amount. For some spare parts fixed replenishment order quantities (larger than one) are ordered at the supplier in case the stock drops below the minimum stock height.

Since over 26,000 different spare parts exist, it is nearly impossible to store all spare parts at the SCW. In case a desired spare part is not available at the SCW, Team Worldwide orders the spare part at either an external supplier or at the factory of Vanderlande. Some spare parts are shipped directly from the supplier to the customer stock, mostly in case both the supplier and the customer are situated abroad, e.g. in China. Figure 3 shows an overview of the spare parts supply chain and the position of the customer centers and team Worldwide.

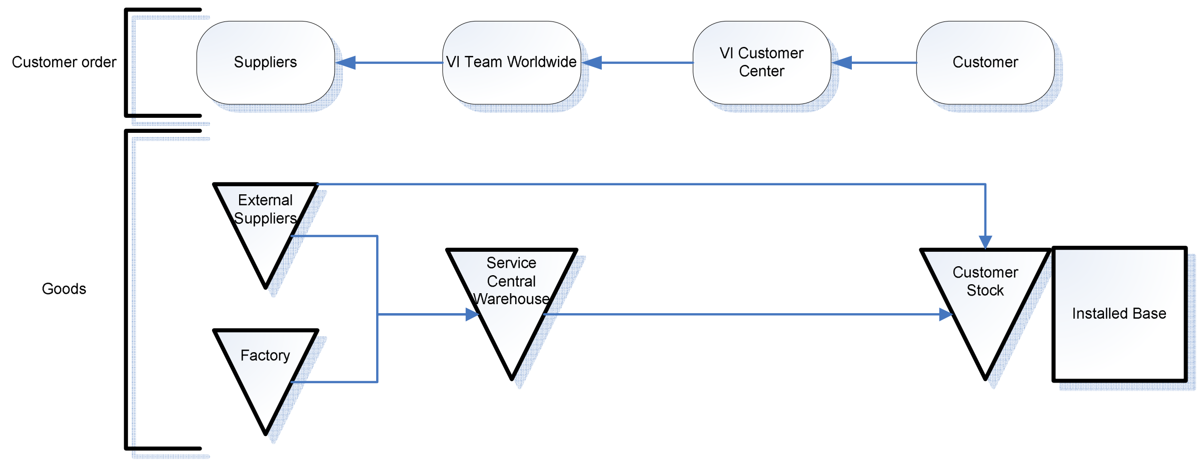

Figure 3 Overview of spare parts supply chain

The SCW is not only used to store spare parts. In case a customer orders multiple spare parts, the complete order is consolidated in the SCW. Partial shipment of orders, in case one of the ordered spare parts is delayed, is only allowed if the customer agrees to it. In many cases the customer wishes to receive only one package and accepts the additional waiting time for the whole order.

The SCW is exploited by an external party that takes care of the picking and packing of orders as well. This party charges Vanderlande for these activities as well as for keeping the spare parts on stock. These holding costs are calculated by looking at the number of pallets and shelves used by Vanderlande. Each month this number is determined.

1.2. Research assignment

Since service is becoming increasingly important within the company, Vanderlande wants to improve in this area. As a part of that strategy, in April 2006 the Supply Chain Management Services (SCMS) department was founded to create a specific department concerned with spare parts supply. In the remainder of this research we refer to this department as the service department or SCMS.

The SCMS department has defined a roadmap to improve all parts of the spare parts supply chain. In 2007, a research was conducted regarding the supply chain downstream to the customer. This research resulted in a spare parts inventory model for the customer and possible forward stock locations (van Sommeren, 2007). This research made a number of assumptions:

- All spare part orders consist of one order line; this means that only one type of spare part is ordered at a time. It is possible that a larger amount than one is ordered. This amount is referred to as the order line quantity.
- Spare part orders need to be filled as fast as possible, since a delay in the supply causes system unavailability. The research assumes that spare parts are only replaced upon breakdown.
- In case the customer or forward stock location does not have the desired spare part(s) on stock, the SCW can always supply the spare part immediately from stock. The research hence assumes an infinite stock of all spare parts at the SCW.

This research focuses on the inventory policy at the SCW. In the next section, the goals, problem statement, and research questions will be amplified.

1.2.1. Research goals

There are two main goals in this research.

- **Provide insight in the spare parts supply chain.**

  To take maximum advantage of the SCW, we need to understand the spare parts supply chain. Which demand streams are applicable at the SCW, and what performance is desired for each stream?
- **Develop a stock policy that optimizes the performance of the SCW.**

  It is the intention of the SCMS department to use the stock at the SCW to both improve the delivery performance – i.e. on time delivery – and to reduce the lead times on customer orders for spare parts. We need to address the following:
  - What performance is exactly desired by customers and Vanderlande?
  - How do the stock levels at the warehouse influence this performance?

  Vanderlande intends to implement the stock policy that results from this research. Part of this goal is hence to deliver a usable tool that applies the developed policy.

1.2.2. Problem statement

Given the research assignment and the research goals, we formulate the problem statement as follows:

> *Develop an inventory policy for spare parts at the service central warehouse of Vanderlande Industries, taking into account the structure of the supply chain. Analyze how the stock levels resulting from the model influence the performance towards customers regarding on-time-delivery and customer order lead time.*

1.2.3. Research questions

To deal with the stated problem, we address a number of research questions. These research questions can roughly be divided into three groups.

I: Analysis of the current supply chain

To be able to improve the performance towards the customers, we need to know exactly who the customers are, what performance is desired by them, and how the demand from customers is composed. To this end, we state the following research questions.

1. How is the spare parts supply chain currently composed, which demand streams can we recognize, and what performance measure is currently used?

2. What are the characteristics of each of the relevant demand streams, and what performance is desired for each demand stream?

II: Modeling an inventory policy for the SCW

After the current and desired situation have been described, we start with a basic inventory model. This basic model simplifies the supply chain and makes some (possibly unrealistic) assumptions. Afterwards we improve this model step by step until a complete model is obtained.

3. On which assumptions is the basic inventory model based and how can we determine optimal stock values from this model?

4. How do we extend the basic inventory policy and how is this advanced model optimized?

III: Applying the inventory model for the SCW

As a result from the previous research questions, we have a complete inventory model for the central warehouse. The final step in our research is to apply this model to the actual situation at Vanderlande.

5. How does the resulting inventory policy perform for real life data? What is the impact on the stock heights, customer order lead time and on time delivery performance?

1.3. Outline of the thesis

We start our research by analyzing the spare parts supply chain at Vanderlande. This analysis is reported in Chapter 2. In section 2.3 we demarcate our research based on the supply chain analysis, after which in Chapter 3 we define in detail the characteristics for the demand streams on which the model is based. A small literature study is carried out in that chapter. Chapter 4 introduces the basic model, based on a number of assumptions. Chapter 5 extends this basic model to a complete and realistic inventory model. In Chapter 6 we discuss how to subtract the data that is needed for our models from the available data at Vanderlande, after which we show and discuss the results of the complete model. In Chapter 7 the implementation of the model at Vanderlande is discussed. Finally in Chapter 8 we formulate our conclusions and recommendations and give suggestions for further research.

2. Supply chain

This chapter answers the first two research questions:

1. How is the spare parts supply chain currently composed, which demand streams can we recognize, and what performance measure is currently used?
2. What are the characteristics of each of the demand streams, and what performance is desired for each demand stream?

2.1. Customers

Most customers have one or more systems that are sold and installed by Vanderlande. Beside the system, a service contract is mostly offered to customers. This contract can include a wide variety of service activities: from periodic maintenance by engineers of Vanderlande to 24/7 assistance, from a few agreements on the supply of spare parts to complete ownership of Vanderlande of the supply of spare parts. In many cases the customer buys an initial package of spare parts upon system completion consisting of the most critical spare parts for maintenance. This small stock of spare parts is held close to the system.

One of the possible service activities Vanderlande can offer to its customers is consignment service. The small stock of critical spare parts is in case of consignment service owned by Vanderlande, and Vanderlande is responsible for the replenishment of this stock. The customer pays a fixed yearly fee for this service plus additional fees for used spare parts.

2.2. Demand streams

Customers need spare parts for a number of reasons. First of all they need an initial spare parts stock upon system completion that they use to maintain their system. If some spare parts in the system need to be replaced, the customer can pick these parts from the stock at the system location if they are available. If the desired spare parts are not available, the customer needs to order the spare parts via the customer center of Vanderlande. Finally customers sometimes desire a large system update to keep their system in optimal condition.

Based on these reasons and other variables such as order characteristics and desired performance we have defined five different demand streams:

- Emergency orders
- Preventive maintenance orders
- Replenishment orders
- Spare part packages
- Revisions, modifications and retrofits (RMR)

Each demand stream is further specified in the next sections, including some remarks on which performance is desired.

2.2.1. Emergency orders

Although most systems are regularly inspected by engineers and a lot of parts in the system are preventively replaced, it can occur that a part in the system suddenly fails, causing the system to stop or function at a lower capacity. Besides, engineers occasionally note that a part will fail in the near future if it is not replaced as soon as possible. If a part fails or is likely to fail, it needs to be replaced immediately.

In this case the customer first turns to the stock at the system location to get the needed spare part. In case the spare part is not available at the system location, the customer orders the spare parts, through the customer center, at the service department of Vanderlande with an emergency order. If the desired part is available at the SCW, it is picked and shipped to the customer as soon as possible. Otherwise the parts are ordered at the suppliers and are shipped directly to the customer to ensure the fastest delivery. Figure 4 shows an overview of the possible flows of an emergency order and the priority of these flows.

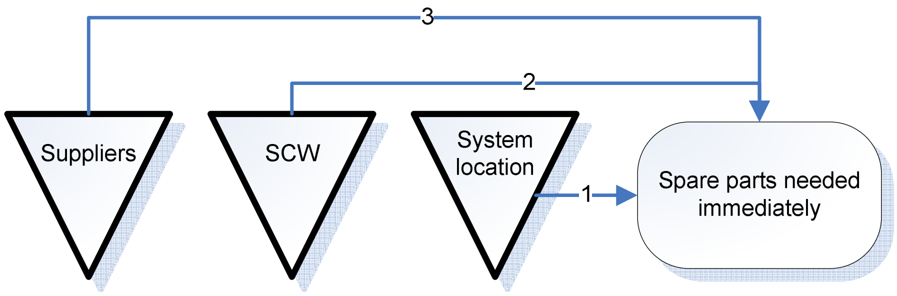

Figure 4 Overview of emergency order supply

The number of spare parts that are needed in this situation is limited. Only when there is a major system crash, a large amount of spare parts is needed by the customer, but in most cases the emergency order consists of only one order line and a relatively low quantity per ordered spare part.

The model of (van Sommeren, 2007) determines the optimal stock heights for the system locations and if applicable forward stock locations, for a given performance regarding emergency orders. The forward stock locations that are included in his research are not included in our research because currently there are no forward stock locations active.

Performance

Downtime is very costly for customers, so this needs to be prevented as much as possible. Ideally, all emergency orders are fulfilled from the stock at the system location, since in that way the downtime is minimized. Our main interest in this research is the SCW. Emergency orders from the customer only arrive here if the system location cannot fill the customer order. For some spare parts the service department has made agreements with the corresponding supplier regarding the supply of emergency orders for these parts; for those parts there is no need to keep parts on stock since the supplier has a stock to cover these emergency orders. For the remaining parts that need to be delivered from the SCW the service department wants to deliver as much complete orders as possible from stock. Within emergency orders all spare parts are immediately needed by the customer, hence it is not useful to deliver an emergency order partially. We therefore strive for a high *order fill rate*. Since most emergency orders consist of one order line, this performance comes close to the *order line fill rate*, i.e. the percentage of order lines that is delivered completely from stock.

2.2.2. Preventive maintenance orders

Most systems are periodically inspected by engineers, either engineers of Vanderlande Industries or the customers' own engineers. At such an inspection, the engineers may conclude that some parts of the system need to be replaced. This does not mean that the system is about to break down, but to improve the quality of the system and to prevent system breakdowns in the future it is advisable to replace these parts within a certain timeframe.

If the customer agrees with the replacement of these parts, and the needed spare parts are not in stock at the system location, an order is placed at the service department of Vanderlande to supply these parts. Based on the lead time given by this department, the customer plans the preventive maintenance action, which may include hiring engineers and/or planning a (partial) system shutdown. In general the engineers foresee the failure of a spare part many weeks in advance, so the preventive maintenance activity can wait until all parts are ordered and delivered within the given lead times.

Most spare parts for a preventive maintenance order are sent to the customer through the SCW; in some cases the parts are directly shipped to the customer, e.g. if the supplier is abroad and close to the customer. Of course this is only possible if all the desired spare parts origin from one supplier. In case the SCW does not have all the desired spare parts on stock, the service department orders these parts at the suppliers, either an external supplier or the factory of Vanderlande. The supplier ships the parts to the SCW, where the customer order is consolidated until all parts are available. Figure 5 shows an overview of the supply of preventive maintenance orders.

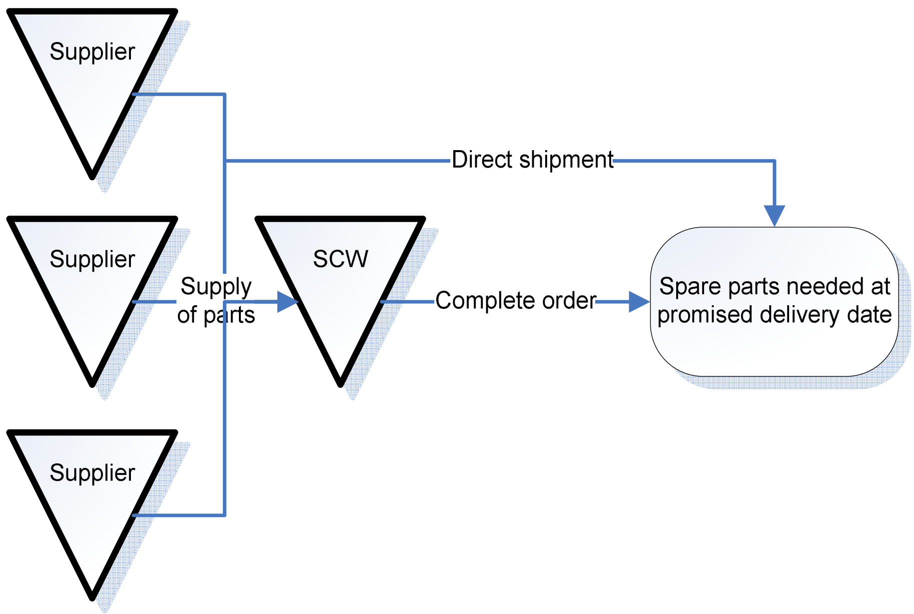

Figure 5 Overview of preventive maintenance orders

Performance

The preventive maintenance activity is planned based on the lead time of the order consisting of the needed spare parts. This means that the activity will take place if all spare parts are delivered at the customer. When the maintenance activity is planned, the customer may take preparations like planning a (partial) system shutdown and/or hiring engineers. Partial delivery of preventive maintenance orders is not useful, since all ordered parts are needed for the maintenance activity. We therefore want to be able to deliver as much orders as possible completely at or before the promised delivery date. This is a special case of the order fill rate performance measure as described in the previous section; we do not strive to fill as many complete orders as possible *from stock*, but we want to fill as much orders *within the given timeframe*. This timeframe is currently dependent on the lead time of the spare parts that are ordered.

2.2.3. Replenishment orders

In the previous sections we have already discussed that most customers have a stock of (the most critical) spare parts at or near the system location. Most customers are themselves responsible for the replenishment of this stock, except in case they have a consignment service agreement with Vanderlande. In the case of a consignment agreement, Vanderlande remains owner of the stock at the customer and takes care of the spare parts management at the customer. Spare parts are picked from the stock at the system location in case they are needed for corrective or preventive maintenance. The customer then decides if he wants to replenish the used parts. Most replenishment orders are the result of one or more maintenance activities, but not all maintenance activities lead to such orders. Customers may also decide to change the basic stock levels of the stock at the system location, but this does not occur often. As for preventive maintenance orders, most replenishment orders are sent to the customer through the SCW. Figure 6 shows an overview of the supply of replenishment orders.

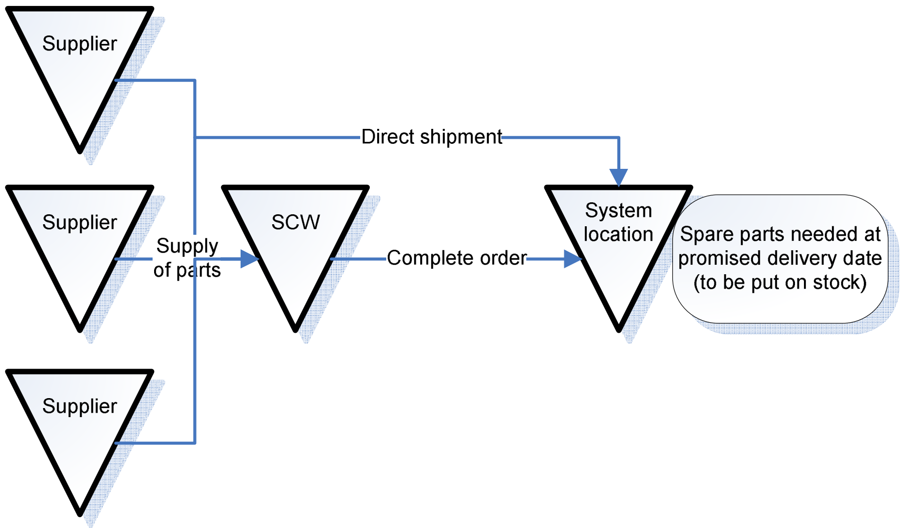

Figure 6 Overview of replenishment orders

Performance

Since replenishment orders are not immediately used by the customer upon delivery, it is not as critical as for preventive maintenance orders to deliver the complete replenishment order at once and in time. The ordered spare parts are put 'on the shelve' upon arrival at the customer location. There only occurs a real problem for the customer if one of the ordered spare parts is needed before it is delivered at the customer. From the point of view of the system, the most important performance is hence to deliver as much requested spare parts as possible at the promised date.

However, the customer expects that all spare parts are delivered at the promised date, and hence he is only satisfied if the complete order is delivered in time. Since the customer usually pays the transportation costs he is not willing to accept a partial delivery in case one or two parts are not available at the promised date.

To meet the desired performance by the customer, we employ the same performance as for preventive maintenance orders, namely order fill rate with respect to the promised delivery date. This means that we want to optimize the percentage of complete orders delivered to the customer at or before the promised delivery date.

2.2.4. Spare part packages

If a new system is sold to a customer, almost always a spare part package is sold. A spare part package consists of the most critical parts of a system. The determination of which parts are critical is executed in two steps. First the research & development (R&D) department selects which parts are classified as spare part. Next, the criticality of these spare parts is determined by a combination of experience and a classification based on the layout of a system made by the engineering department. The engineering department classifies subsystems as A (high priority), B (medium priority) or C (low priority), and this classification is copied to all the underlying items. Since a system consists of critical and non-critical items (e.g. an engine is critical, but a cover plate is not) the uniform classification of a system has the disadvantage that also non-critical items can have an A classification. This also means that spare parts can be classified as A in one system and C in another system. We clarify the classification and the sketched problem with a small example:

> **Example 1**
>
> Consider a check-in for a baggage handling system. The system consists of three check-in desks connected to one main conveyor that carries the bags to the sorting system. Because there are three check-in desks, each desk is classified as B, medium priority. All spare parts belonging to these desks adopt this classification. The main conveyor has priority A because the system is completely down if this loop breaks down. This means that all parts from the main loop have priority A. There are however some parts that are not critical for the functioning of the conveyor. If a classification per spare part would be made, this spare part would receive priority C, but since it is part of the main conveyor it has priority A.

The spare part packages are composed by the service department, in consultation with the customer. In general, items for spare part packages are not picked from stock. This only occurs if there is sufficient stock left to cover other service orders or when the package is needed by the customer short notice.

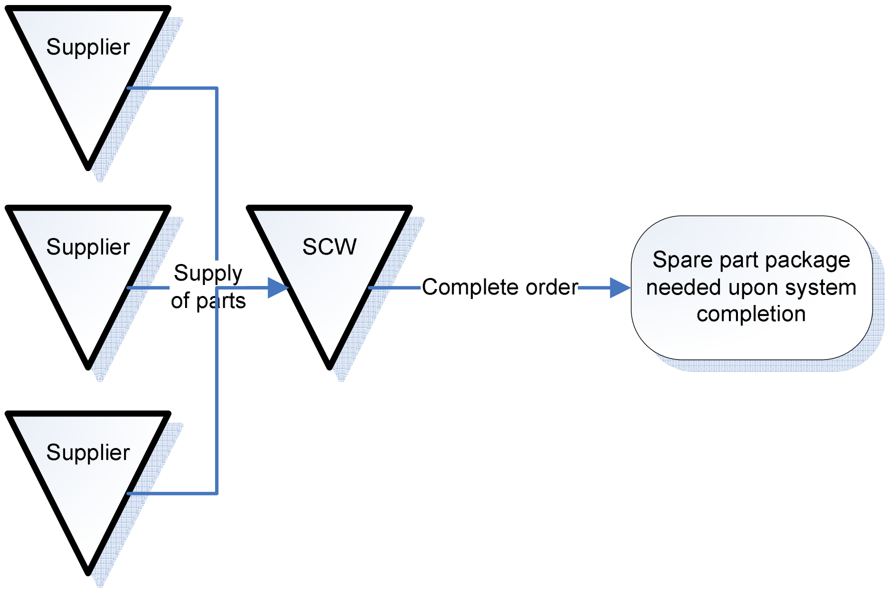

Figure 7 Overview of spare part packages orders

Performance

The most important performance for spare part packages is the availability of the package to the customer upon system completion. When the customer starts to use the system, he wants to have disposition of the most critical spare parts, so that he can quickly replace parts in case a failure occurs. Partial delivery of a package is not a problem as long as the last shipment arrives at the customer at or before the system completion. As for the previous demand streams, order fill rate with respect to a certain date is the most important performance measure, but in most cases the timeframe is considerably longer than for the previous demand streams since spare part packages are generally ordered well before system completion.

2.2.5. Revisions, modifications and retrofits (RMR)

The final demand stream consists of three types of large maintenance activities. Most of these activities are carried out outside the service department, but since they are service related we include them in this section. Below each category of RMR activities is discussed in more detail. In general RMR activities are characterized by large sets of spare parts to review (a part of) the system of the customer. They are often ordered well in advance and hence the needed spare parts are not picked from stock but ordered at the suppliers and consolidated at the SCW.

Revisions

Revisions are maintenance activities during which a large amount of spare parts are replaced in one activity. Usually this concerns moving parts or parts that can wear out, such as belts. For a couple of large customers, engineers of Vanderlande plan revisions together with the customer, but in most cases the customer requests a revision at its customer center of Vanderlande. Revisions are mostly part of a maintenance plan and are planned well in advance, hence the spare parts can be ordered at the supplier or the factory without delaying the planning and do not need to be supplied from a stock.

Modifications

If a customer wants to extend their system or change the functionality, Vanderlande can offer a modification of the system. Modifications can be seen as small projects, since a whole (sub-)system needs to be delivered. Since modifications are generally considered as small projects, the required parts are supplied by the operations department of Vanderlande, because that department takes care of normal projects as well. The service department and hence the central service warehouse are not involved in this flow.

Retrofits

Retrofits are carried out when parts in a system become obsolete; this means that the supplier does not supply the old parts anymore. During a retrofit, these old parts are replaced by new parts and if necessary the system is adapted to these new parts. Ideally the suppliers indicate months in advance that a spare part will become obsolete, so that Vanderlande has enough time to select alternatives.

If the alternative for an obsolete item is "Form-Fit-Function", i.e. it can be installed in the system instead of the old item without any further adaptations needed, Vanderlande may decide not to execute a full retrofit but gradually replace the old items (upon failure or preventive maintenance) with the new items. In case the system has to be adapted, a retrofit is offered to the customer. In that case, all relevant old items are replaced by the new items in one activity.

Performance

For all categories holds that the delivery of the desired spare parts coincides with a service activity in which the parts are actually being replaced in the system. This can be compared with preventive maintenance orders, so after a delivery date is promised, the customer may take preparations like the planning of a (partial) system shutdown to carry out the maintenance activity. To prevent that the customer shuts down its system to no purpose, it is important to deliver all parts in time, i.e. before the planned RMR activity. Hence, order fill rate with respect to the planned date of the RMR activity is the important performance measure. We must note that this date is usually known several months in advance, so there is enough time to order all desired spare parts.

2.3. Demarcation

Our main interest in this research is the SCW. Hence we mainly focus on the demand streams which performance we can improve by using this warehouse.

2.3.1. Demand streams

We have already mentioned that orders for RMR activities do not need to be fulfilled from stock, since they are ordered months in advance. For spare part packages this usually applies as well. Only in some cases – i.e. when there is sufficient stock left and the package is needed on short notice – a spare part for a package is picked from stock. Taking the number of spare part packages into account (around 60 to 80 per year) we conclude that this does not influence the stock heights significantly since the total number of spare part orders per year is more than 3,500. Therefore we do not include RMR activities and spare part packages in the remainder of this research.

Emergency orders

Emergency orders are preferably supplied from the SCW, but these orders are very rare. Based on experience and some historical data we have found out that at most two or three emergency orders per month occur. This is less than 0.5% of the total amount of orders, hence these orders merely influence the total demand. However, if an emergency orders arrives at the SCW, we want to deliver it from stock immediately to limit the possible downtime of the customers system. Therefore we have analyzed the emergency orders further.

We have some data available on emergency orders in 2007 and 2008. In total, we have 34 emergency orders, all of one order line. Only one spare part has occurred twice in an emergency order, all other spare parts are demanded once in an emergency order. The target performance is order fill rate, which comes very close to order line fill rate, as discussed in section 2.2.1.

Since the emergency orders are quite rare, and it is not possible to predict which spare parts are demanded in these orders based on historical data, we believe it is not possible to model this demand stream. Even if we would limit the spare parts that could be demanded in emergency orders to 1,000, we need to put (possibly more than one) spare part of nearly each type on stock in order to be able to achieve an order (line) fill rate of close to 100%. Since in general not more than 30 to 40 different spare parts are demanded in emergency orders per year, this means that the majority of this stock is never used.

In practice we will ignore emergency orders as if they never occur. In case an emergency order does occur, there are three options:

1) The spare part is on stock (following the model based on the other demand streams) in a sufficient amount. Even if the stock is allocated to another order, the spare part is picked for the emergency order and immediately shipped to the customer.
2) The spare part is not on stock at the service central warehouse, but on stock at either the distribution center of Vanderlande or at the supplier. In both cases the spare part is immediately shipped to the customer.
3) Both the service central warehouse and any other stock location do not have the spare part on stock. The spare part is ordered at the supplier with great emergency and is shipped directly to the customer as soon as possible; in the mean time, temporary solutions are found or (in worst case) the customers system is down until the spare part arrives.

Only in the third option the emergency order fill rate is 0%. Analysis of past emergency orders shows that around 50% of all emergency orders can be filled according to option 1. Data on fill rates of the other options is not available.

Service orders

The focus in the remainder of this research is therefore on the preventive maintenance and replenishment orders, which we call service orders when observed together. These demand streams form the largest part of orders handled by the service department, and the SCW can be used to improve the performance on these demand streams. In case an emergency order occurs, we use the available stock to try to fill the emergency order, but we do not include these orders in our model.

Currently the two main demand streams are not separated at Vanderlande. Detailed analysis of some service orders shows that in many cases both demand streams occur in a single service order. This means that at an inspection, some parts are replaced and then ordered for replenishment, and other parts are ordered for a future preventive maintenance activity.

Due to the way the preventive maintenance and replenishment orders are created, the composition of orders of these two types is much alike. Because no separate data on these demand streams is available as well, we base our model on the two demand streams combined, denoted as service orders. In section 6.5 we discuss the consequences of separating the service order demand stream into preventive maintenance and replenishment orders.

2.3.2. Performance

Currently, the most important performance measure at Vanderlande is on-time-delivery of customer orders, with respect to the complete order. The delivery date of the order (on which the on-time-delivery is based) is determined through the replenishment lead time of the spare parts of that order; if the longest replenishment lead time of all spare parts is e.g. four weeks, the delivery date is set at least four weeks from now; in the actual delivery date the shipment time is included as well.

From the previous sections we have learned that in all demand streams we strive to fill as much orders as possible completely. The percentage of orders that is filled completely is referred to as "order fill rate". Order fill rate calculates the percentage of orders that is filled completely *from stock*. However, for the two most important demand streams we strive to fill as much orders as possible completely *within the given timeframe*. We define this performance measure as *time-based* order fill rate with respect to a given timeframe. We define one desired timeframe for all orders within a demand stream. Therefore the timeframe becomes independent of the replenishment lead time of the spare parts.

Besides, we define the timeframe as the time between the placement of the order at the team Worldwide by the customer center and the moment that the complete order is ready to be shipped to the customer at the SCW. Hence, we do not take the handling time of the customer center and the transportation time from the SCW to the customer into account. In practice this implies that the actual timeframe – so including delivery to the customer – for customers in Europe is smaller than for customer in e.g. China due to longer shipment times.

These two adjustments enable us to define the desired performance of the central service warehouse as:

> *Ensure that a certain percentage of all service orders (e.g. 80%) is ready to be shipped to the customer within a fixed timeframe (e.g. three weeks), measured from the moment the customer center places the order.*

2.3.3. Two types of spare parts

Within the spare parts supply chain we recognize two types of spare parts. Around 60% of all spare parts are 'Vanderlande items', which means that they have a unique number within the company. This number enables us to retrieve data for each spare part on e.g. past orders. The remaining 40% are 'resale items'. For these items no unique number exists within the company so it is nearly impossible to retrieve historical data on them. Resale items are always bought at external suppliers; Vanderlande items are either produced at the factory or bought at external suppliers.

Because we are not able to retrieve historical information on resale items, we do not include these items in this research. At this moment Vanderlande is starting a project to define unique numbers for resale items. These items can be included in our model in the future.

2.4. Current performance

Currently, Vanderlande only measures the on time delivery for service orders, with respect to the promised delivery date to the customer. Figure 8 shows a graph of this on time delivery from January 2007 to May 2009. The straight line represent the trend in on time delivery.

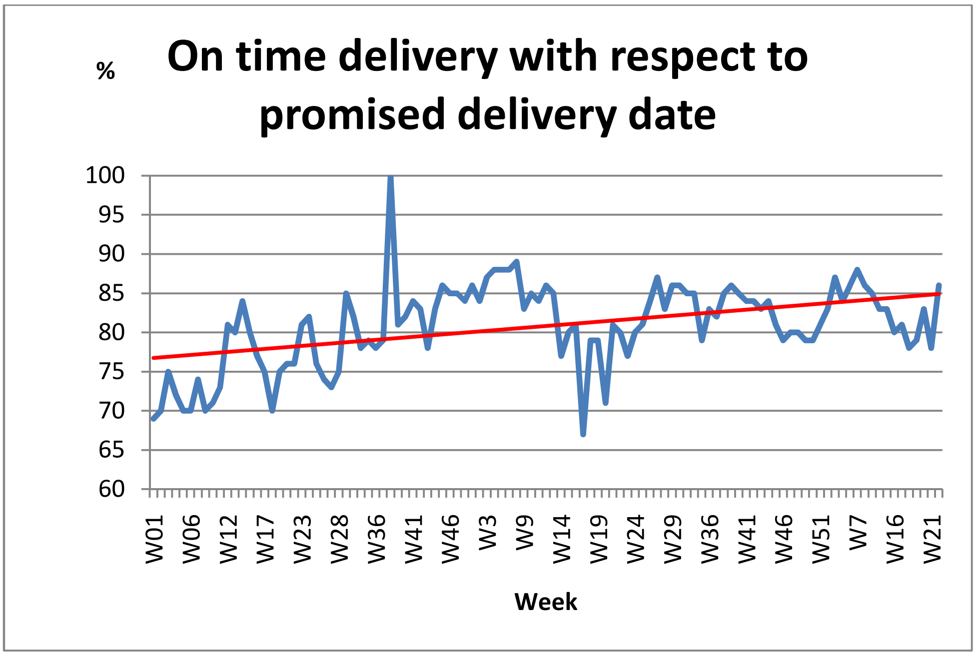

Figure 8 On time delivery of service orders

The average on time delivery is around 80%, but from this graph we cannot deduce within which timeframe this on time delivery is achieved. From April 2008 onwards more detailed information on each service order is recorded. More than 5,300 service orders were shipped from April 2008 to August 2009. The on time delivery, calculated by comparing the promised ship date to the actual ship date, is 83.3% in this period. 60.5% of all service orders was shipped within fifteen workdays from the placement of the order. Other percentages are shown in Table 1.

Timeframe	Percentage of complete orders within Timeframe
10 days (2 weeks)	50.5%
15 days (3 weeks)	60.5%
20 days (4 weeks)	69.9%
25 days (5 weeks)	79.3%
30 days (6 weeks)	85.7%
40 days (8 weeks)	93.5%
50 days (10 weeks)	96.9%

Table 1 Percentage of service orders delivered within timeframes

This performance is achieved partly due to the available stock at the central service warehouse. Currently 632 items have a positive stock height, corresponding to a total value of nearly € 185,000. The annual holding costs were around € 37,000 in financial year 2009 – i.e. April 2008 to March 2009 – but this includes temporary holding of spare parts that are usually not held on stock as well.

2.5. Conclusion

Based on the characteristics of each demand stream and the desired performances we have selected two demand streams – preventive maintenance and replenishment orders – for which the SCW can be effectively used to improve the performance of spare parts supply to customers. Due to data unavailability and the fact that those two demand streams are much alike, we take these two demand streams together and call them service orders.

We have defined our desired performance as time-based order fill rate, which means that we want to maximize the percentage of customers orders that is completely ready to be shipped to the customer within a given timeframe. The timeframe that Vanderlande has in view is three weeks (fifteen workdays), currently around 60% of all service orders is ready to be shipped to the customer within this timeframe.

In the next chapter we define more characteristics of the demand stream we observe, and of the spare parts supply chain in general. Using this extensive list of characteristics we conduct a literature research to find suitable models.

3. Modeling

Before we start defining a basic inventory model, we define the characteristics of the spare parts supply chain and the demand streams that we observe. Based on these characteristics we have conducted a literature research to find suitable inventory models. The next section discusses the most important characteristics. In section 3.2 a report of our literature research is given. Section 3.3 discusses how we intend to apply the found models to the supply chain of Vanderlande.

3.1. Supply chain characteristics

In (Hoving, 2008) an overview of characteristics of spare parts inventory control problems is defined. We use his approach to define the characteristics of the spare parts supply chain of Vanderlande. The complete list of characteristics is found in Appendix F. In this section the most important characteristics are discussed.

Demand characteristics per spare part

The demand for spare parts is characterized by a low frequency and generally low order quantities per spare part. Over 65% of all known spare parts has not been ordered since 2001. Around 85% of all spare parts that has been ordered, has been ordered at most eight times per year on average since 2001. This means that only around 5% of all known spare parts has been ordered more than eight times per year on average.

Each spare part can be ordered by more than one customer. Of course there may be certain parts that are only present in one or a few systems, but theoretically there is no limit on the number of customers that can order a spare part.

Order characteristics

In many cases, customers do not order a single spare part, but they order larger quantities of a type of spare part and/or multiple types of spare parts. Each type of spare part creates an order line in a customer order. On average a customer order consists of 2.1 order lines, but over 55% of all orders consists of only one order line.

The customer order quantity per order line depends on the type of spare part; small parts like washers or screws are obviously ordered in larger quantities by customers than motors and belts. Table 2 shows the percentage of all order lines that has at most a certain order quantity.

Order line quantity at most:	(Cumulative) percentage
1	31.2%
2	47.5%
3	52.6%
4	58.4%
5	64.3%
10	76.1%
25	84.5%
50	90.9%

Table 2 Customer order quantities

The average order quantity for all spare parts equals 41, but this includes large order quantities up to 50,000. As shown in the table, over 90% of all order lines has an order quantity of at most 50.

Other spare part characteristics

We make a number of assumptions on the characteristics of spare parts, because there is no detailed information available or to simplify modeling.

- All spare parts are critical; i.e. failure of any spare part causes failure of the customers' system, and hence all spare parts have the same priority in the delivery process. Because we focus on service orders and not on emergency orders it is acceptable to assume an equal priority to all customer orders.
- There is no obsolescence. This means that parts do not become obsolete through revised versions or replacements with other, similar, spare parts. In practice spare parts can be (and are) replaced by newer versions, possibly causing the remaining inventory to be useless.
- We assume a single-indenture model, hence we do not take the relation between spare parts following construction of a system into account.

Supply chain characteristics

In our research we observe a single-echelon, single-location supply chain in a static environment. We only look at the service central warehouse and observe the demand as if it origins from one customer. The demand rate for spare parts does not change over time due to e.g. seasonal influences.

3.2. Literature

Based on the characteristics as described in the previous section and the desired performance as mentioned in section 2.3.2 we executed a literature study for suitable models. The most important characteristic on which we selected suitable models was the order fill rate that forms the desired performance measure. The approach of the literature study is found in Appendix G (in Dutch), below we discuss the most interesting articles.

The main performance measure in our research is the (time-based) order fill rate, therefore we have focused our literature research on order fill rate models. These models are mostly described as assemble-to-order models, since an order of multiple different spare parts can be compared to an end-product that needs to be assembled from different parts. One of the first articles that discusses order fill rate is (Song, 1998). It considers a base-stock inventory system, i.e. an inventory system where every used product is immediately reordered at the supplier, with constant lead times and a multivariate compound Poisson process to model the demand of orders. The possible order compositions are known in advance, as are the parameters for their arrival process. The order quantity per order line is determined by a positive integer random variable.

In (Song, 2000) the same model is observed, but now for a batch-ordering inventory policy at the warehouse. This means that items are replenished by an integer multiple of a fixed order quantity if the inventory level drops below the reorder point. In (Song, 1998) the Poisson arrival process of orders is used to compute the order fill rate exactly; in (Song, 2000) these results are used to obtain the order fill rates in case of batch ordering.

Both articles of Song assume constant lead times for the replenishment of spare parts at the warehouse. With the same model assumptions – i.e. compound Poisson arrival process for demand and order line quantity by a random integer variable – (Lu, et al., 2003) expands the base-stock inventory policy model with stochastic lead times and (Zhao, 2008) treats the batch ordering inventory policy.

(Song, et al., 2002) and (Yao, et al., 2006) include stochastic replenishment lead times but both assume unit demand per order line. (Song, et al., 2002) observe a single product and a base-stock policy for each part, while (Yao, et al., 2006) look at both a base-stock and a batch ordering replenishment policy in a system with multiple products.

(Lu, 2008) presents a generalization of the model of (Lu, et al., 2003) by not assuming a compound Poisson arrival process, but a general renewal process for the arrivals of orders. Like in (Lu, et al., 2003) and (Zhao, 2008) the lead times for replenishment at the warehouse follow a general distribution. (Lu, 2008) gives approximations on order fill rates and average inventory.

All observed articles assume backlogging of the whole order in case an order cannot be filled completely from stock. Parts that are available are committed until the whole order is complete. Both orders and backorders are filled on a first-come-first-served basis. All articles observe products (i.e. orders) whose compositions in terms of parts are known in advance, together with the arrival rates of all products. Parts can be included in multiple products or orders in case multiple products are allowed in the model.

3.3. Approach

The main disadvantage of the models described in the previous section is that they all assume a finite set of given order compositions. This means that we know in advance which combinations of parts can be ordered together in one order. In our situation we do not have this knowledge and due to the large amount of spare parts, we suspect that it is very hard to determine fixed order compositions. Analysis of the data on service orders shows that a specific order composition seldom occurs more than once. Therefore all these models cannot be used in practice in our situation.

For our basic and advanced model we therefore turn to models that observe the items individually. Combining the performance of individual spare parts with some general characteristics of customer orders, we deduce a performance measure based on customer orders. For individual spare parts there are two relevant performance measures that can be compared with our desired performance measure – time-based order fill rate – as defined in subsection 2.3.2.

- **Item fill rate:** The percentage of all ordered items that is immediately available from stock when the order arrives. The *time-based* item fill rate gives the percentage of all ordered items that is available within the given timeframe.
- **Order line fill rate:** The percentage of all order lines that is completely and immediately available from stock when the order arrives. The *time-based* order line fill rate gives the percentage of all order lines that is completely available within the given timeframe.

To create a complete inventory model we start with a basic model that uses order line fill rate as performance measure. This basic model is introduced in Chapter 4. This model is extended to a complete inventory model best suited to the situation at Vanderlande. This complete model is explained in Chapter 5.

4. Order line fill rate model

In this chapter we design a basic model where order line fill rate is the performance measure of use. We answer the third research question:

3. On which assumptions is the basic inventory model based and how do we determine optimal stock values from this model?

First we give some basic definitions and notations in section 4.1. Section 4.2 discusses the most important assumptions, after which we define the model in section 4.3. The basic model is verified in section 4.4.

4.1. Definitions and notations

We introduce some general definitions and notations for inventory models:

- **Backorders (BO):** The total amount of spare parts from customer orders that have not been filled. We assume that if a customer order cannot be filled immediately, it is completely backlogged.
- **On hand inventory (OH):** The total amount of spare parts that is physically on stock.
- **Inventory level (IL):** The inventory level is equal to the on hand inventory minus all backorders, i.e. IL=OH-BO. If the inventory level is positive, it denotes the amount that is freely available to fill new customer orders. If it is negative, it means that the number of backorders – i.e. the total amount of spare parts not yet delivered to customers – is larger than the on hand inventory.
- **Outstanding orders (OO):** The total amount of spare parts 'on order', i.e. the amount of spare parts in replenishment orders that is not available yet.
- **Inventory position (IP):** The amount of spare parts that is available if all outstanding orders from suppliers have arrived, i.e.: IP=IL+OO.

Throughout our research these are all stochastic variables. The notation Pr(IL=j) hence stands for the probability that the stochastic variable IL (denoting the inventory level) equals j. Besides, in our model, the following notations are used. When applied to a spare part i, a subscript i is added to the parameter notation.

- **Reorder point (R):** The reorder point determines when to place a replenishment order at the supplier. If the inventory position has a value of R or lower, a replenishment order is placed.
- **Replenishment order quantity (Q):** If a replenishment order is placed, the size of this order is equal to Q or an integer multiple of Q.
- **Demand rate (λ):** The daily demand rate of a spare part denotes the average number of customer order lines per day that contain this spare part. Note that this demand rate is not equal to the demand rate of customer orders, since more than one spare part exist and a customer order can consist of more than one order line.
- **Customer order quantity distribution (F):** This is a stochastic variable that describes the probability distribution of the quantity that a customer orders of a spare part. It is a discrete distribution which is determined empirically based on historical data. In section 6.1.2 is worked out into detail how this distribution is determined.

- **Replenishment lead time (L):** In the basic model this is a constant, indicating the lead time of a replenishment order in days. In the next chapter, L becomes a stochastic variable to include variability in the replenishment lead time.
- **Demand during replenishment lead time (D(L)):** In our model we are interested in the total demand for spare parts during the replenishment lead time. This is denoted with D(L) where L is (the stochastic variable denoting) the replenishment lead time.
- **Holding costs (h):** All costs for keeping one spare part on stock for one day are combined in the holding costs. The holding costs are measured in Euros per day.
- **Fixed order costs (C):** For each replenishment order, certain fixed costs are made like the handling of an incoming order. These costs are given in Euros.

4.2. Modeling assumptions

In defining our basic model we make a number of assumptions. Some assumptions are valid through the whole research, others are relaxed or changed in the next chapter. In this section we explain for each assumption why we assume this and, in case the assumption is relaxed or changed in the advanced model, how we intend to do so.

4.2.1. General assumptions

The assumptions mentioned in this subsection are valid through the whole research.

- *First-Come-First-Serve policy for customer orders:* We assume customer orders are handled following a First-Come-First-Serve (FCFS) policy. This means that customer orders are filled in the same order as they have arrived. This is clarified in Example 2. In practice only emergency orders are favored compared to orders that have arrived earlier, but less than 0.5% of all customer orders is an emergency order.

> **Example 2**
>
> The available stock for a spare part is 10 and an arriving customer demands 15 pieces of this spare part. All subsequent customers have to wait until the 'large' customer order has been fulfilled, even if these subsequent orders for this spare part are smaller than 10. After arrival of a sufficiently large replenishment order, the large customer is delivered first. Any subsequent customer order is filled if the remaining stock is large enough.

- *Compound Poisson arrivals for each spare part:* We assume that customer orders for a spare part arrive following a compound Poisson process. This means that the inter arrival times – i.e. the time between two consecutive customer orders for this spare part – are exponentially distributed with parameter λ and that the order quantity of each customer order is defined by the stochastic variable F which follows a discrete distribution. Exponential inter arrival times is a widely accepted assumption for modeling spare parts demand. We assume that the arrival rate of customers, λ, is static; this means that it is not changed due to seasonal or other influences.
  Since nearly 70% of all order lines has an order quantity of more than one, we need to include the customer order quantity in our arrival model. To do so we include the stochastic variable F that denotes the customer order quantity.
- *Holding costs per spare part are calculated based on cost price:* We assume that the holding costs per spare part are calculated as a percentage of the cost price. In practice this is only

partly true. The holding costs are actually calculated as a combination of a percentage of the cost price and a fee for the used shelve space in the warehouse. The latter is not related to the cost price. In section 6.1.2 we further explain how we approximate the holding costs by assuming that all holding costs are related to the cost price.

- *(R,Q)-policy for replenishment:* In the replenishment of the central service warehouse by the suppliers we assume an (R,Q)-policy. This means that when the inventory position drops to or below the reorder point R, an integer multiple of Q parts is ordered at the supplier such that the resulting inventory position is larger than R. Q is called the replenishment order quantity.
  Currently for nearly all spare parts a base-stock policy is used, which means that Q=1 and each customer order implies a replenishment order of the same amount to the supplier. Since this may lead to a large number of replenishment orders, Vanderlande wants to apply the (R,Q)-policy in the future to reduce the number of replenishment orders. Within the (R,Q)-policy, Q=1 is still possible so the base-stock policy is covered by this assumption.

Restrictions on reorder point and replenishment order quantity

The replenishment order quantity Q is subjected to some restrictions by Vanderlande and/or the supplier of the spare part:

- Q has to be at least equal to the minimum order quantity (MOQ), i.e. the minimum value that needs to be ordered at the supplier. This quantity is imposed by the supplier. If no constraints on the minimum order quantity exist, we set MOQ=1.
  $$Q \geq MOQ$$
- Q has to be an integer multiple of the fixed order quantity (FOQ). A fixed order quantity is imposed by the supplier e.g. in case spare parts are only sold in boxes containing more than one part. If no constraints on the fixed order quantity exist, we set FOQ=1.
  $$Q = n \cdot FOQ, \quad n \in \mathbb{N}^+$$
- To keep the total inventory at an acceptable level and to prevent obsolescence of spare parts – which is not explicitly included in our model – Vanderlande strives to order a spare part at least once per three months (i.e. thirteen weeks, 65 workdays). This means that the replenishment order quantity may not be too large.
  We impose these restrictions through the expected customer order quantity (E[F]) and the arrival rate (λ). Per day, on average λE[F] spare parts are ordered. Based on our restrictions as formulated above, we impose the following restrictions on the replenishment order quantity Q:
  $$Q \leq 65\lambda E[F]$$

Beside the restriction on the replenishment order quantity, Vanderlande also wants to imply a restriction on the reorder point R.

- If no stock is held for a spare part, its reorder point theoretically equals –Q. This means that the inventory position is never positive. For values of Q larger than 1 this can result in the situation where there is no stock (the inventory level is always at most equal to the inventory position, so it is never positive as well) but an arriving customer order does not imply a replenishment order since it does not cause the inventory position to drop below the reorder point. In reality, any customer order that cannot be filled from the available stock or any outstanding replenishment orders, implies a new replenishment order. This implies that the reorder point is always at least equal to -1, so if the inventory position becomes negative, immediately a replenishment order is placed. With regard to the replenishment order quantity Q, we distinguish a number of scenarios in case a customer arrives and no stock is available:

- The customer is asked to order (a multiple of) the replenishment order quantity, ensuring that there are no spare parts left that need to be put on stock.
- A smaller amount is ordered at the supplier, against (in most cases) higher costs.
- The customer order quantity is picked from the stock at the distribution center of Vanderlande (normally used for manufacturing and non-service related orders), hence no replenishment order is placed and no stock is left.
- The remainder of the replenishment order with size Q is put on stock at the distribution center of Vanderlande.
- The remainder of the replenishment order with size Q is put on stock at the service central warehouse.

In the first four scenarios, the customer or replenishment order quantity is adapted such that there is no remaining stock or the remainder of the replenishment order is sold to a third party. These scenarios show most resemblance with a base stock policy with R=-1, i.e. an (-1,1) policy. In such a policy no stock is held and each customer order triggers a replenishment order of the same size as the customer order.

The fifth scenario is modeled as an (R,Q)-policy where Q is the originally calculated value of the replenishment order quantity and R is at least equal to -1; hence each customer order is either filled by stock or outstanding replenishment orders, or it triggers a new replenishment order.

Since it is not guaranteed that one of the first four scenarios is possible, we need to assume a (R,nFOQ)-policy with R≥-1. The determination of the replenishment order quantity reduces to finding the integer value n such that:

$$\begin{aligned} nFOQ &\geq MOQ \\ nFOQ &\leq 65\lambda E[F] \end{aligned} \tag{4.1}$$

The method to determine this value n and the reorder point R is discussed in section 4.3.2.

4.2.2. Basic model assumptions

Beside the general assumptions mentioned in the previous section, we make a number of additional assumptions in our basic model. Each of these assumptions is changed or released in the advanced model.

- *Fixed lead time per spare part:* For each spare part a basic (fixed) replenishment lead time is defined in the data. In practice this lead time varies because suppliers are not completely reliable in their deliveries or the basic replenishment lead time cannot be met due to e.g. a temporary lack of capacity at the supplier. Overall around 80% of all replenishment orders placed by Vanderlande is delivered in time, so most orders are delivered within the fixed lead time that is assumed in our basic model. In our advanced model we include the unreliability of the suppliers by looking at the probability distribution of the delay in the delivery of replenishment orders per supplier.

- *Suppliers deliver replenishment orders on all days:* Our basic model is a continuous review model. This means that both customer orders and replenishment orders can arrive at any time. To fit this model we need to assume that suppliers can deliver at any time. In practice 10 to 20% of all suppliers only deliver replenishment orders once or twice per week; all other suppliers deliver once per day. In the next chapter we discuss how to include these fixed delivery days in the model.
- *Order line fill rate 'off-the-shelf' as performance measure:* Our basic model is based on order line fill rate as performance measure, and we do not take the desired delivery timeframe into account. We hence calculate an 'off-the-shelf' fill rate, i.e. the percentage of customer order lines that is filled immediately from stock. This fill rate is calculated for each spare part individually. From this basic model we deduce a time-based order line fill rate for each spare part in the next chapter. This advanced model is then used to derive a model for calculating the time-based order fill rate for customer orders in general.

4.3. Calculations

Our basic model is based on calculating the 'off-the-shelf' order line fill rate for each spare part individually in a continuous review model with fixed lead times and no restrictions on the deliveries of replenishment orders by suppliers. In the basic model, all customer orders consist of one single order line, so a customer order refers to an order of one order line.

Before we start with our model we define a method to reduce the calculation times throughout the model. Therefore we reduce the spare part data as much as possible without deleting any data. In case a spare part is always ordered in e.g. multiples of 100, we create a 'new' spare part through the following procedure:

1) Divide all possible customer order quantities by 100
2) Divide the MOQ en FOQ by 100
3) Multiply the price of the spare part by 100

One unit of the new spare part hence equals 100 units of the old spare part.

In the remainder of this section we define the basic model. The first step is to define an expression that calculates the order line fill rate.

4.3.1. Calculating the order line fill rate

As stated in section 4.2.1, we assume an (R,Q)-inventory policy for the service central warehouse. The values of the reorder point R and the replenishment order quantity Q determine the order line fill rate we are interested in. This order line fill rate is denoted with OLFR(R,Q).

A customer order line can be delivered completely and immediately from stock if, at the moment the order arrives, the inventory level is at least equal to the ordered quantity by the customer. Hence if the inventory level has an arbitrary value j, a customer order can be filled if its quantity is at most this value j. The order line fill rate for a given inventory level j is:

$$OLFR\left(R,Q \mid IL=j\right)=\Pr\left(F\leq j\right) \tag{4.2}$$

To find the order line fill rate we sum over all possible values of the inventory level (Thorstenson, et al., 2008):

$$
\begin{aligned}
OLFR(R,Q) &= \sum_{j=-\infty}^{\infty} OLFR(R,Q \mid IL=j)\Pr(IL=j) \\
&= \sum_{j=-\infty}^{\infty} \Pr(F \le j)\Pr(IL=j)
\end{aligned}
\tag{4.3}
$$

The probability distribution of the stochastic variable F is known. The probability distribution of the inventory level is related to the probability distributions of the inventory position (IP) and the demand during replenishment lead time L – denoted with D(L). A full derivation is found in Appendix B.1.1, the resulting expression for the order line fill rate is:

$$
OLFR(R,Q) = \sum_{j=1}^{R+Q} \Pr(F \le j) \sum_{k=\max\{R+1,j\}}^{R+Q} \Pr(IP=k)\Pr(D(L)=k-j) \tag{4.4}
$$

Since the probability distribution of the inventory position is a uniform distribution (Axsäter, 2006) – as explained in Appendix B.1.1 – we only need to determine the probability distribution of the demand during replenishment lead time, denoted with the stochastic variable D(L). We have developed two methods to determine this probability distribution.

The first method is based on the exact calculation of the compound Poisson probability of the demand during replenishment lead time. This is achieved through convolutions of the probability distribution of the customer order quantity. The y-fold convolution of this distribution in the point k, indicating the probability that y customers order k spare parts in total, is denoted with $f^y(k)$. We use a small case f to avoid confusion with the stochastic variable F.

The derivation of the order line fill rate is found in Appendix B.1.1, the resulting expression is:

$$
OLFR(R,Q) = \frac{1}{Q}\sum_{j=1}^{R+Q} \Pr(F \le j) \sum_{k=\max(R+1,j)}^{R+Q} \left[\sum_{y=0}^{\infty} \frac{(\lambda L)^y e^{-\lambda L}}{y!} f^y(k-j) \right] \tag{4.5}
$$

We refer to this calculation method as the “Compound Poisson method”.

The second method we developed is an approximative method that avoids the – in most cases extensive – calculations of the convolutions $f^y(k)$. It is based on the mean and variance of the customer order quantity. These values are used to determine the parameters of a suitable probability distribution. The probability distribution of choice is determined by looking at the variance-to-mean ratio of the demand during replenishment lead time: the variance divided by the mean. If this is smaller than one, a binomial distribution is best suitable (Law, et al., 2000). For a variance-to-mean ratio close to or equal to one, the Poisson distribution is suitable (Axsäter, 2006; Law, et al., 2000). Finally, for larger variance-to-mean ratios, the negative binomial distribution is best suitable (Law, et al., 2000). We refer to this method as the ‘two-moments-method’, since we use the first two moments of the demand during replenishment lead time.

The basis of the order line fill rate calculation in this method is equation (4.4), since we only change the probability distribution of the demand during replenishment lead time compared to the previous method. For all three possible distributions we need to calculate the mean and variance of the demand during replenishment lead time, E(D(L)) and var(D(L)) respectively. In Appendix B.1.2 we deduce how the mean and variance of D(L) are calculated using the mean and variance of the customer order size, the customer arrival rate and the replenishment lead time. This calculation method is valid for all three probability distributions.

To find the actual expression of the order line fill rate we relate this mean and variance of D(L) to the input parameters of the probability distribution of choice. This calculation of the probability distribution of the demand during replenishment lead time is worked out in Appendix B.1.2.

Conclusion

In this subsection we have developed two methods to calculate the order line fill rate for given values of the reorder point R and replenishment order quantity Q. In the next section we discuss how to determine these values of R and Q so that a given target order line fill rate is achieved. The verification of both the compound Poisson method (equation (4.5)) and the two-moments method (using equation (9.1), (9.3) or (9.5) for the calculation of the probability distribution of the demand during replenishment lead time) is discussed in section 4.4.

4.3.2. Determining reorder point and replenishment order quantity

The goal is now to find the values of the reorder point R and the replenishment order quantity Q such that a given order line fill rate is achieved, preferably against the lowest total costs. This order line fill rate is denoted with $OLFR^{target}$. The costs consist of holding costs – the costs for keeping a spare part on stock – and order costs – the fixed costs per replenishment order, e.g. administrative costs.

C(R,Q) = Tot.holding costs + Tot.order costs

In determining the values for R and Q we hence want to solve the following optimization problem:

$$\begin{aligned}&\min C(R,Q)\\&\text{s.t. } OLFR(R,Q)\geq OLFR^{\text{target}}\\&Q\in\mathbf{Q}, R\in\mathbf{R}\end{aligned}\tag{4.6}$$

The set **Q** contains all integer values that satisfy the constraints on the replenishment order quantity, as stated in (4.1). The set **R** contains all integer values ≥ -1.

We discuss a number of approaches to solve this optimization problem:

- Enumerate all possible combinations of R and Q
- Add backorder costs and minimize total costs without the given constraints
- Relax the optimization problem
- Use a sequential method, i.e. first determine Q and use this value to determine R

Enumeration

Enumeration means that we enumerate all possible combinations of R and Q and find the combination that has the lowest costs. Especially for spare parts with a high demand rate and high customer order quantities, the number of possible combinations can be very large (up to 100 million for some cases at Vanderlande). Of course, many of these combinations are not interesting; e.g. if the average customer order quantity is 100, we intuitively know the combination of a reorder point of 1 and a replenishment order quantity of 1 is not interesting. However, in many cases there are still a lot of combinations left, for each of which the order line fill rate and corresponding holding costs have to be calculated.

A second disadvantage of this method is that it is not suitable to optimize the joint line fill rate of multiple items. If we want to calculate the reorder points and replenishment order quantities for more than one spare part, where the objective is to find the lowest costs at which a given total line fill rate is achieved, we need to observe all the combinations of all combinations of each spare part.

Add backorder costs

In many models (e.g. (Federgruen, et al., 1992), (Chang, et al., 2005)) the order line fill rate constraint is replaced by an addition to the cost function. For each spare part that cannot be delivered at the desired moment, penalty costs are awarded. These penalty, or backorder, costs are added to the cost function, and the objective becomes to minimize the total costs, with only constraints on the values of R and Q – i.e. $Q \in \mathbf{Q}, R \in \mathbf{R}$.

This method is not applicable in this case, since at Vanderlande no penalty costs exist when a customer order is not filled within the desired timeframe. With some customers there are agreements on the level of service, but there are no direct costs in case an order cannot be filled.

Relaxation

Another method to find an approximation of the optimal solution of an optimization problem is to find a relaxation of the optimization problem. Relaxation means that constraints are made less strict or are left out completely. A possible relaxation method is Lagrangian relaxation, as discussed (amongst others) in (Fisher, 1981). Lagrangian relaxation adds the constraints to the objective function using a Lagrangian multiplier, denoted with µ. The Lagrangian relaxation of the problem stated in (4.6) is:

$$\begin{aligned} &\min C(R,Q) - \mu\left(OLFR(R,Q) - OLFR^{\text{target}}\right) \\ &Q \in \mathbf{Q}, R \in \mathbf{R} \end{aligned} \tag{4.7}$$

The exact expression of the Lagrangian relaxation is found in Appendix B.3. Lagrangian relaxation is solved by taking the derivative of the new minimization function with respect to the original variables R and Q (Zwillinger, 2003). This derivative is hard to determine due to the composition of the cost and order fill rate functions. Besides, Lagrangian relaxation yields only an exact solution in case there are no restrictions on R and Q (i.e. $R, Q \in \mathbb{R}$). In our case, solving the Lagrangian relaxation could yield an infeasible result and further steps are needed to find a feasible solution, i.e. a solution of R and Q such that $Q \in \mathbf{Q}, R \in \mathbf{R}$.

Sequential method

All methods described above try to determine the reorder point and the replenishment order quantity simultaneously. This causes the methods to be laborious and/or complex, and possibly hard to execute when we want to observe multiple spare parts. We therefore developed a sequential method. In this method we first determine the replenishment order quantity Q, and use this value to determine the optimal value of the reorder point R with respect to that Q.

Determining the replenishment order quantity Q

The replenishment order quantity Q determines the amount that is ordered at the supplier in case the inventory position becomes equal to or less than the reorder point R. For some spare parts, the purchase (or cost) price of the spare part depends on the ordered quantity. In general any cost advantages of larger order quantities are passed on to the customer. Therefore we do not take these economies of scale into account in the calculation of the replenishment order quantity Q, and we only calculate a single value for the replenishment order quantity Q.

To seek a good value for the replenishment order quantity Q within the constraints imposed in (4.1), we use the Economic Order Quantity (EOQ) formula. This is a widely accepted method to determine the replenishment order quantity. It was developed by (Harris, 1913) and (Wilson, 1934). They prove that if the fixed ordering costs C, the rate of demand λ, the replenishment lead time L, and the cost holding costs of the spare part h are all constant, and no partial deliveries are allowed, the EOQ is the optimal replenishment order quantity.

The EOQ formula is given by:

$$Q^* = \sqrt{\frac{2C\lambda EF}{h}} \tag{4.8}$$

However, the EOQ formula does not necessarily lead to an integer value Q* and thus does not lead directly to a suitable replenishment order quantity Q. To achieve this, we use the following procedure to determine the replenishment order quantity Q. Remember that we actually seek for an integer n that satisfies the restrictions of (4.1):

1) Calculate the EOQ Q* using (4.8)
2) Round Q* to the nearest positive integer multiple of the fixed order quantity FOQ:
   $$n = \max\left\{1, \mathrm{round}\left(Q^* / FOQ\right)\right\}$$
3) If $nFOQ > 65\lambda EF$, $n := \left\lfloor 65\lambda EF / FOQ \right\rfloor$
4) If $nFOQ < MOQ$, $n := \left\lceil MOQ / FOQ \right\rceil$
5) Set $Q := nFOQ$ (4.9)

Determining the reorder point R

For a given replenishment order quantity Q, the corresponding reorder point R has to guarantee that the target order line fill rate is achieved. We therefore seek for a value R for which:

$$OLFR(R,Q) \geq OLFR^{\text{target}}$$

Since our objective is to minimize the total costs given that we obtain the target order line fill rate, we set the reorder point R to the smallest value that satisfies the inequality, because the total costs increase when R increases for a fixed value of Q. This is proven in Appendix B.2. Furthermore, in this Appendix we deduce the following recursive method to find this value R:

1) Set R = -Q, and define OLFR(-Q,Q) = 0
2) Increase R by 1 and calculate the new OLFR(R,Q) by:

$$OLFR(R,Q) = OLFR(R-1,Q) + \frac{1}{Q}\left[\sum_{j=1}^{R+Q} \Pr(F \le j)\Pr(D(L) = R+Q-j) - \sum_{j=1}^{R} \Pr(F \le j)\Pr(D(L) = R-j)\right] \tag{4.10}$$

3) If OLFR(R,Q)≥$OLFR^{target}$ and R≥-1, STOP, else go to Step 2.

We start with R=-Q since for this reorder point the order line fill rate is 0 and we can easily apply the recursive relation of (4.10) in this way. The third step ensures that the final reorder point is at least equal to -1.

4.4. Model verification

In the previous section we have designed a basic model to determine the order line fill rate for a given target fill rate. We have defined two possible methods to approximate the probability distribution of the customer demand during replenishment lead time, which is used in the calculation of the order line fill rate. These methods are the compound Poisson method, using equation (4.5), and the two-moments method, using either (9.1), (9.3) or (9.5), depending on the variance-to-mean ratio of the customer order quantity distribution.

In this section we verify both calculation methods by comparing the calculated order line fill rates to the resulting order line fill rates from a simulation study for a set of test instances. First we define how this set of test instances is composed in section 4.4.1. In section 4.4.2 we show how we carry out the simulation. In section 4.4.3 the results of our experiments are shown.

4.4.1. Test instances

Based on historical service orders we have compiled a set of test instances that covers most of the spare parts that are ordered by customers. We have defined four characteristics of spare parts for which we vary the parameter values:

- *Daily demand rate:* This indicates the daily arrival rate of customers, denoted with λ. This rate is the mean in the Poisson arrival process that describes the arrivals of customers.
  We choose seven values varying from 1 arrival per year (λ=1/260) to 221 arrivals per year (λ=0.85). These values are chosen based on analysis of the historical data on service orders such that they represent a wide range of spare parts. The arrival rates are shown in Table 3.

Daily demand rate
0.85 (=221/260)
0.75 (195/260)
0.35 (91/260)
0.25 (65/260)
0.077 (20/260)
0.027 (7/260)
0.0038 (1/260)

Table 3 Arrival rates for verification study

- *Replenishment lead time:* This indicates the basic replenishment lead time from the supplier. In the basic model, all replenishment orders have a fixed replenishment lead time which is equal to this value. Again, the used values are chosen based on historical data. We have four values as shown in Table 4.

Basic replenishment lead time
65 days
28 days
15 days
5 days

Table 4 Replenishment lead time for verification study

- *Customer order quantity:* The customer order quantity is determined by a discrete probability distribution. We have designed ten fictional probability distributions that are more or less similar to some empirical probability distributions that we have encountered in the historical data. An overview of the distributions is shown in Table 5, the complete distributions are clarified in Appendix C.1.

Distribution #	Mean	Standard deviation	Min	Max
1	100	0.00	100	100
2	3.83	3.53	1	20
3	6.09	3.97	1	20
4	10.52	3.32	2	20
5	9.51	12.93	1	100
6	20.97	24.13	1	100
7	145.94	107.50	1	500
8	113.61	160.79	10	1000
9	109.22	123.73	10	500
10	723.29	508.91	100	2000

Table 5 Customer order quantity distributions for verification study

- *Cost price:* The cost price influences the replenishment order quantity via the EOQ formula. We choose two different cost prices, €5 and €1000. These values are chosen to make a clear distinction between cheap and expensive spare parts. The fixed order costs are constant, € 20,- for all experiments. This value is based on the average handling costs of incoming replenishment orders, as charged by the third party that exploits the service central warehouse where these orders are received.

We calculate the reorder point and the replenishment order quantity following respectively (4.10) and (4.9), for three target order line fill rates: 90%, 95% and 99%. Each combination of demand rate, replenishment lead time, customer order quantity distribution, cost price and target fill rate is observed; hence we execute 1680 experiments in the basic model.

To calculate the order line fill rate we verify the two calculations methods mentioned earlier in this section – the compound Poisson method and the two-moments method – by comparing their results to the outcomes of the simulation study.

4.4.2. Simulation study

In our simulation studies we generate spare part orders with an exponential inter arrival time – using the demand rate of the spare part. The customer order quantity for each arriving customer is drawn from the corresponding customer order quantity distribution. When the inventory position drops below the reorder point, the calculated replenishment order quantity is ordered and arrives after the fixed lead time of the spare part. We simulate around 20,000 replenishment orders per spare part. The first 100 replenishment orders are considered as the 'warming up period' (Law, et al., 2000) and are not included in the performance measurement. Because we have encountered relatively large variations in the outcomes of the simulation study, we have decided to execute five independent simulation runs and take the average order line fill rate as outcome for our simulation study.

4.4.3. Results

The detailed results can be found in Appendix D.1. Here we discuss the most important results of the basic model verification. As stated in the previous subsection, we have verified two calculation methods. We first analyze the results per method, then we conclude with the comparison of the two methods.

We analyze the results by looking at the difference between the order line fill rate resulting from the calculation method and the order line fill rate that results from the simulation studies. We subtract the simulation result from the calculated order line fill rate, so if the outcome is negative, the calculation underestimates the simulation. A positive result means that the order line fill rate is overestimated by the calculation; this is less desirable, since the objective is to at least achieve the target order line fill rate. We refer to this as "the difference between simulation and calculation".

We also look at the absolute difference, to calculate the average deviation of the calculation compared to the simulation. Both the normal and the absolute differences are measured 'absolute', i.e. opposite to the relative difference; hence, if the calculation results in an order line fill rate of 80% and the simulation gives 85% as a results, the difference we speak about is -5%, with an absolute value of 5%. Furthermore we look at the standard deviation of the absolute difference as well. This is the 'average deviation from the average difference' between the calculation and the simulation results. Finally, we have constructed percentiles on the absolute difference between the order line fill rates from the calculation and the simulation to provide more insight in the spread of the differences.

In the remainder of this section we provide the results of all experiments taken together, and we discuss notable results of our detailed analysis. We have analyzed the results per parameter value – e.g. the results for all spare parts with a replenishment lead time of 65 days are grouped together. Our main performance measure is the absolute difference in order line fill rate between the calculation method and the outcome of the simulation. The quality of each calculation method is determined mainly based on this performance measure. We accept a calculation method if the average absolute difference is at most 0.1% and the maximum absolute difference is less than 1.0%. Due to the large number of experiments, the 95% confidence intervals on the difference between simulation and calculation all have a width of at most 0.004%. Therefore we have not included them in the discussion of the results below.

Measure	Compound Poisson method	Two moments method
Absolute difference simulation-calculation	0.054%	0.237%
Standard deviation on absolute difference	0.063%	0.315%
90%-percentile absolute difference	0.127%	0.485%
Maximum absolute difference	0.634%	4.395%

Table 6 Key results of basic model verification

The results of the compound Poisson method are much better than the results of the two moments method. At the compound Poisson method, the highest absolute difference between calculation and simulation occurred for a spare part with high replenishment lead time, high demand rate, high costs and customer order quantity distribution 7. In this order quantity distribution, the difference between the smallest and the largest possible order quantity is 500. The probability that the quantity equals 500 is small, but such an order has a large impact on the performance. If in the simulation study the resulting empirical probability of having an order quantity of 500 is slightly higher or lower, the order line fill rate is affected.

At the two moments method we observe very large differences of over 4% between the calculated order line fill rate and the result from the simulation study. The customer order quantity distribution mainly determines the quality of this method; for five distributions, the maximum difference is less than 1% - although the average difference is in all but one case still over 0.1% - and in four of the remaining distributions the maximum difference is at most 1.6%. The largest differences all occur for customer order quantity distribution 4. This is explained by the fact that the variance-to-mean ratio of this distribution is close to 1. Therefore the Poisson distribution is chosen as the probability distribution of the demand during replenishment lead time. Figure 9 shows the probability mass functions of both the original customer order quantity distribution and the Poisson distribution resulting from the two moments method.

From the graph it becomes clear that the approximative Poisson distribution differs from the original distribution. Especially the probabilities of high order quantities are larger in the original distribution. Thence the probability of having a high cumulative demand – i.e. multiple customers all order high quantities resulting in a high cumulative demand – is much smaller in the two-moments method than in practice. However, as the demand increases, this problem decreases. If more customers arrive during replenishment lead time, the probability of having all large order quantities is very small in the empirical distribution as well. Differences between this distribution and the approximative distribution are then less noticeable.

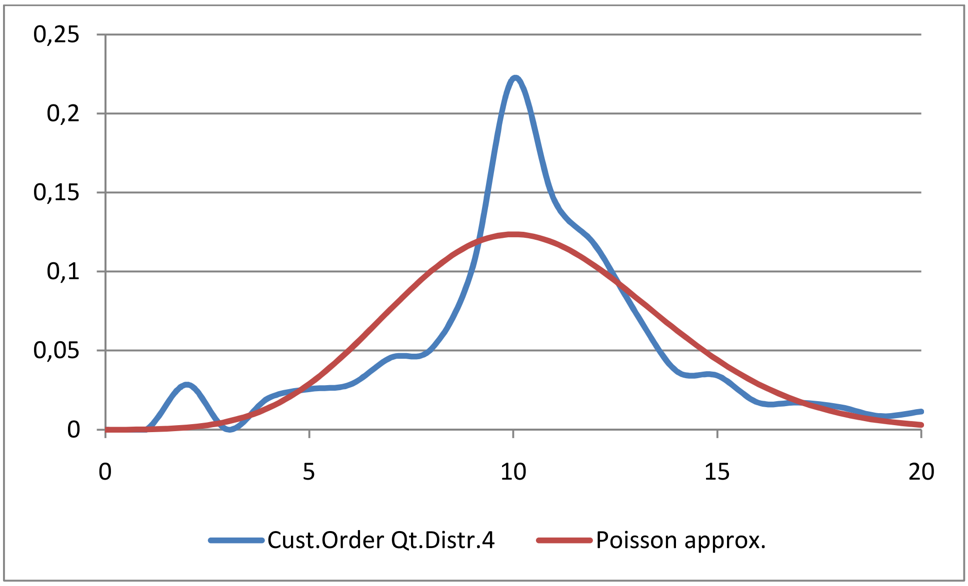

Figure 9 Probability mass function of customer order quantity distribution 4 and Poisson approximation

For the other customer order quantity distributions that score large differences the same occurs; the probability that a high order quantity occurs is higher in the empirical distribution than in the distribution following the two-moments method.

During our verification study we have measured the calculation time of both methods by looking at the time it took to calculate the probability distribution of the demand during replenishment lead time. For the compound Poisson method this took nearly 8 seconds. The two-moments method scored around 0.5 seconds faster. We expected a larger difference in calculation time, but apparently the calculation of the convolution of the customer order quantity distribution is handled well.

4.5. Conclusion

As we would expect, the compound Poisson method scores very good compared to the simulation. This is logical since in our simulation we use the same arrival rate and order quantity distribution as in this calculation method. The two moments method scores worse. Both the average absolute difference as the maximum absolute difference between simulation and calculation are above our threshold for acceptable solutions. However, the worst results are caused by a single customer order quantity distribution that performs much worse than other distributions in this method.

In this chapter we have defined a basic model that calculates the 'off-the-shelf' order line fill rate per spare part in a continuous review model using an (R,Q)-replenishment policy where the replenishment lead time is fixed. We have observed two methods to calculate the probability distribution of the demand during replenishment lead time; based on respectively the compound Poisson and an approximative probability distribution based on the mean and variance of the demand during lead time. The determination of the reorder point and replenishment order quantity is based on a sequential method where first the replenishment order quantity is determined, which is then used to calculate the optimal value of the reorder point. There exist simultaneous methods to determine optimal values for the reorder point and replenishment order quantity, but these methods are either too complex or not suitable for the situation at Vanderlande.

We have verified the model by comparing the order line fill rate resulting from each calculation method with the order line fill rate resulting from a simulation study. Both the compound Poisson calculation method (4.5) and the two moments method ((9.1), (9.3) or (9.5)) calculate order line fill rates that are close to the simulated order line fill rate; the compound Poisson methods scores better with respect to the absolute difference between the calculation and the simulation, but the two moments method shows acceptably small differences as well in some cases.

In the next chapter we expand our basic model by relaxing a number of assumptions we made in this chapter. We include variable replenishment lead times and a delivery timeframe, and include the possibility of fixed delivery days by suppliers as well. We continue verifying both the compound Poisson method and the two moments method to calculate the order line fill rate, but we reconsider this choice after each extension.

5. Order fill rate model

This chapter answers the fourth research question:

4. How can we extend the basic inventory policy and how is this advanced model optimized?

We extend the basic inventory policy from the previous chapter in four steps:

1) Include variable replenishment lead time
2) Include a timeframe
3) Include fixed supplier delivery days
4) Include order fill rate as performance measure

In the previous chapter we have concluded that the best method to model the customer demand is by using the compound Poisson method from equation (9.1) in calculating the customer demand during replenishment lead time. Since the results of the two-moments method of equations (9.1), (9.3) and (9.5) are in some cases acceptable as well, we include this method in the first extension.

5.1. Including variable replenishment lead time

In the basic model we assumed a fixed replenishment lead time per spare part, based on the lead time as mentioned in the data provided by Vanderlande. In reality suppliers do not always deliver the replenishment order at the promised date. We measure this unreliability in days too late per supplier. In practice it can occur that an order arrives before the promised delivery date, but in most cases this is only a few days too early. We consider these orders to have no delay. This means that in our model we:

1) slightly underestimate the order line fill rate; the real replenishment lead time is shorter than the actual replenishment lead time. Shorter replenishment lead times lead to higher order line fill rates.
2) slightly underestimate the total holding costs; if an order is e.g. two days early, the inventory level is at most two days higher than assumed.

The unreliability of the supplier is defined by a discrete probability distribution and has to be included in the lead time of the spare parts. This probability distribution is discrete since the replenishment lead time is measured in whole days. Suppliers deliver at most only once a day, so if an order is delayed, it is delayed by a complete day.

The lead time hereby becomes a discrete stochastic variable, denoted by L. In the next section we discuss a number of modeling options. In subsection 5.1.2 we discuss the verification of these modeling options and the results are given in subsection 5.1.3.

5.1.1. Modeling options

The general expression for the order line fill rate calculation, as given in (4.4), is:

$$OLFR(R,Q)=\sum_{j=1}^{R+Q}\Pr(F\leq j)\sum_{k=\max\{R+1,j\}}^{R+Q}\Pr(IP=k)\Pr(D(L)=k-j) \tag{5.1}$$

Including variable replenishment lead time means that the most right part of the expression, Pr(D(L)=k-j), changes. Using the compound Poisson method, we consider the following two modeling options to include the discrete probability distribution of the lead time in this expression.

1) Calculate the mean of the lead time and incorporate this mean as the fixed lead time in our basic model. This method is based on (Feeney, et al., 1966). The mean of the replenishment lead time is denoted with EL and it is calculated as:

$$EL = \sum_t t \Pr(L = t) \tag{5.2}$$

In the calculation of the probability distribution of the demand during replenishment lead time, we replace L by EL. The calculation of the order line fill rate in this extension hence becomes (compare equation (4.5)):

$$\begin{aligned} OLFR(R,Q) &= \sum_{j=1}^{R+Q} \Pr(F \le j) \sum_{k=\max\{R+1,j\}}^{R+Q} \Pr(IP = k) \Pr(D(EL) = k - j) \\ &= \frac{1}{Q} \sum_{j=1}^{R+Q} \Pr(F \le j) \sum_{k=\max(R+1,j)}^{R+Q} \left[\sum_{y=0}^{\infty} \frac{(\lambda EL)^y e^{-\lambda EL}}{y!} f^y (k-j) \right] \end{aligned} \tag{5.3}$$

2) Incorporate each possible lead time separately as fixed lead time in the basic model, and calculate the weighted average of the resulting order line fill rates. This method is based on conditioning on the possible replenishment lead times, as suggested by (Axsäter, 2006). Incorporating the replenishment lead time distribution in the calculation of the probability distribution of the demand during replenishment lead time leads to:

$$\begin{aligned} \Pr(D^*(L) = k - j) &= \sum_t \Pr(L = t) \Pr(D(L) = k - j \mid L = t) \\ &= \sum_t \Pr(L = t) \left[\sum_{y=0}^{\infty} \frac{(\lambda t)^y e^{-\lambda t}}{y!} f^y (k-j) \right] \end{aligned} \tag{5.4}$$

In the calculation of the order line fill rate (5.1) we have to replace Pr(D(L)=k-j) by Pr(D*(L)=k-j):

$$\begin{aligned} OLFR(R,Q) &= \sum_{j=1}^{R+Q} \Pr(F \le j) \sum_{k=\max\{R+1,j\}}^{R+Q} \Pr(IP = k) \Pr(D^*(L) = k - j) \\ &= \frac{1}{Q} \sum_{j=1}^{R+Q} \Pr(F \le j) \sum_{k=\max(R+1,j)}^{R+Q} \left(\sum_t \Pr(L = t) \left[\sum_{y=0}^{\infty} \frac{(\lambda t)^y e^{-\lambda t}}{y!} f^y (k-j) \right] \right) \end{aligned} \tag{5.5}$$

In the two moments method, we incorporate the variability of the lead time in the calculation of the mean and variance of the demand during replenishment lead time. For this calculation we need the mean and variance of the replenishment lead time, which can easily be determined from the (known) discrete probability distribution. The third modeling option we observe is hence:

3) Incorporate the mean and variance of the replenishment lead time in the two moments method in the calculation of the probability distribution of the demand during replenishment

lead time. In Appendix B.1.2. we discuss how the mean and variance of the replenishment lead time are included in the determination of the mean and variance of the demand during lead time, ED and varD respectively. The order line fill rate calculation is as in (5.1) where the method to calculate the probability distribution of the lead time demand is given by either (9.1), (9.3) or (9.5), depending on the variance-to-mean ratio of the lead time demand.

5.1.2. Verification

To verify our methods we consider an instance with seven suppliers, each with a different delay probability distribution – i.e. the discrete probability distribution that defines the delay on deliveries of this supplier. This delay distribution is combined with the basic lead time of each spare part to define the replenishment lead time probability distribution. The delay distributions of the suppliers are taken from real-life data on deliveries, where we have set the maximum delay to be seven weeks, or 35 workdays. The complete delay distributions are shown in Appendix C.2, Table 7 shows the key characteristics of these distributions.

Supplier	Mean	Standard deviation	Min	Max
A	0.35	2.33	0	35
B	0.06	0.60	0	13
C	4.89	4.91	0	35
D	1.40	5.74	0	35
E	4.48	7.50	0	35
F	0.11	0.56	0	5
G	0.21	1.26	0	12

Table 7 Supplier delay distributions for model verification (workdays)

In the basic model we executed 1,680 experiments for each method. Each experiment is applied to all seven suppliers and all three modeling options. This means we have (3*7*1,680=) 35,280 experiments. For each experiment, the calculated order lien fill rate is compared to the outcomes of a simulation study, similar to the simulation study of the verification in section 4.4.2.

Additional lead time adaptation in simulation

Due to the variable replenishment lead time, it can theoretically occur that from two outstanding replenishment orders the order that was ordered later arrivers earlier. The next example clarifies this:

Example 3

At an arbitrary time t (say t_A=0) a replenishment order (order A) is placed at the supplier. The standard lead time for this spare part is 10 days, but the delay on this specific order is 15 days. Hence, the total replenishment lead time for order A (L_A) is 25 days. At t_B=5, a second replenishment order (order B) is placed, on which there is no delay, so L_B is 10 days. Order B arrives at t_B+L_B=day 15, while order A arrives at t_A+L_A=day 25, so although order A is placed earlier, it arrives later.

In Example 3 the replenishment orders cross. In practice, this does not occur. This means that if an order is delayed, all subsequent orders are delayed until the moment the first delayed order arrives. In the example, order B is delayed such that is arrives together with order A on day 25.

In our simulation study, the replenishment lead time for a replenishment order is determined independent of the outstanding replenishment orders, based on the original delay distribution of the supplier. This means that the delay on the order is determined without taking into account that orders cannot cross. The actual delay of the order can be longer due to delays on preceding outstanding orders. This means that the actual (output) delay distribution is different from the original (input) delay distribution in our simulation study.

To verify our calculation, we need the delay distribution of the simulation to be equal to the distribution that is used in the calculations. We therefore execute the simulation study first and use the output delay distribution of the simulation study as delay distribution in our calculation. In practice we do not need to make this additional adaptation. In the data on which the delay distribution is based, this adaptation is already included.

5.1.3. Results

Table 8 shows the key results of the verification study for the three modeling options mentioned in section 5.1.1.

Measure	Option 1)	Option 2)	Option 3)
Absolute difference simulation-calculation	0.846%	0.056%	0.421%
Standard deviation on absolute difference	1.816%	0.065%	0.627%
90%-percentile absolute difference	2.705%	0.132%	1.082%
Maximum absolute difference	16.342%	0.669%	5.018%

Table 8 Key results for verification of modeling options to include variable lead times

There are large differences in the key results of the modeling options. Only the second modeling option scores acceptable, the other two options show much worse results. The results of the second modeling option – where conditioning on the replenishment lead time distribution is applied – show many similarities with the results of the verification study of the compound Poisson method in the basic model. The largest differences again occurred at customer order quantity distribution 7, where the difference between the smallest and largest possible order quantity is highest.

The first modeling option, where the mean of the replenishment lead time is used as fixed lead time, shows unacceptable differences, up to over 16%. The worst results occur for spare parts with low replenishment lead times, high arrival rates and a supplier that has a high mean and variability of the delay distribution. This is explained by the fact that for low replenishment lead times, high delays have more effect on the lead time than for high replenishment lead times. For high replenishment lead times the differences are unacceptable as well, especially for high variability delay distributions. For a spare part with a lead time of 65 days, in case the maximum delay is e.g. 35 days, but we incorporate the mean delay of 4.5 as additional lead time, we assume that the lead time is always 69.5 days, while it can be 100 days in practice.

The third modeling option is the adaptation of the two-moments method from the previous chapter by incorporating the mean and variance of the replenishment lead time as well. These results show some similarities with the results of the verification of the two-moments method in the basic model, in general the results in this section are slightly worse. This is caused by the fact that the delay distribution is an empirical distribution as well. We use only the mean and variance of this distribution to incorporate in the calculation method, hence the delay distribution is incorporated through an approximation. Most empirical distributions that we used have a relatively high probability of a maximum delay. Like we encountered in the basic model for customer order quantities with high probabilities of high order quantities, the approximation of the replenishment lead time distribution does not cope well with these high probabilities.

Looking at the calculation times, we observe large differences between the three methods. The first method took over 2 minutes to calculate the probability distribution of the demand during replenishment lead time for all 35,280 experiments. The third method (two-moments method) was nearly 30 seconds faster. The second method, which showed acceptable results with respect to the calculation of the order line fill rate, took nearly 17 minutes to calculate the order line fill rate. This is due to the fact that for each possible replenishment lead time the complete probability distribution has to be calculated.

Conclusion

Based on the verification study in this section we conclude that the only suitable modeling option is the second modeling option – i.e. the compound Poisson method, where the variability in replenishment lead time is included through conditioning on the lead time. We hence choose this method as the single method to which we apply the remaining extensions, starting by including the timeframe to the performance measure. This method takes considerably longer calculation time than the other methods. Since we have executed a very large amount of verification experiments, we expect that in practice the calculation time is not problematic.

5.2. Including delivery timeframe in performance measure

Until now we have considered the 'off-the-shelf' order line fill rate as performance measure. Vanderlande is interested in the percentage of complete orders that can be delivered within a certain timeframe. In this section we define how the timeframe is added to the model and how we calculate the time-based order line fill rate with respect to this timeframe. We build further on the previous extension, which means that we include variable replenishment lead times following equation (5.5).

5.2.1. Modeling approach

The goal is to fill a customer order within a timeframe of length τ. For our performance is does not matter in this case if the customer order is filled immediately upon arrival or just before the timeframe expires. With respect to the time-based order line fill rate, filling a customer order within a timeframe of length τ is hence equal to filling a customer order a period of length τ after customer arrival.

The basic approach is as follows: at the arrival of the customer, we immediately update the inventory level by adding the order to the backorders; thereby the inventory position is updated as well. If necessary – i.e. if the inventory position drops at or below the reorder point – a replenishment order is placed. We then 'hold' the customer order for a period of length τ, after which we fill the order in case the on hand inventory is high enough. Both the on hand inventory and the backorders are decreased with the customer order quantity at this moment.

In our model, we include the timeframe in a slightly different way. Instead of holding the customer, we reduce the replenishment lead times, and apply 'off-the-shelf' fill rate again. This means that our goal remains to fill a customer immediately upon arrival, but the lead time of all replenishment orders is reduced with τ. In case the replenishment lead time originally is less than the timeframe, the new lead time equals zero, so if the customer arrives and the available stock is not high enough, a replenishment order is placed and immediately becomes available. Hence the customer order is filled immediately.

Using this approach, we only need to alter the replenishment lead time distribution, compared to the previous extension. The new replenishment lead time probability distribution, indicated with the stochastic variable L^τ where τ denotes the length of the timeframe, is constructed as follows:

$$\begin{aligned}
&\Pr\left(L^\tau = t\right) = \Pr\left(L = t + \tau\right) \qquad t > 0 \\
&\Pr\left(L^\tau = 0\right) = \sum_{t \le \tau} \Pr\left(L = t\right)
\end{aligned} \tag{5.6}$$

A replenishment lead time of zero leads to an order line fill rate of 100%, since a customer order is either filled from stock upon arrival, or the replenishment order that is needed to fill the customer order arrives immediately and the customer order is still filled upon arrival. We incorporate this in the order line fill rate calculation from equation (5.5) as follows:

$$\begin{aligned}
OLFR(R,Q) &= \frac{1}{Q}\sum_{j=1}^{R+Q} \Pr\left(F \le j\right) \sum_{k=\max\{R+1,j\}}^{R+Q} \left(\sum_{t} \Pr\left(L^\tau = t\right) \left[\sum_{y=0}^{M} \frac{(\lambda t)^y e^{-\lambda t}}{y!} f^y (k-j) \right] \right) \\
&= \Pr\left(L^\tau = 0\right) + \\
&\qquad + \frac{1}{Q}\sum_{j=1}^{R+Q} \Pr\left(F \le j\right) \sum_{k=\max\{R+1,j\}}^{R+Q} \left(\sum_{t>0} \Pr\left(L^\tau = t\right) \left[\sum_{y=0}^{M} \frac{(\lambda t)^y e^{-\lambda t}}{y!} f^y (k-j) \right] \right)
\end{aligned} \tag{5.7}$$

Note that this only holds in case the reorder point is at least equal to -1, which is the case in our model as explained in section 4.2.1. For reorder points lower than -1, it is not ensured that a customer order that cannot be filled from the on hand inventory or the outstanding orders yields a replenishment order (with lead time zero). Hence, it is possible that, although the lead time equals zero, the customer order is not filled immediately, and thereby the order line fill rate in this case does not equal 100%.

5.2.2. Verification

We use the same set of experiments as in the previous verification in section 5.1.2, since we also include variable replenishment lead times in this model. Following the results of the verification study in section 5.1.3, we only verify one method; hence we have (1,680*7=) 11,760 experiments. In each experiment, the calculated time-based order line fill rate is compared to the outcomes of a simulation study containing five independent simulation runs of 20,000 replenishment orders each. The simulation runs are based on the basic approach mentioned in section 5.2.1, so customers are held for the length of the timeframe before filling the customer order. The timeframe we apply in our verification study is 15 workdays.

5.2.3. Results

The key results for this verification study are found in Table 9.

Measure	Value
Absolute difference simulation-calculation	0.040%
Standard deviation on absolute difference	0.057%
90%-percentile absolute difference	0.105%
Maximum absolute difference	0.836%

Table 9 Key results for including timeframe

All values in Table 9 lie within our range of acceptable solutions, since the absolute average is under 0.1% and the maximum is less than 1%. From the detailed results we conclude that the results sho many similarities with the results of the verifications of the compound Poisson method in the basic model and the previous extension.

Conclusion

Adding the timeframe to the model resulting from the previous section yields acceptable results. The average absolute difference between calculation and simulation is even less than in the previous verification study, so adding the timeframe has improved our calculation method slightly. This is explained by the fact that – in case the replenishment lead time is lower than the timeframe – the order line fill rate is 100% in both the calculation and simulation; hence, there is no difference at all.

5.3. Adding fixed supplier delivery days

At Vanderlande there are a number of suppliers that only deliver replenishments on fixed days of the week, e.g. only on Tuesdays and Thursdays. All other suppliers can deliver replenishments each day, but only once a day.

Until now we have assumed a continuous review policy. This means that replenishments can arrive at all moments, and that it is possible that there are more deliveries on one day. In practice there is at most one delivery per day per supplier, and for a number of suppliers only one or two per week. This is best modeled by a periodic review model, as we clarify in the following section.

5.3.1. Periodic review model

We first assume that a supplier can deliver at all days, but only once per day. Without loss of generality we assume that the supplier delivers at the end of the day. As mentioned in section 5.1.1 the replenishment lead time is measured in whole days, so if we place a replenishment order at the end of each day – in case the inventory position is at or below the reorder point – the order arrives at the end of the day as well.

Placing replenishment orders only at the end of each day (if necessary) is modeled by a periodic review model with a review period of one day. At the end of each review period, so at the end of each day, the inventory position is observed and a replenishment order is placed.

Now observe a supplier that only delivers once per week, i.e. once per five workdays. The next example shows that we can model such a supplier with a periodic review model with a review period length of five days:

Example 4

The supplier of a spare part delivers only on Thursdays. The replenishment lead time of the spare part is equal to eight days; for now we assume that this supplier always delivers on time. This means that a replenishment order is delivered at the first Thursday after the lead time of eight days is expired. An order that is placed on Tuesday in week X is delivered on Thursday in week X+2, since eight days later is Friday week X+1, so the first possible delivery day is Thursday X+2. An order that is placed on Monday in week X+1 is delivered on Thursday in week X+2 as well, since eight days later is exactly this Thursday.

From the above it becomes clear that ordering a replenishment order on Tuesday in week X equals ordering a replenishment order on Monday in week X+1. In this case we can hold the ordering of replenishment orders until Monday due to the lead time of eight days. This is modeled by a periodic review model where the 'review moment' (i.e. the end of the review period) is each Monday, hence the review period length is five days.

So in case the supplier delivers once per five days, the review period length is equal to five days. Similarly, in case a supplier would deliver once per two days, the review period length equals two days. In case a supplier delivers twice a week, we have two options. We could 1) define the review period length as 2.5 days, or 2) calculate the performance for both a review period length of two and three days, and take the average of the outcomes of those two models.

From this point on we denote the review period length with T. As stated above, in a periodic review model the inventory position is not continuously observed, but only at the end of fixed time intervals of length T. For the order line fill rate calculation, we need to observe the inventory level. We observe the inventory level just after the arrival of a replenishment order. From this point, the inventory level is not increased by any replenishment order for a period of length T, since only once per T days a replenishment order can arrive.

Consider that at time t a replenishment order has just been placed, so we are at the start of a new review period. The replenishment order arrives at time t+L, where L is the replenishment lead time. As in the basic model, the inventory level at time t+L equals the inventory position at time t (just after the replenishment order is placed) minus the demand during the period (t,t+L). So:

$$\begin{aligned}\Pr\left(IL(t+L)=j\right) &= \sum_{k}\Pr\left(IP(t)=k\right)\Pr\left(D(t,t+L)=k-j\right) \\ &= \sum_{k=\max\{R+1,j\}}^{R+Q}\Pr\left(IP(t)=k\right)\Pr\left(D(t,t+L)=k-j\right)\end{aligned} \tag{5.8}$$

The inventory level at the time t+L is used to fill all orders during the period (t+L,t+L+T), the moment at which the next replenishment order arrives. To find the order line fill rate, we need to calculate the probability that y customers arrive in this period times the probability that only x orders are filled in case the starting inventory level is j, for all possible values of j (0 to R+Q), x (1 to y) and y (1 to infinity). The calculation of this probability is clarified in the next example:

Example 5

The inventory level at time t+L is equal to 5. In case one customer arrives during the review period length, the order line fill rate is 100% if the customer order quantity is at most 5, and 0% in case the customer orders more than 5. If two customers arrive, the order line fill rate is 100% if the cumulative demand is at most 5, it is 50% if the order quantity of the first customer is at most 5 and the order quantity of the second customer is more than the remaining inventory level. The order line fill rate is 0% if the first customer orders more than 5.

In general, the order line fill rate in case y customers arrive during the period t+L,t+L+T equals $x/y*100\%$ if the cumulative demand of the first x customers is at most the inventory level, and the customer order quantity of the next customer is more than the remaining inventory level after the first x customer orders have been filled.

The calculation of the order line fill rate following Example 5 is found in Appendix B.4, the resulting expression is:

$$\begin{aligned} OLFR(R,Q) = \sum_{z=0}^{R+Q} \Pr\left(IL(t+L)=z\right) \Bigg\{ \sum_{y=1}^{\infty} \Pr\left(C(T)=y\right) \cdot \ldots \\ \ldots \cdot \left[\sum_{x=0}^{y-1} \frac{x}{y} \sum_{j=0}^{z} f^{x}(j) \Pr\left(F > z-j\right) + \sum_{i=0}^{z} f^{y}(i) \right] + \Pr\left(C(T)=0\right) \Bigg\} \end{aligned} \quad (5.9)$$

Here, C(T) is a stochastic variable denoting the number of customers that arrive during the review period of length T.

The problem with the calculation of the (time-based) order line fill rate in a periodic review method is that we do not know the probability distribution of the inventory position in equation (5.8). In the previous models, the inventory position was independent of the time t and uniformly distribution between R+1 and R+Q. in this case, the inventory position is not independent of the time t anymore, and the probability distribution of the inventory position at the start of a review period is very hard to determine. We have therefore developed an approximation of the periodic review model which we describe in the remainder of this section. Nevertheless we have carried out a small verification of the order line fill rate calculation of equation (5.9). This verification is worked out in Appendix D.4.2. From this verification results that the equation is indeed not suitable to calculate the order line fill rate in a periodic review model.

Consider again Example 4. The actual lead time of the order that is originally placed on Tuesday is not eight days but twelve days. Similarly the replenishment lead time for an order placed on Thursday is ten days. The probability that an order is placed on Tuesday is equal to 20%; due to the exponential arrival process of customers, the probability that a customer arrives on Tuesday is equal to the probability that it arrives on other days and hence is 20% for all days of the week.

We combine the probability that an order is placed at a certain day with the replenishment lead time that follows if an order is placed at this day, and use this to create an additional ‘delay’ in the replenishment lead time. In Example 4, assuming a fixed lead time of eight days, we apply the following delay probability distribution to the replenishment lead time:

$$\Pr(\text{delay} = 0 \text{ days}) = 0.2$$
$$\Pr(\text{delay} = 1 \text{ days}) = 0.2$$
$$\Pr(\text{delay} = 2 \text{ days}) = 0.2$$
$$\Pr(\text{delay} = 3 \text{ days}) = 0.2$$
$$\Pr(\text{delay} = 4 \text{ days}) = 0.2$$

This delay distribution is applied to the original replenishment lead time distribution. If the lead time is fixed, say L*, the resulting replenishment lead time distribution is given by:

$$\Pr(L = L^* + t) = \Pr(\text{delay} = t)$$

In case the original replenishment lead time is variable, the delay distribution is applied to all possible replenishment lead times, as clarified in the next example.

Example 6

The replenishment lead time is 10 days with 90% probability, and 13 days with 10% probability. The supplier of this spare part delivers once per week, so we need to apply the additional delay distribution as sketched above.

The probability that the adapted lead time equals 10 days is 90%*20% (the probability of no delay). The same holds for lead times of 11 and 12 days. The probability of a replenishment lead time of 13 days equals 90%*20%+10%*20%; it consists of the original lead time of 10 days with a delay of 3 days, and the lead time of 13 days with zero days delay.

The resulting replenishment lead time distribution becomes:

$$\begin{aligned}
&\Pr(L=10)=0.18; \quad &&\Pr(L=11)=0.18; \quad &&\Pr(L=12)=0.18 \\
&\Pr(L=13)=0.20; \quad &&\Pr(L=14)=0.20; \quad &&\Pr(L=15)=0.02 \\
&\Pr(L=16)=0.02; \quad &&\Pr(L=17)=0.02
\end{aligned} \tag{5.10}$$

This method is an approximative method since we assume a variable replenishment lead time where a fixed lead time is the case. In Example 4, if an order is placed on Thursday, the lead time is 10 days with 100% probability, instead of either 8, 9, 10, 11, or 12 days, all with 20% probability.

Using this approximation we have ruled out the fixed delivery days of suppliers, but we still remain with the fact that suppliers only deliver once per day, modeled by a periodic review model with a review period of one day. Since the order line fill rate cannot be calculated in the periodic review model, we need to approximate this model by a continuous review model.

Consider the periodic review model with a review period of one day. Replenishment orders are placed at the end of each day (if necessary), but customers arrive during the day. The distribution of arriving customers over a day is uniform, due to the exponential arrival process. Hence the average arrival time of a customer is exactly halfway through the day. For this 'average' customer the replenishment lead time is half a day higher than it is at the moment the order is actually placed. If we assume the replenishment order to be made immediately after the arrival of a customer (as in a continuous review model) the replenishment lead times are on average 0.5 days higher. We add this

0.5 day to all possible replenishment lead times resulting from the adaptation that is described above. The replenishment lead time probability distribution of (5.10) hence becomes:

$$\begin{aligned}
&\Pr(L=10.5)=0.18; \quad \Pr(L=11.5)=0.18; \quad \Pr(L=12.5)=0.18 \\
&\Pr(L=13.5)=0.20; \quad \Pr(L=14.5)=0.20; \quad \Pr(L=15.5)=0.02 \\
&\Pr(L=16.5)=0.02; \quad \Pr(L=17.5)=0.02
\end{aligned}$$

5.3.2. Verification

To verify this model, we have selected two of the delay probability distributions of our variable replenishment lead time verification studies as mentioned in section 5.1.2; those of suppliers B and C. The parameters for these suppliers are shown in Table 10.

Supplier	Mean	Standard deviation	Min	Max
B	0.056	0.60	0	13
C	4.885	4.91	0	35

Table 10 Parameters of suppliers B and C in verification study

Each supplier is subjected to three situations:

1) delivering once per week
2) delivering twice per week on non-consecutive days
3) delivering once per day

This creates six scenarios. Each scenario is subjected to all 1,680 experiments, creating a total 10,080 experiments in total. For each experiment, we calculate the order line fill rate and compare this to the result of a simulation study. In this study we simulate the periodic review method we want to approximate. For the second situation, delivering twice per week on non-consecutive days, we have set the review period length to 2.5 days. In all experiments, variable lead times are included. We have verified both the model without timeframe and including a timeframe of 15 days.

5.3.3. Results

The following table summarizes the results of the verification of the approximative continuous review model for including the fixed delivery days of suppliers, without a timeframe and including a timeframe of 15 days.

Measure	Without Timeframe	Including Timeframe
Absolute difference simulation-calculation	0.059%	0.047%
Standard deviation on absolute difference	0.065%	0.059%
90%-percentile absolute difference	0.139%	0.119%
Maximum absolute difference	0.614%	0.669%

Table 11 Key results for including fixed delivery days

All values of Table 11 lie within our range of acceptable solutions, since the absolute average is under 0.1% and the maximum is less than 1%. The average absolute difference is slightly smaller for the experiments including a timeframe of 15 days. This is similar to the results of the verification studies in sections 5.1.3 and 5.2.3; the experiments including timeframe showed slightly better results.

Conclusion

Adding the fixed delivery days of suppliers through an adaptation of the replenishment lead time in a continuous review method yields acceptable results. Compared to the model without this adaptation, the average difference between simulation and calculation has only increased a bit.

5.4. Calculation of order fill rate

The previous sections have resulted in a model to calculate the time-based order line fill rate including variable supplier lead time and a desired delivery timeframe. The final step is to use these time-based order line fill rates to determine the time-based order fill rate, i.e. the percentage of orders that is completely filled within the timeframe.

5.4.1. Calculating order fill rate from order line fill rates

Before we model the calculation of the order fill rate, we investigate the composition of orders. We do this by looking at the correlation between spare parts. If two spare parts have a high positive correlation (i.e. close to 1), this means that if one spare part is ordered, the probability that the other spare part is ordered as well is high, and vice versa. When the correlation is close to -1, this means that if one spare part is ordered, the probability that the other spare part is NOT ordered is high, and vice versa. The smaller the absolute value of the correlation between two spare parts, the more independent these spare parts are of each other. We classify the correlations between spare parts following Table 12:

Abs.value of correlation	Class
>0.90	Very high
0.70-0.90	High
0.40-0.70	Medium
0.20-0.40	Low
<0.20	Very low

Table 12 Classification of correlations

We have observed all spare parts that have been ordered at least ten times between July 2003 and April 2009 and all orders of at least two order lines from that period. There are 1,277 spare parts that satisfy these criteria. For each spare part we have calculated the correlation with all other spare parts and made an overview of the largest correlation with any other spare part, the second largest correlation with any other spare part, etcetera. The first row in Table 13 shows how many spare parts have a largest correlation with any other spare part this is at least high (first column) and at least medium (second column). The second row shows for how many spare parts the second largest correlation is at least high and medium, in other words how many spare parts have at least two other spare parts with which they are highly of medium correlated.

X^{th} largest correlation	>0.70	>0.40
1^{st}	119	385
2^{nd}	33	154
3^{rd}	15	77
4^{th}	7	35
5^{th}	6	29
6^{th}	4	25

7th	0	22
---------------	----------	--------------
15th	0	7
16th	0	0

Table 13 Number of high or medium correlated spare parts

The fifteen spare parts that have at least three other spare parts with which they are highly correlated (the highlighted entry in the first column, third row of Table 13) can be divided into three groups of spare parts that are mutually highly correlated. These groups consist of e.g. left and right cover plates, it is obvious that those spare parts are mostly ordered together.

It appears that there are no spare parts that are at least medium correlated to more than fifteen other spare parts. Since the number of spare parts that is highly correlated to more than one other spare part is considerably low as well, we assume that spare part arrivals are not correlated and hence customer orders are composed arbitrarily. For (groups) of spare parts that are highly correlated, there are two options:

1) Create a new 'spare part' consisting of all spare parts that are highly correlated to each other. If this spare part is put on stock at a certain amount, each underlying spare part is put on stock in that amount.
2) After the stock heights have been determined, ensure that for each group of highly correlated spare parts for all underlying parts the same decision is made; hence, if one of these spare /parts is put on stock, put all spare parts on stock.

In both options we may assume that spare part arrivals are not correlated. Assuming this has some interesting consequences. The order fill rate for an order of multiple order lines can now be calculated easily from the order line fill rates. First of all, the order fill rate for orders consisting of one order line, denoted by OFR_1, is equal to the weighted average of the order line fill rates of all items, so: $OFR_1 = \sum_{i \in I} p_i OLFR_i$ where p_i is the probability that item i is ordered: $p_i = \lambda_i \Big/ \sum_{j \in I} \lambda_j$ and I the set of all items. Now consider the following example:

Example 7

An order consists of two order lines, items A and B with respective calculated time-based order line fill rates $OLFR_A$ and $OLFR_B$. The order fill rate of this order is at least equal to $OLFR_A$*$OLFR_B$ (Song, 1998). The probability that this order actually occurs is, because orders are composed arbitrarily, equal to p_A*p_B, the product of the respective arrival probabilities. The contribution of the order (A,B) to the total order fill rate is hence equal to p_A*p_B*$OLFR_A$*$OLFR_B$.

Summing over all possible order compositions – which is equal to all combinations of two items – gives an expression for the general order fill rate of orders of two order lines. Because the product of the individual order line fill rates is a lower bound to the order fill rate (Song, 1998), we get a lower bound for the general order fill rate.

$$OFR_2 \geq \sum_{i \in I} \sum_{j \in I, j \neq i} \left(OLFR_i p_i OLFR_j \frac{p_j}{1-p_i} \right) \geq \sum_{i \in I} \sum_{j \in I, j \neq i} \left(p_i p_j OLFR_i OLFR_j \right) \tag{5.11}$$

Because we have a lot of items, determining the right side of the inequality takes a lot of calculations. We therefore simplify these calculations as follows:

$$\begin{aligned} OFR_2 &\geq \sum_{i\in I}\sum_{j\in I, j\neq i}\left(p_i p_j OLFR_i OLFR_j\right) \\ &\geq \left(\sum_{i\in I} p_i OLFR_i\right)\left(\sum_{j\in I} p_j OLFR_j\right) - \sum_{i\in I} p_i p_i OFLR_i OLFR_i \\ &\geq \left(OFR_1\right)^2 - \sum_{i\in I} p_i^2 OLFR_i^2 \geq \left(OFR_1\right)^2 \end{aligned} \tag{5.12}$$

The difference $\sum_{i\in I} p_i^2 OLFR_i^2$ is extremely small since $p_i \ll 1$ for all spare parts. In the same way we find that $OFR_3 \geq \left(OFR_1\right)^3$ etc., so in general:

$$OFR_i \geq \left(OFR_1\right)^i \tag{5.13}$$

Using the lower bound as deduced above in, we calculate our total order fill rate TOFR as follows:

$$TOFR \geq \sum_{i=1}^{\infty} q_i \left(OFR_1\right)^i = \sum_{i=1}^{\infty} q_i \left(\sum_{j\in I} p_j OLFR_j\right)^i \tag{5.14}$$

Here q_i is the probability that an arbitrary order consists of i order lines. The values of q_i are based on an empirical distribution deduced from the historical data on service orders. $(OFR_1)^i$ gives the total order fill rate for orders that consist of i order lines.

We define a maximum number of order lines – denoted as M – based on the historical data. The 2% largest orders are discarded from this data set to exclude very large orders that are probably not service orders at all. Given this number M and a desired total order fill rate we find our target order line fill rate (i.e. the order fill rate for orders consisting of one line) as the solution X of the following equation.

$$\sum_{i=1}^{M} q_i X^i \geq TOFR^{\text{target}} \tag{5.15}$$

We call X the target order line fill rate. Since the total order fill rate increases when X increases, it suffices to solve the inequality as an equality. Next we determine the order line fill rate per item such that $\sum_{i\in I} p_i OLFR_i \geq X$. This solution will never overestimate the total order fill rate since:

$$TOFR \geq \sum_{i=1}^{M} q_i \left(\sum_{j\in I} p_j OLFR_j\right)^i \geq \sum_{i=1}^{M} q_i X^i \geq TOFR^{\text{target}} \tag{5.16}$$

5.4.2. Verification

As for the order line fill rate models, we have verified the order fill rate model of section 5.4.1 by comparing the calculation with the outcomes of a simulation study for a set of fictional spare parts.

We used the same set as in section 5.3.2. Since taking all spare parts of that set into account would result in extremely high order arrival rates, we chose a number of spare parts randomly from this set. We have executed the verification for sets of 100 spare parts, 150 spare parts and 300 spare parts.

This verification set only contains data on spare parts, not on orders. We have designed five (fictional) order size probability distribution, which determines the number of order lines in an arriving order:

1) All orders consist of one order line
2) All orders consist of two order lines
3) Orders consist of one to five order lines, each with equal probability (i.e. 20%)
4) Orders consist of one to twelve order lines, using the following probabilities. This distribution is based on the actual order data from April 2008 to August 2009.

#order lines	1	2	3	4	5	6
Probability	59.2%	18.0%	8.3%	4.6%	3.3%	1.8%
#order lines	7	8	9	10	11	12
Probability	1.3%	1.0%	0.8%	0.7%	0.5%	0.5%

This leads to (3*4=) 12 combinations of spare part set and order size probability distribution. Each combination is subjected to target order fill rates of 60% and 80%, hence creating 24 experiments in total.

As in all previous verification studies, we have compared the calculation method to the average of five simulation runs. In each simulation run, we generated orders following the calculated order arrival rate. For an arriving order, the number of order lines was determined following the order size distribution of choice. To each order line a spare part was assigned through the demand probabilities of the spare parts; we ensured that the different order lines within one order contained different spare parts. In the simulation, an order is only filled if all spare parts are available on time.

5.4.3. Results

In Appendix D.5 the complete results are shown. Table 14 shows some key results of the verification study.

Measure	Value
Absolute difference simulation-calculation	0.085%
Standard deviation on absolute difference	0.078%
Maximum absolute difference	0.285%

Table 14 Key results of order fill rate model verification

The key results all lie within our threshold for acceptable solutions (<0.1% average, <1% max as stated in section 4.4.3).

From the detailed results we conclude that in all cases over 50% of all spare parts is put on stock, and in case the target order fill rate equals 80%, nearly 90% of all spare parts are put on stock in some cases. In all experiments, the average order line fill rate of the spare parts that are put on stock is near 90%; this means that if a spare part is put on stock, it is put on stock such that most customer orders are filled. We hence encounter few spare parts with an order line fill rate of e.g. 50%.

5.4.4. Validation

From the previous section becomes clear that the order fill rate calculation of section 5.4.1 suits the set of fictional spare parts. However, these spare parts are all completely independent so we would have expected the results of the verification to be positive. In this section we validate the order fill rate model using real life data of Vanderlande as input.

In section 2.4 we have discussed the performance regarding service orders between April 2008 and August 2009. The order fill rate with respect to a timeframe of three weeks was 60.2% in this case, with total yearly holding costs of around €37,000. Using the stock levels and minimum order quantities from the database, the model leads to an order fill rate of 52.8% and total yearly costs of €44,136. These costs are based on the physical holding costs only; the indirect holding costs are not taken into account here. In this study, we have assumed an order fill rate of 0% for all spare parts that do not have a positive stock level in the current situation, unless of course the replenishment lead time is shorter than the timeframe and (a portion of) the replenishment orders are delivered within the timeframe.

However, for a number of spare parts, minimum order quantities at the supplier exist. In section 4.2.1 we have seen that the reorder point should be always at least equal to -1, also if the order quantity is larger than 1 (which is the case for these spare parts). This leads to a situation for spare parts with larger replenishment order quantities, that a stock is held and hence the order line fill rate is larger than 0% for spare parts that are initially not put on stock. We have set the reorder points at least equal to -1 for all spare parts and executed the validation study again. This resulted in an order fill rate of 54.5% and total costs of €54,910 per year.

This validation study is not exact either. The minimum order quantity for spare parts is in some cases based on the minimum order quantity that was agreed with the project organization of Vanderlande. In general this department orders parts in very large amounts, so the quantities are not always suited for the service department. In some cases, the service department therefore does not order these parts in very large amounts, but uses some of the stock of the project organization, or they order a smaller amount at the supplier. As shown in section 4.2.1 this is best modeled by a replenishment order quantity of 1. In the final validation study hence we have set all replenishment order quantities to 1. Consequently, all reorder points are at least equal to -1. This results in an order fill rate of 50.1% and total costs of €27,508 per year.

Table 15 summarizes the results of the different validation studies.

Study	Order fill rate	Total costs
Actual performance	60.2%	€37,000
Basic	52.8%	€44,136
All R≥-1	54.5%	€54,910
All Q=1, R≥-1	50.1%	€27,508

Table 15 Results of validation studies

In practice the handling of the large minimum order quantities lies somewhere between the second and third validation study, so in some cases the large order quantity is ordered at the supplier and the remainder is put on stock, but in most cases a smaller amount is ordered at the supplier or the needed spare parts are taken from the warehouse of the project organization of Vanderlande. With respect to the costs, the validation study comes relatively close to the actual holding costs of €37,000, especially considering the fact the these holding costs are in practice calculated differently from our method.

The order fill rate lies in all validation studies well below the actual order fill rate of 60.2%. There are a number of explanations for this difference:

- In April 2008, the stock levels were slightly higher for some spare parts compared to the current stock levels. We have executed the validation study again with these 'old' stock levels, but the increase in order fill rate was at most 0.1% in all cases.
- In some cases, it is possible for the supplier to deliver spare parts in fewer days than the replenishment lead time that is assumed in the database. This occurs when there is an emergency order or when the supplier has the spare part available in stock.
- The service department largely profits from the stock levels in the warehouse of the project organization of Vanderlande. Numerous parts that are ordered by customers of the service department are held on stock in this warehouse as well, and if these spare parts are freely available, the service department is allowed to use this stock. Since this warehouse is at the terrain of Vanderlande itself, the replenishment lead time of these spare parts, originally assumed to be e.g. 20 days, becomes only a few days, namely the time it costs to ship it from the Vanderlande warehouse to the service central warehouse.

The third situation occurs very often, according to the employees of the Team Worldwide, the team that is responsible for the supply of spare parts. Their estimate is that it is certainly possible that for 10% of all service orders, the total lead time of the order was reduced to less than 15 days due to the fact that one or more of the ordered spare parts were available at the warehouse of the project organization. Especially the past few months this warehouse is trying to reduce its stock levels, causing more spare parts to become available for the service department.

Although the validation outcomes showed large differences with the performance in practice, with respect to the order fill rate, employees of the service department believe that this difference can be explained by the possibility to strongly reduce the replenishment lead times of spare parts by taking these spare parts from the warehouse of the project organization of Vanderlande. Therefore we conclude that the model not only suits the fictional set of spare parts, but is applicable to the real life data of Vanderlande as well.

The model results in a lower bound of the actual order fill rate, since in practice shorter replenishment lead times than assumed are possible. However, we may not assume that these lead times can always be met. Therefore we do not apply any correction to the order fill rate.

5.4.5. Determine reorder points

We now have a model to calculate the order fill rate, based on the order line fill rate per item. In this section we show how we find the reorder points for all spare parts such that the target order fill rate is achieved against minimal holding costs. For each spare part i, the following parameters are given or calculated:

- F_i, the customer order quantity distribution of the spare part.

- L_i, the replenishment lead time distribution, calculated by the basic replenishment lead time added to the delay probability distribution of the supplier and if necessary additional adaptation for handling the fixed delivery days of the supplier.
- λ_i, the demand rate of the spare part.
- p_i, the demand probability of the spare part, calculated by dividing the demand rate by the total demand rate (i.e. the sum of all demand rates)
- h, the daily holding costs per spare part.
- Q_i, the replenishment order quantity of the spare part, determined by (4.9)

The decision variable for each spare part is the reorder point R_i. As stated in section 5.4.1, we need to determine these reorder points such that the joint target order *line* fill rate β is achieved:

$$\sum_{i \in I} p_i OLFR_i \left(R_i\right) \geq \beta \tag{5.17}$$

Besides, we want to minimize the total holding costs. These costs are determined through the inventory level. At an inventory level of j, the holding costs are jh, with h the daily holding costs per spare part. The total holding costs are given by summing over all possible values of the inventory level:

$$C\left(R\right) = \sum_{j=1}^{R+Q} jh \Pr\left(IL^R = j\right) \tag{5.18}$$

The superscript R at the inventory level is to emphasize that the inventory level depends on the reorder point R, since:

$$\begin{aligned} \Pr\left(IL^R = j\right) &= \sum_{k=\max\{R+1,j\}}^{R+Q} \Pr\left(IP = k\right) \Pr\left(D\left(L\right) = k - j\right) \\ &= \frac{1}{Q} \sum_{k=\max\{R+1,j\}}^{R+Q} \Pr\left(D\left(L\right) = k - j\right) \end{aligned} \tag{5.19}$$

The optimization problem we want to solve is hence formulated as follows:

$$\begin{aligned} &\min \sum_{i \in I} C_i\left(R_i\right) \\ &\text{s.t. } \sum_{i \in I} p_i OLFR_i\left(R_i\right) \geq \beta \\ &\qquad R_i \geq -1 \qquad \forall i \in I \end{aligned} \tag{5.20}$$

Filling in the expression for the holding costs and order line fill rate, we get:

$$
\begin{aligned}
&\min \sum_{i\in I}\frac{h_i}{Q_i}\sum_{j=1}^{R_i+Q_i} j \sum_{k=\max\{R_i+1,j\}}^{R_i+Q_i} \sum_t \Pr\left(L_i = t\right)\sum_{y=0}^{\infty}\frac{\left(\lambda_i t\right)^y e^{-\lambda_i t}}{y!} f_i^{\,y}\left(k-j\right)\\
&\text{s.t. } \sum_{i\in I}\frac{p_i}{Q_i}\sum_{j=1}^{R_i+Q_i}\Pr\left(F_i \le j\right)\sum_{k=\max\{R_i+1,j\}}^{R_i+Q_i}\sum_t \Pr\left(L_i = t\right)\sum_{y=0}^{\infty}\frac{\left(\lambda_i t\right)^y e^{-\lambda_i t}}{y!} f_i^{\,y}\left(k-j\right)\ge \beta \qquad (5.21)\\
&\quad R_i \ge -1 \qquad \forall i\in I
\end{aligned}
$$

As becomes clear from the formulation of (5.21) this optimization problem is very complex. We therefore need to find an alternative to find the reorder points for all spare parts that lead to the desired order fill rate against minimal holding costs.

We choose to find the reorder points for all spare parts using a greedy heuristic, also referred to as marginal analysis. This method is – amongst others – described in (Sherbrooke, 2004). A greedy heuristic is a method where per spare part a function is defined that determines how good increasing the reorder point of this spare part is with respect to the total order line fill rate and holding costs, taking the current reorder point as starting point. The spare part with the highest function value is chosen to achieve 'the biggest bang for the buck' and for this spare part the reorder point is increased.

The function that determines how good improving the reorder point is, is called the delta function and is denoted with Δ. The delta function value for spare part j when increasing the reorder point to R from R-1 is denoted with $\Delta_j(R)$.

(Sherbrooke, 2004) proves that in case the delta functions for all spare parts are non-increasing, the greedy heuristic provides an optimal solution. We hence need to define a delta function that is non-increasing to find an optimal solution.

In our case, the delta function for each spare part has three elements:

- The increase in order line fill rate in case the reorder point is increased by 1. This is given by (based on the calculations in Appendix B.2:

$$
\begin{aligned}
OLFR\left(R\right)-OLFR\left(R-1\right)=\frac{1}{Q}\Bigg[&\sum_{j=1}^{R+Q}\Pr\left(F\le j\right)\Pr\left(D=R+Q-j\right)\\
&-\sum_{j=1}^{R}\Pr\left(F\le j\right)\Pr\left(D=R-j\right)\Bigg]
\end{aligned}
$$

- The increase in holding costs in case the reorder point is increase by 1, given by (based on combining (5.18) and (5.19))

$$
C\left(R\right)-C\left(R-1\right)=\sum_{i=1}^{R} h\Pr\left(D\left(L\right)=R-i\right)+\sum_{j=1}^{R+Q}\frac{jh}{Q}\Pr\left(D\left(L\right)=R+Q-j\right)
$$

- The probability that the spare part is demanded, given by:

$$
p_i=\frac{\lambda_i}{\sum_{j\in I}\lambda_j}
$$

We have combined these elements to define the delta function as follows:

$$\Delta_j(R) = p_j \frac{OLFR(R) - OLFR(R-1)}{C(R) - C(R-1)} \tag{5.22}$$

This delta function calculates the increase in order line fill rate divided by the increase in holding costs, weighed with the demand probability of the spare part. This means that it calculates the increase in total order line fill rate (=p_j*(OLFR(R)-OLFR(R-1))) divided by the additional holding costs that need to be made (C(R)-C(R-1)). Thereby it calculates the 'bang for the buck' and the next step in the heuristic is to choose the 'biggest bang', i.e. the spare part with the largest delta value.

In (Sherbrooke, 2004) the holding costs are linear; increasing the reorder point with 1 means additional holding costs of h, with h the (daily) holding costs per spare part. This means that only the order line fill rate function determines if the delta function is non-increasing. For this, the difference OLFR(R)-OLFR(R-1) needs to be non-increasing in R. If the order line fill rate function is concave, this is the case. Hence, in case the holding costs are constant, a concave order line fill rate function leads to an optimal solution.

Our delta function from equation (5.22) is non-increasing if the following holds:

$$\begin{aligned}
&\Delta_j(R) \le \Delta_j(R-1) \\
&\Leftrightarrow p_j \frac{OLFR(R) - OLFR(R-1)}{C(R) - C(R-1)} \le p_j \frac{OLFR(R-1) - OLFR(R-2)}{C(R-1) - C(R-2)} \\
&\Leftrightarrow \frac{OLFR(R) - OLFR(R-1)}{C(R) - C(R-1)} \le \frac{OLFR(R-1) - OLFR(R-2)}{C(R-1) - C(R-2)}
\end{aligned}$$

$$\Leftrightarrow \frac{OLFR(R) - OLFR(R-1)}{OLFR(R-1) - OLFR(R-2)} \le \frac{C(R) - C(R-1)}{C(R-1) - C(R-2)} \tag{5.23}$$

We first investigate our cost function. The cost function is based on the inventory level of the spare part; at a certain reorder point R the holding costs are given by equation (5.18). For values of k larger than one holds that $\Pr\left(IL^R = k\right) = \Pr\left(IL^{R-1} = k-1\right)$. The next equation shows that if the probability that the inventory level is zero is zero, the increase in holding costs is h when the reorder point is increased by 1.

$$\begin{aligned}
C(R) - C(R-1) &= \sum_{k=1}^{R+Q} kh \Pr\left(IL^R = k\right) - \sum_{j=1}^{R+Q-1} jh \Pr\left(IL^{R-1} = j\right) \\
&= h\left[\sum_{k=1}^{R+Q} k \Pr\left(IL^R = k\right) - \sum_{j=1}^{R+Q-1} j \Pr\left(IL^{R-1} = j\right)\right] \\
&= h\left[\sum_{k=1}^{R+Q} k \Pr\left(IL^{R-1} = k-1\right) - \sum_{j=1}^{R+Q-1} j \Pr\left(IL^{R-1} = j\right)\right] \\
&= h\left[\sum_{k=1}^{R+Q} \Pr\left(IL^{R-1} = k-1\right)\right] = h
\end{aligned}$$

If the probability that the inventory level equals zero is positive, the above does not hold anymore, if $\Pr(IL^{R-1}=0)$ does not equal $\Pr(IL^{R}=1)$. In this case, the holding costs increase by:

$$C(R)-C(R-1)=h\left[\sum_{k=2}^{R+Q}\Pr\left(IL^{R-1}=k-1\right)+\Pr\left(IL^{R}=1\right)\right]$$

If $\Pr(IL^{R-1}=0)$ equals $\Pr(IL^{R}=1)$, the part between square brackets equals 1 again, and the holding costs increase by h. Else, $\Pr(IL^{R}=1)<\Pr(IL^{R-1}=0)$ and the holding costs increase by less than h. Hence, the holding costs are linear in case the reorder point is large enough. In other cases, we only know that the increase in holding costs is less than h if the reorder point is increased by 1.

We return to the investigation of our delta function. Regarding the holding costs, there are two options:

1) Cost function is linear: in this case, the delta function is non-increasing if:
$$\frac{OLFR(R)-OLFR(R-1)}{OLFR(R-1)-OLFR(R-2)}\leq 1$$
This is the case if:
$$OLFR(R)-OLFR(R-1)\leq OLFR(R-1)-OLFR(R-2)$$
Or equivalently:
$$OLFR(R)+OLFR(R-2)-2OLFR(R-1)\leq 0$$
The last inequality holds if the order line fill rate function is (linear or) concave.
2) Cost function is not linear: in this case, the delta function is non-increasing if:
$$\frac{OLFR(R)-OLFR(R-1)}{OLFR(R-1)-OLFR(R-2)}\leq\frac{C(R)-C(R-1)}{C(R-1)-C(R-2)}$$
Since we do not know the value of the right side, we cannot define when this holds.

Due to the way we calculate the order line fill rate and the holding costs, we cannot say immediately if inequality (5.23) is met for all reorder points of all spare parts in the real life data of Vanderlande. We have therefore calculated the values of the delta functions for reorder points up to an order line fill rate of 99.99% for 1000 spare parts, randomly chosen from the real life data set.

From these 1000 spare parts, for 672 spare parts the delta function was non-increasing for all reorder points until an order line fill rate of 99.99% was achieved. For 115 spare parts, it occurred more than five times – i.e. when increasing the reorder point by 1 – that the delta function increased instead of decreased. For the remaining 213 spare parts this occurred at most five times. In many cases, if the delta function increased, the increase was very small.

Based on these results we conclude that the greedy heuristic will lead to close to optimal solution for the real life data of Vanderlande. The following section formulates this greedy heuristic.

5.4.6. Greedy heuristic

The starting point of the greedy heuristic is a minimal reorder point for all spare parts. In section 4.2.1 we haves seen that this means that all reorder points are -1 in the starting situation, despite the replenishment order quantity. From this start situation we calculate the delta function for all spare parts and choose to increase the reorder point for the spare part with the highest delta function. This procedure is repeated until the target order fill rate is achieved. The heuristic is formulated as follows:

1) Set the joint order line fill rate TOLFR:=0, and the total costs TC:=0
2) For all spare parts $i \in I$:
    a. set R_i=-1
    b. Calculate the order line fill rate $OLFR_i(R_i)$ and add $p_i OLFR_i(R_i)$ to TOLFR
    c. Calculate the total holding costs $C_i(R_i)$ and add $C_i(R_i)$ to TC
    d. Calculate $\Delta_i(0)$ following (5.22)
3) If TOLFR≥X, STOP.
4) Choose the spare part j such that: $\Delta_j\left(R_j+1\right) \geq \Delta_i\left(R_i+1\right) \qquad \forall i \in I$
5) Set R_j:=R_j+1, update $OLFR_j(R_j)$ and TOLFR, $C_j(R_j)$ and TC.
6) If TOLFR≥X, STOP, else calculate $\Delta_j(R_j+1)$ and go to step 4).

The results of this heuristic are discussed in the next chapter, where real life data is applied to the model.

5.5. Conclusion

In this chapter we have extended the order line fill rate model of chapter 4 by including variable replenishment lead times, fixed delivery days of suppliers and a timeframe. By conditioning on the replenishment lead time, the compound Poisson method as developed in the basic model yields order line fill rates that are close to the outcomes of a simulation study in all extensions.

With respect to spare parts orders, we assume that spare parts are uncorrelated; this means that the probability that a spare part is demanded in a customer order is independent of the other spare parts this customer orders. This assumption creates the possibility to approximate the order fill rate by using the joint order line fill rate of all spare parts, combined with the probability distribution of the number of order lines in a customer order. Using a fictional set of spare parts we have shown that this approximation is close to the actual order fill rate.

We have also validated the order fill rate model with real life data from Vanderlande, from the period April 2008 – August 2009. The outcome of the model differs significantly from the actual order fill rate performance, but this difference is fully explained by the fact that in many cases replenishment lead times are shorter than assumed due to the availability of spare parts in the warehouse of the project organization of Vanderlande.

In the next chapter the order fill rate model is applied to the complete set of data of Vanderlande and different settings of e.g. the target order fill rate and timeframe are analyzed.

6. Results and analysis

Before we can generate and analyze the results based on the real-life data, we need to subtract this data from the available information. Section 6.1 discusses which data is available and how the desired data is deduced from the available data. In section 6.2 the results of the order fill rate model with real life data are discussed.

6.1. Data

We discuss which data is available in section 6.1.1. Section 6.1.2 describes how the desired data is deduced from this.

6.1.1. Available data

Nearly all service related orders at Vanderlande are labeled as “SI-orders”. These SI-orders consist of all emergency, replenishment and preventive maintenance orders, plus all the orders for RMR activities that are handled by the service department; the latter mostly concerns orders for spare parts for revisions. Orders for spare part packages are labeled differently and hence these data can be separated from the other demand streams. Within the SI-orders there is currently no separation between the different demand streams, therefore it is hard to determine to which demand streams an SI-order belongs. Emergency orders can sometimes be recognized on a manually added description (e.g. “urgent” or “emergency”) and the way they are shipped – mostly with a fast courier. Orders for RMR activities usually consist of larger quantities but in most cases we do not know for sure if an order is for an RMR activity or that it is a large replenishment order. The majority of spare parts for RMR activities is not ordered through service orders, hence the number of service orders related to RMR activities is very small.

Since both the emergency orders as the orders for RMR activities are quite rare, and orders for spare part packages are not included in the data on SI-orders, we know that almost all SI-orders are either a replenishment order or a preventive maintenance order. A division between these two demand streams has not been made by Vanderlande, and the demand streams are too similar to make this division based on the order characteristics. On top of that, detailed analysis – i.e. looking into the correspondence between the service department and the local engineer – of some SI-orders has shown that in many cases a SI-order consists of both order lines for preventive maintenance and for replenishment. Data from over 25,000 SI-orders is available, which makes it undoable to manually divide these orders in preventive maintenance and replenishment orders based on these correspondences. In section 6.5 we discuss the consequences of a separation of those demand streams under the assumption that both demand streams have the same characteristics.

The data of each SI-order consists, amongst others, of the following elements:

- All ordered spare parts in this order, with per order line the ordered quantity
- Date on which the order has been placed by the customer
- Original delivery date as agreed with the customer
- Actual date on which the order is shipped
- Customer to which the order should be delivered

Beside these data on orders, we have some data on each spare part as well. These include:

- Cost price

- Supplier
- Lead time from supplier to central service warehouse, as given by the supplier.
- If applicable, minimum or fixed order quantity, as agreed with the supplier

Finally, we have some information on the replenishment orders that arrive at the service central warehouse:

- Supplier of the replenishment order
- All spare parts that have been delivered in the replenishment order
- Date on which the replenishment order is placed by the service department
- Original delivery date as agreed with the supplier
- Actual delivery date at the central service warehouse

6.1.2. Desired data

We can divide the desired data into three groups:

- Data per spare part
- Data per supplier
- Data on service orders in general

In the next subsections we define the desired data per group, and indicate how we extract these from the available data.

Spare part data

In our model a number of parameters is defined per spare part.

- **Demand rate (λ):** This is the arrival rate of orders for a spare part. We measure this rate per day. To calculate this rate, we use data on the number of orders that included this spare part in the past years in the following way:
  - Number of orders in the past twelve months: N_1
  - Number of orders placed 13 to 24 months ago: N_2
  - Number of orders placed 25 to 36 months ago: N_3
  - Number of orders per twelve months, up to 96 months ago: N_4 to N_8

  To each N_i we assign a weight w_i, where the weights are non-increasing, so:

  $w_1 \geq w_2 \geq w_3 \geq \ldots \geq w_8 \geq 0$ and $\Sigma w_i = 1$

  The yearly demand rate is then calculated as the weighted average of the amounts N_i. The daily demand rate follows from this by dividing with the number of workdays in a year, 260:

$$\lambda = \frac{\sum_{i=1}^{8} w_i N_i}{260}$$

  The weights are determined in consultation with employees of the service department. The weights are chosen equal for all spare parts. Initially the following weights are chosen:

$$w_1 = 0.40;\ w_2 = 0.30;\ w_3 = 0.15;\ w_4 = 0.10;\ w_5 = 0.05;\ w_6 = w_7 = w_8 = 0$$

  In section 6.3 we analyze what the influence of changing these weights is for the decisions made in the model.

A second option would be to not assign weights, but to use the values of N_i to find a trend in the demand rate and forecast the demand rate based on this trend. However, detailed analysis of the spare parts data has shown that currently there is no real trend within the demand rate of the spare parts. Hence this option is not worked out in our research.

- **Basic lead time (L):** Each spare part has a basic lead time that is obtained if the supplier delivers its replenishment in time. This lead time is copied from the available data and is measured in workdays. The variability in the lead time, leading to the probability distribution of the stochastic variable L, is incorporated in the delay distribution of the supplier. This delay distribution is subtracted from the data on replenishment orders. The basic replenishment lead times cannot be subtracted from this data set due to the following:

  In many replenishment orders, multiple spare parts are ordered. The lead time of the replenishment order is based on the lead times of the individual spare parts. This means that in some cases spare parts with short basic lead times have a longer actual replenishment lead time, due to the fact that they are shipped together with spare parts with higher lead times. Besides, in case a customer requests a spare part in eight weeks, in most cases the replenishment order is requested in e.g. seven weeks. Including such orders in the determination of the lead time distribution results in longer basic replenishment lead times.
- **Order line quantity distribution (F):** Per spare part we need the distribution of the order line quantity. We derive the distribution from the historical data through an empirical distribution. This distribution is composed as the weighted average of the empirical distributions of the customer order quantity per twelve months, using weighing factors like in the calculation of the demand rate. However, we use different weighing factors than we did for the determination of the demand rate. At the determination of the demand rate, we only look to historical orders of at most five years ago. If a spare parts is demanded for the last time six years ago, the probability that it will be ordered again is (near) zero, so we set the demand rate to zero. However, the ordered quantity of a spare part within an order of eight years ago is still of value to determine the order quantity distribution. The weights we use hence are:

  $$w_1 = 0.40;\ w_2 = 0.20;\ w_3 = 0.15;\ w_4 = 0.10;\ w_5 = 0.05;\ w_6 = 0.05;\ w_7 = 0.03;\ w_8 = 0.02$$

  These weights are adjusted in case there are years with no orders for this spare part, such that the total weight sums up to 1 again. If e.g. there are only orders in the second and fifth period, the new weights are w_2=0.80 and w_5=0.20. The other weights are 0.

  For a large number of spare parts we only have data available on a few orders. The characteristics of spare parts are however such that grouping spare parts is not possible. We therefore use an empirical distribution for these spare parts also. This means that if we only have data on one order, and the order quantity in this order is e.g. 5, the probability that for this spare part the order quantity is 5 equals 1.
- **Holding costs per unit (h):** The price to keep a spare part in stock, measured in Euros per day. The holdings costs are calculated as a percentage of the cost price. This percentage is based on both indirect costs (investment costs) and direct costs (costs per pallet of shelve in the warehouse) and is determined in consultation with financial experts within Vanderlande. In practice the indirect costs are 15% of the total value of the inventory per year; the direct costs are not determined per spare part but per pallet or shelve. We do not have information about the size of each spare part so the direct costs need to be estimated. Based on the total

amount that is calculated by the warehouse owner for storage space and other handling costs – as stated in section 2.4 – we estimate the direct costs to 15% of the cost price per year. The total holding costs hence are 30% of the cost price per year, or 30/260≈0.115% of the cost price per day.

- **Minimum and fixed order quantity (MOQ,FOQ):** For some spare parts, agreements with the supplier have been made regarding order quantities; a minimum quantity and/or only integer multiples of a certain amount may be ordered. The minimum and fixed order quantity are copied from the available data.
- **Supplier:** Each spare part has its unique supplier, which is known from the available data on spare parts. Based on the characteristics of this supplier we define the delay on the lead time of the spare part.

Supplier data

- **Fixed delivery days:** Some suppliers only deliver replenishment orders on fixed days of the week. These days are known within the service department.
- **Delivery delay distribution:** Suppliers do not always achieve the lead time they promise. The delivery delay distribution indicates what the probability is that a supplier delivers e.g. one day too late. This data is subtracted from the data on replenishment orders, as the difference between the promised and the actual delivery data of replenishment orders for each supplier.

Service order data

- **Customer order size distribution:** Orders may consist of more than one order line. The order size distribution gives the probability that an order has a certain amount of order lines. This distribution is derived as a weighted average of the number of order lines in orders of the past eight years, using the same weighing factors as for the determination of the demand rate.
- **Fixed replenishment order costs:** To handle a replenishment order, costs are made, e.g. personnel costs for the employees that enter the order into the system and costs for picking and packing the order. The major part of these costs can be deduced from the handling costs that are calculated by the warehouse owner. Per order line received by the warehouse, the average costs in 2008 were around €15,-. The costs of handling a replenishment order by the service department are hence estimated on €20,- per order line.

6.2. Results

We can now apply the model to real life data. We first investigate the outcomes of the model for a setup similar to the current situation. This setup is defined in section 6.2.1. In section 6.2.2 we investigate if the outcomes meet our expectations, or adaptations to the model need to be made.

6.2.1. Basic scenario

The basic scenario is defined as follows:

- The timeframe equals 15 (work)days
- The target time-based order fill rate equals 60.2%, equal to the current performance
- The weights to determine the demand rate are:

$$w_1 = 0.40;\ w_2 = 0.30;\ w_3 = 0.15;\ w_4 = 0.10;\ w_5 = 0.05;\ w_6 = w_7 = w_8 = 0$$

- The weights to determine the customer order size distribution are:

$$w_1 = 0.40;\ w_2 = 0.20;\ w_3 = 0.15;\ w_4 = 0.10$$
$$w_5 = 0.05;\ w_6 = 0.05;\ w_7 = 0.03;\ w_8 = 0.02$$

Entering the basic scenario in the model yields the following results:

Measure	Value
Order fill rate	60.21%
Holding costs	€42,847
#Spare parts on stock	1,375

Table 16 Key results of basic scenario

The holding costs and the number of different spare parts are high, more than twice the number of spare parts that is currently on stock. Detailed analysis – shown in Appendix E.1 – shows that from the 1,375 spare parts that are put on stock, 675 items are demanded three times or less in total in the past eight years. This is mainly caused by replenishment order quantities that are larger than one; since the reorder point is at least -1, an order quantity of at least 2 means that the spare part is held on stock, even if the demand rate of the spare part is very low. Especially spare parts that have a minimum order quantity larger than one are hence put on stock and contribute heavily to the total holding costs. An example of three spare parts largely contributing to the total holding costs is given in Table 17.

Spare part	Reorder Point	(minimum) Order Quantity	Number of orders since 2001	Yearly holding costs
Item A	-1	1000	1	€1,846
Item B	-1	1000	3	€1,910
Item C	-1	1000	1	€1,878

Table 17 Spare parts with high holding costs

6.2.2. Model adjustments

In the previous section we have seen that high minimum order quantities can lead to undesired spare parts on stock and high holding costs as a consequence. In this section we look for a solution to this problem.

In section 4.2.1 we have discussed that in many cases alternatives exist to prevent or reduce the large replenishment order; consequently these spare parts do not need to be kept on stock. This allows us to change the starting point of our heuristic. Instead of starting with spare parts on stock, since the reorder point R equals -1 and the replenishment order quantity Q is larger than 1, we now start with no spare parts on stock; consequently, R:=-Q. In case this implies a reorder point of less than -1, the first delta function that is calculated in step 2 of our heuristic is defined as follows:

$$\Delta = \frac{OLFR(-1) - OLFR(-Q)}{C(-1) - C(-Q)}$$

Hence, if we decide to put the spare part on stock, we immediately set the reorder point to -1. If the heuristic terminates because the target is reached and the reorder point is equal to –Q (with Q>1), we define R:=-1, Q:=1, independent on any constraints on the replenishment order quantity Q.

In our model, we process this as follows. In general the reorder point for all spare parts equals -Q, as described above. However, for spare parts that are demanded more frequently, say $\lambda \geq \mu$, and have a minimum order quantity larger than 1, we set the reorder point initially to -1, like we did in the original model. The value of μ can be changed by the user. By setting $\mu=0$ all spare parts that have a minimum order quantity larger than 1 have a reorder point of -1 as starting point; $\mu=\infty$ corresponds to the situation where all reorder points are initially set to -Q. Our basic value of μ is 12/260, i.e. only for spare parts that are demanded on average once per month, the reorder point is initially set to -1.

Note that this only holds for spare parts that have a minimum order quantity, so for spare parts that do not have a minimum order quantity and the calculated replenishment order quantity is larger than 1 we do not apply this. These spare parts all have -Q as starting point for the reorder point. In case the reorder point is -Q at the termination of the heuristic, the order quantity is set to 1 and the reorder point to -1.

This adaptation of our model has some consequences for the results of our basic scenario. Applying the basic value of μ (12/260) to the model yields the following results.

Measure	Value
Order fill rate	60.21%
Holding costs	€23,830
#Spare parts on stock	1,332

Table 18 Key results of alternative model

The holding costs have been strongly reduced by this adaptation of the model. The lower the value of μ, the higher the holding costs will be. In section 6.3.1 we analyze the results of different values for μ.

The number of spare parts on stock is still high in the situation described above, more than twice the amount that is currently held on stock. In Appendix E.1 we show that 665 spare parts of the 1,332 spare parts in total are demanded at most three times in the past eight years. We therefore suggest a second adaptation of our model, by excluding spare parts that have a demand rate below a certain threshold δ. If $\delta>0$ all spare parts that have a demand rate of less than δ receive a delta value of 0 in the heuristic and hence they are never put on stock. For these spare parts the reorder point equals -1 and the replenishment order quantity equals 1. Our basic value of δ is 1/260; in section ⍰ we analyze other values of δ, and we discuss the consequences for the maximum order fill rate that can be obtained by excluding low demand spare parts.

The results of the model including the second adaptation, with $\delta=1/260$, are as follows.

Measure	Value
Order fill rate	60.21%
Holding costs	€29,066
#Spare parts on stock	706

Table 19 Key results of alternative model + excluding low demand spare parts

The costs have increased with over €5,000, but the number of spare parts on stock has decreased to a more acceptable value, close to the current situation. The increase in holding costs is explained by the fact that for low demand spare parts – that are excluded now – only a few spare parts need to be put on stock to achieve a high order fill rate; hence the holding costs are low. By excluding these spare parts, other spare parts are put on stock in larger amounts, hence the holding costs increase.

In these adaptations, the minimum demand rate for inclusion (δ) overrules the minimum demand rate to remain the original order quantity (μ). This means that if μ=0 and δ=1/260, all spare parts that have a demand rate lower than 1/260 are not put on stock, even if their minimum order quantity is larger than 1 and hence the value of μ would imply that the spare parts should be put on stock.

The value of the parameter δ influences the outcomes of the model. If we exclude spare parts due to their low demand rate, the maximum order fill rate that can be achieved becomes less than 100%. Due to short replenishment lead times, some spare parts that are not held on stock can still be delivered within the timeframe, since the lead time is less than this timeframe. Hence, the timeframe influences the maximum achievable order fill rate as well. Table 20 shows the maximum order fill rates that can be achieved for different values of these parameters. The final column shows the number of spare parts included in the model following the value of δ.

δ\Timeframe	0	10	15	20	25	#Spare parts
0	100.00%	100.00%	100.00%	100.00%	100.00%	7,454
0.5	76.95%	80.93%	84.33%	97.21%	99.28%	2,582
1	61.03%	67.51%	73.28%	91.91%	95.74%	1,387
2	47.36%	56.23%	65.20%	87.86%	93.89%	749
3	39.00%	48.98%	59.76%	84.03%	92.02%	485

Table 20 Maximum time-based order fill rate that can be achieved for different values of δ and the timeframe

By allowing all spare parts, theoretically an order fill rate of 100% is possible. The large step between maximum order fill rates between the timeframes 15 and 20 workdays is explained by the fact that many spare parts have a replenishment lead time of 15 days. As becomes clear from the final column of Table 20 nearly 5,000 spare parts have a demand rate lower than 0.5/260.

Summarizing, the initial model of section 5.4.6 results in more spare parts on stock than currently, including some spare parts with relatively high replenishment order quantities compared to the yearly demand or low demand rates. We have introduced two parameters to control these problems. In the next section we investigate how the outcomes of the final model relate to the current situation with respect to which spare parts are held on stock.

6.2.3. Comparison to current stock levels

We have compared the results of the final model (Table 19) to the current situation at Vanderlande. Currently, there are 632 spare parts that are permanently held on stock at the service central warehouse. From these spare parts, 353 spare parts are held on stock in the proposed situation as well, hence 279 spare parts are not held on stock anymore. This is caused by a number of reasons:

- 200 spare parts have a replenishment lead time of at most 10 workdays. These spare parts do not have to be kept on stock because in many cases the spare part can be delivered by the supplier within the timeframe of 15 days. The exact order line fill rates for these spare parts of course depend on the delay distribution of the corresponding suppliers.
- An additional 33 spare parts have a replenishment lead time of 15 workdays. For these spare parts only a few parts have to be held on stock.

- For 23 of the remaining 46 spare parts, the minimum order quantity exceeds the average yearly demand; for 7 of these spare parts the average yearly demand is even less than 1 spare part.
- The remaining spare parts are in most cases relatively expensive or their demand rate is too low and hence they are uninteresting to put on stock.

The 353 spare parts that are not put on stock in the current situation, but are put on stock in the proposal, are partly characterized by:

- 37 spare parts have been demanded more than 20 times in the past eight years.
- 36 spare parts have a replenishment lead time of more than 25 days.
- 100 spare parts have a cost price of less than €5, and hence are relatively cheap to put on stock.
- 57 spare parts have a supplier with either always a positive delay or a probability of on-time delivery of less than 80%.

6.3. Analysis of parameters

We have a number of parameters that can be changed in our model. We want to analyze how changing these parameters influences the outcomes of our model in terms of time-based order fill rate, holding costs and number of spare parts that are held on stock. In each of the next sections we vary the value of one of the parameters μ, δ, the timeframe and the weights to determine the demand rate. The other parameters are equal to the value in the basic scenario discussed in section 6.2.1.

6.3.1. Minimum demand rate to remain original order quantity

The basic value for this parameter equals μ=12/260. Executing the model for different values of this parameter has led to the following results:

- As the value of μ increases, the holding costs decrease, as shown in Figure 10.

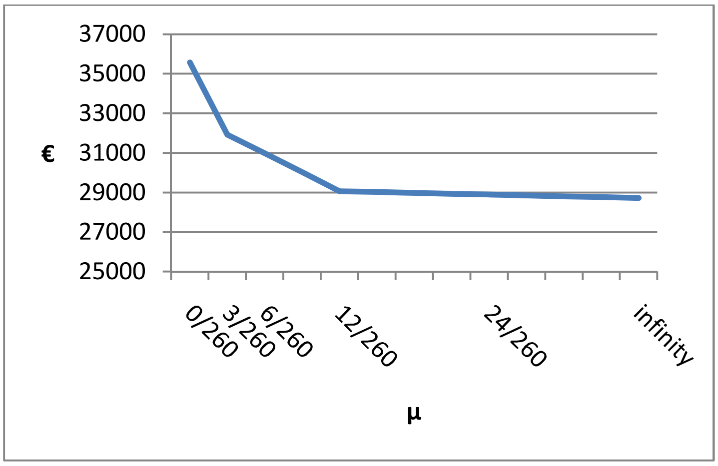

Figure 10 Holding costs at different values of μ for target fill rate of 60.2%

This decrease in holding costs is explained by the fact that the less spare parts with high replenishment order quantities we put on stock initially, the more freedom we have in choosing the ‘biggest bang for the buck’, i.e. the increase in order fill rate against the lowest costs. For μ=0 we might put spare parts on stock in larger amounts than its yearly demand (see Table 17) causing the holding costs to be higher.

- At a target time-based order fill rate of 70%, the holding costs merely decrease; at μ=0 the holding costs are €104,869, at $\mu=\infty$ these costs are €104,271. This means that in this case, nearly all spare parts with a minimum order quantity larger than 1 are put on stock, either initially (μ=0) or during the heuristic (μ>0).

6.3.2. Minimum demand rate for inclusion

The basic value for this parameter equals δ=1/260. Varying this value has led to the following results.

- As the value of δ increases, the holding costs increase and the number of spare parts decrease for a fixed value of the order fill rate. At a target order fill rate of 50%, for δ=0 the holding costs are equal to €12,709 and 694 spare parts are held on stock; for δ=3/260 the holding costs are €15,201 and 254 spare parts are held on stock. The decrease in number of spare parts on stock is explained by the fact that less spare parts are allowed to be held on stock if δ increases; the holding costs increase because high demand spare parts require a higher stock than low demand spare parts. Excluding low demand spare parts means that larger amounts of spare parts are held on stock, hence the holding costs increase.
- At a fixed budget of €100,000, the time-based order fill rate that is achieved with respect to the timeframe of 15 workdays decreases from 80.42% (δ=0) to 59.05% (δ=3/260). The higher the value of δ, the lower the maximum order fill rate that can be achieved, as shown in Table 20. At high values of δ, the order fill rate at holding costs of €100,000 come close to the maximum order fill rate that is theoretically possible. This means that nearly all available spare parts are held on stock in such an amount that the order line fill rate of each spare part is close to 100%.

6.3.3. Timeframe

The timeframe determines the period within which we want to optimize the performance. The higher the timeframe, the more spare parts can be delivered by the supplier within the timeframe and hence the easier it becomes to achieve a high order fill rate. In case spare parts are excluded due to low demand rates (i.e. δ>0), the timeframe influences the maximum order fill rate that can be achieved, as shown in section 6.2.2. Until now we have focused on the time-based order fill rate with respect to a timeframe of 15 workdays. We have varied the length of the timeframe, leading to the following results:

- At an equal order fill rate, the holding costs decrease as the timeframe increases. At a timeframe of 0 days, the performance is purely based on the ability to deliver from stock. Hence to achieve a high order fill rate, many spare parts needs to be held on stock. For higher timeframes, there are a number of spare parts for which the replenishment lead time is shorter than the timeframe; for these parts the order line fill rate is 100% without having to take any spare part on stock. The holding costs thereby decrease for increasing timeframes. At a target order fill rate of 60%, the holding costs at a timeframe of 0 days are nearly €800,000, at a timeframe of 15 days these costs are roughly €28,500 and for timeframes of 20 days and longer the order fill rate after initially putting spare parts on stock is already more than 60%.

We have to note that comparing the time-based order fill rates for different timeframes is hard. A time-based order fill rate of 90% with respect to a timeframe of 25 days does not necessarily have to be a better performance than an order fill rate of 70% with respect to a timeframe of 15 days.

6.3.4. Demand rate weights

In consultation with members of the Team Worldwide basic values for the weights that determine the demand rate of spare parts are agreed. In this section we enter some different weights to analyze the influence of these weights to the outcomes of the model. In Table 21 the values of these weights are shown.

Weight	Basic scenario	Alternative 1	Alternative 2	Alternative 3
w_1	0.40	0.60	0.25	0.20
w_2	0.30	0.30	0.20	0.20
w_3	0.15	0.10	0.15	0.20
w_4	0.10	0	0.10	0.20
w_5	0.05	0	0.10	0.20
w_6	0	0	0.10	0
w_7	0	0	0.05	0
w_8	0	0	0.05	0

Table 21 Scenarios for analysis of demand rate weights

From our analysis we conclude that – as we would expect – the weights mainly determine which spare parts are held on stock; either spare parts that are demanded in the past one or two year (alternative 1) or spare parts that have also been demanded six to eight years ago (alternative 2). In the second and third alternative spare parts that have been demanded three, four or even more times in the past year are not held on stock anymore due to the low weights that the past two years have received in those alternatives.

6.4. Scenarios

Until now we have varied only one parameter at a time, to give an idea of the influences of these parameters to the model. We are also interested in the outcomes of the model and deeper analysis of these results for some realistic scenarios, i.e. values of the parameters that are relevant for the spare parts supply chain at Vanderlande. For each scenario we investigate how the order fill rate relates to the holding costs, and what – for given values of the reorder points and replenishment order quantities – the time-based order fill rate becomes for other timeframes.

We have chosen five scenarios to analyze into detail. Each scenario is characterized by values of the timeframe, δ and μ. The weights that determine the demand rate are fixed and equal to the basic values from section 6.2.1. The scenarios we investigate are:

A) The basic scenario, so the timeframe is 15 days, $\delta=1/260$, $\mu=12/260$.
B) Delivery from stock for medium- and high-demand spare parts: timeframe=0 days, $\delta=3/260$, $\mu=3/260$.
C) One week improvement compared to basic scenario: timeframe=10 days, $\delta=1/260$, $\mu=12/260$.
D) No constraints on spare parts: timeframe=15 days, $\delta=0$, $\mu=0$.
E) Longer timeframe and less constraints: timeframe=20 days, $\delta=0.5/260$, $\mu=3/260$.

The detailed results of these scenarios are found in Appendix E.2. The main conclusions are:

- From yearly physical holding costs of €70,000 onwards, increasing the holding costs with €10,000 leads to an increase of the order fill rate of less than 1%. From €100,000 onwards, the holding costs need to be increased with more €100,000 to achieve an increase in order fill rate of 1%. Hence, very high holding costs only lead to slightly better order fill rates. For relatively low holding costs (i.e. around €50,000) we achieve an order fill rate close to the maximum value that can be obtained. The only exception is scenario D, where we allow all spare parts to be held on stock. In this scenario, the holding costs even increase with nearly 1% when increasing the holding costs from €300,000 to €400,000. This is caused by the fact that all spare parts can be held on stock, so to achieve the maximum order fill rate, the holding costs need to be high.
- The chosen timeframe influences the time-based order fill rates for other timeframes as well. Executing the model for a timeframe of 0 leads to different time-based order fill rates for other timeframes as executing the model for a timeframe of 20 days. In the latter case, spare parts with low replenishment lead times are not held on stock, and increasing the timeframe from 0 to e.g. 30 days yields an increase from 0% to 100% for the order line fill rate of these spare parts. At an initial timeframe of 0 days, the order line fill rate of these low replenishment lead time spare parts may already be more than 0%, so the increase in order line fill rate is less when increasing the timeframe.

6.5. Separation of service orders

Until now we have assumed one demand stream for service orders. In section 2.3.1 we have discussed that these service orders consist of orders for preventive maintenance and orders for replenishment. We are interested in the consequences of separating those demand streams with respect to the performance that can be achieved for each demand stream.

According to Vanderlande, 30% of all order lines of service orders are related to preventive maintenance. We have analyzed the influence of separating the demand streams in two ways:

- We investigate the performance for preventive maintenance orders only using the reorder points and replenishment order quantities from the original model (i.e. using all service orders) as input. The demand rate for preventive maintenance orders is calculated as 30% of the demand rate for service orders.
- We calculate the reorder points and replenishment order quantities for the preventive maintenance demand stream only, and compare this to the outcomes of the model including all service orders. The difference of those stock levels is the applied to the replenishment demand stream to observe the performance for that stream.

The second analysis only gives an indication of the actual performance if we would separate the demand streams, since we now assume independent stocks for both demand streams. In practice, separating these demand streams will not imply two separate stocks, and the performance of at least one of the demand streams is higher since both streams make use of the combined stock of spare parts.

Detailed results are shown in Appendix E.3. From the analysis we have concluded the following:

- If we allocate the complete stock to preventive maintenance, the performance only increases marginally (at most 3% at 50% order fill rate). This is caused by the fact that if a spare part is held on stock, the order line fill rate of that spare part in the original model is already close to 100%. Reducing the demand rate of that spare part only slightly improves the order line fill rate for this spare part and hence the total order fill rate is not changed by much either.
- If we calculate the optimal stock levels for preventive maintenance, the remaining stock leads to an order *line* fill rate that is in most cases lower than the order fill rate, so the performance for replenishment orders is decreased heavily. As stated above, in practice the performance will be higher for one or both demand streams but the performance of replenishment orders will definitely be worse than in the original model where both demand streams are taken together.

The performance for replenishment orders decreases heavily by separating the demand streams, in case we set the desired order fill rate for preventive maintenance orders equal to the original desired order fill rate. One of the goals of separating demand streams is that we are able to increase the performance of the most important demand stream, preventive maintenance in this case. If we would increase the performance of preventive maintenance orders, the performance of replenishment orders will decrease even more. We therefore do not think it is profitable to separate the demand streams.

6.6. Conclusion

In this chapter we have applied the model resulting from the previous chapter to the actual situation at Vanderlande. To be able to apply the model we have described how the desired data for the model is subtracted from the available data. The basic model yielded relatively high holding costs and many spare parts with low demand rates on stock. We have introduced two parameters to control the holding costs and number of spare parts on stock.

We analyzed the influence of the different parameters to the outcomes of the model. The most important conclusions are:

- By excluding spare parts with low demand rates (i.e. demand rates below δ) the maximum order fill rate that can be achieved decreases. This maximum is also dependent of the timeframe. As the value of δ increases, the number of spare parts decrease but the holding costs increase if the target order fill rate remains equal. The increase in holding costs is caused by the fact that for medium- and high-demand spare parts more parts are need to be held on stock to achieve a similar fill rate, hence the holding costs for these spare parts are higher.
- A number of spare parts have a minimum order quantity, imposed by the supplier. If the demand rate of such a spare part is at least equal to the parameter μ, we put this spare part on stock initially. As μ increases, the holding costs decrease, since in this case less spare parts that might not be interesting to put on stock are initially put on stock. If the target order fill rate is close to the maximum order fill rate that can be achieved, the holding costs merely decrease as μ increases. In this case all spare parts with minimum order quantities are put on stock, either initially or in the model.
- Increasing the timeframe increases the order fill rate for the same stock levels, and if the order fill rate is constant, increasing the timeframe decreases the holding costs. The longer

the timeframe, the more spare parts can be delivered within the timeframe by the supplier. These spare parts do not have to be held on stock and hence the holding costs are reduced. We have to note that comparing to time-based order fill rates with different timeframes is hard. A time-based order fill rate of 90% with respect to a timeframe of 25 days does not necessarily have to be a better performance than an order fill rate of 70% with respect to a timeframe of 15 days.

We have chosen five realistic scenarios to analyze further. In all scenarios, order fill rates close to the maximum value that can be obtained are achieved for relatively low holding costs (i.e. less than €70,000). Increasing the holding costs further leads to a very slow increase of the order fill rate, less than 1% per €100,000 additional holding costs for all but one scenario. The choice of the timeframe influences the performance with respect to the order fill rates for other timeframes; initially choosing a timeframe of 0 days leads to some spare parts on stock with low replenishment lead times, while at a large timeframe mostly high replenishment lead time spare parts are preferred.

For each of the scenarios we have investigated the consequences of separating the preventive maintenance and replenishment order demand streams. Allocating the complete stock to preventive maintenance orders only increases the order fill rate marginally, due to the fact that for most spare parts that are held on stock, the order line fill rate is close to 100%; the performance hence cannot be increased by much if the demand rate is reduced. Holding separate stocks for each demand stream is not advisable either; at an equal performance for preventive maintenance orders (compared to the original performance for service orders) the order *line* fill rate for replenishment orders is much lower. If we increase the performance for preventive maintenance orders, the performance for replenishment orders becomes even worse.

7. Implementation

One of the goals of this research is to deliver a ready-to-use end product, so that the service department is able to apply the model to the real-life data at any moment. To achieve this, we have implemented the following elements:

- Templates to subtract relevant data from the database of service orders of Vanderlande.
- A spreadsheet macro that transforms the rough data to input files for the model, through xml-files in our case.
- An executable program ('Vanderlande service inventory tool') that uses the xml-files as input and calculates the optimal reorder points and replenishment order quantities following our model.

The goal is to make the process from subtracting the data to evaluating the inventory policy as easy as possible. At best, the user only needs to make a couple of actions to read the data into the model. The process of subtracting the data from the database using the templates and applying the spreadsheet macro to these files is described in manuals.

7.1. Software

7.1.1. Applying the model

The eventual users make use of two interfaces. The first interface is in a spreadsheet and is used to determine the input data of the model. The data extracted from the database is here processed to the input files. Before these input files are created, the user has the possibility to adapt the input data, e.g. by updating holding and orders costs or deleting orders that were not service orders.

The executable program has an interface as well. In this interface the user can execute the model by pressing a single button. However, the user has the possibility to set some parameters manually to other values. The parameters that can be changed by users are:

- The timeframe
- The target (time-based) order fill rate
- The weights that determine the demand rate
- The minimum demand rate that spare parts need to have to be included
- The minimum demand rate that spare parts need to have to remain the original replenishment order quantity

The final two parameters are based on the basic model adjustments, as discussed in section 6.2.2. Beside these parameters, the user can also decide to optimize the order fill rate for a given budget, instead of optimizing the costs for a given order fill rate.

7.1.2. Categorization of spare parts

Beside applying the model, the executable program has a second functionality as well. This functionality is not meant to optimize the order fill rate and holding costs, but to calculate the order fill rate and holding costs based on input given by the user. By allowing the user to manually assign performances to certain groups of spare parts, we hope to create a higher acceptance and understanding of the calculation model we use.

In this program, the user defines categories of spare parts. These categories are based on two criteria: holding costs and demand rate. Per criterion, the user selects up to three threshold values. Based on these threshold values, categories are created, as shown in the next example:

Example 8

The user selects two threshold values for the holding costs (h_1 and h_2) and three threshold values for the demand rate (λ_1, λ_2 and λ_3). These values are increasing, so $h_1<h_2$ and $\lambda_1<\lambda_2<\lambda_3$. Figure 11 shows how the categories are created from these threshold values.

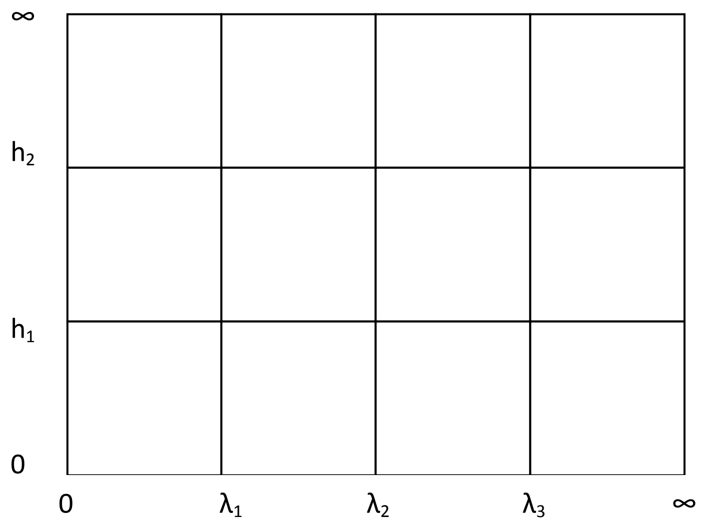

Figure 11 Example of categorization of spare parts

The user manually assigns an order line fill rate to each of the categories. In most cases, the categories for which the daily demand rate is below λ_1 and the categories for this the daily holding costs exceed h_2 receive an order line fill rate of 0%, but the user is free to choose differently.

If the user has assigned order line fill rates to all categories, the program can calculate the time-based order fill rate for a given timeframe (entered by the user as well). For each spare part, the reorder point is calculated for which the order line fill rate of the category the spare part belongs is achieved. The holding costs are calculated as well. The next section discusses the results of a number of case studies.

7.1.3. Results

We discuss four case studies in this section. In all cases we use the basic values of the weights to determine the demand rates (as mentioned in section 6.2). The value of μ equals 12/260. For all spare parts that are put on stock holds that the reorder point should be at least equal to -1. The specific characteristics of these case studies are as follows:

A) We create one category containing all spare parts; the target time-based order line fill rate for this category is 50%, with respect to a timeframe of 15 workdays.
B) We create nine categories; the borders of the categories are: λ_1=1/260, λ_2=12/260, h_1=(0.3*5)/260, h_2=(0.3*25)/260. The thresholds of the holding costs are based on the cost price of the spare parts (€5 and €25 respectively).

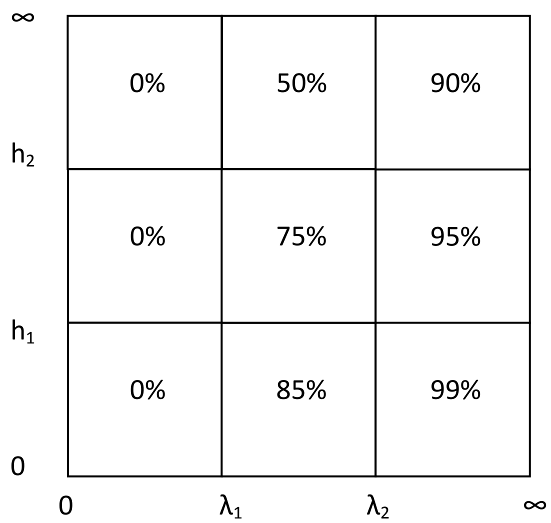

C) We create sixteen categories, with the following thresholds: λ_1, h_1, h_2 as in case study B), λ_2=6/260, λ_3=18/260, h_3=(0.3*100)/260. We assign the following time-based order fill rates, with respect to a timeframe of 20 workdays, to the categories:

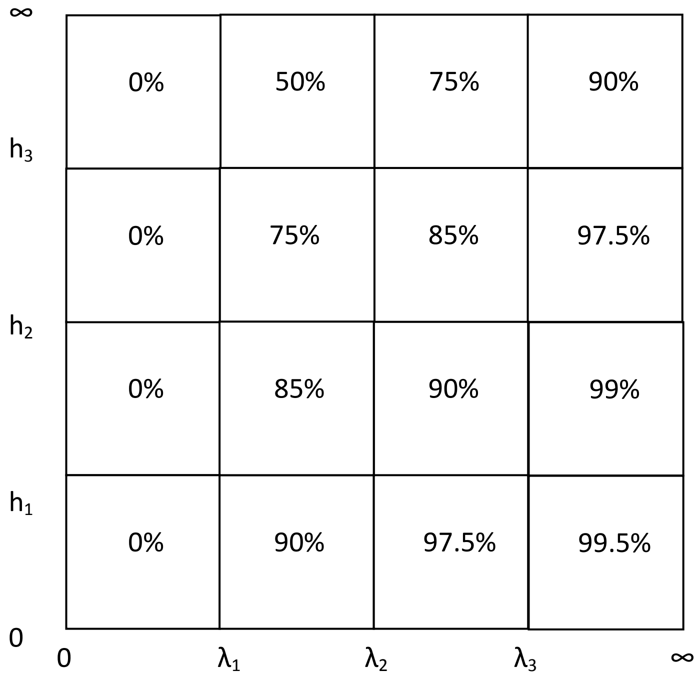

D) We create eight categories, with the following thresholds: λ_1=1/260, λ_2=6/260, λ_3=12/260, h_1=(0.3*50)/260. We assign the following order fill rate, with respect to a timeframe of 0 days, to the categories:

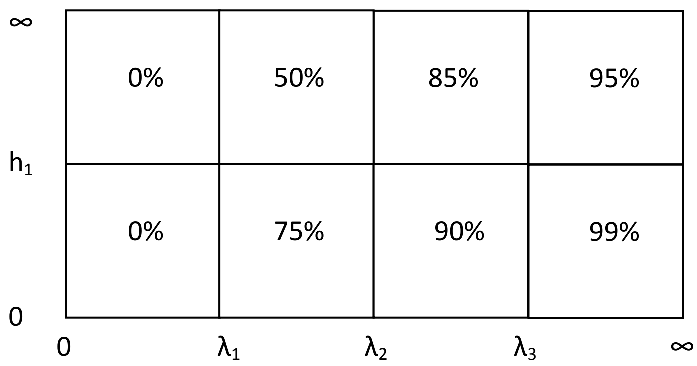

The next table shows the key results of these four case studies.

Measure	Case A)	Case B)	Case C)	Case D)
Order fill rate	73.37%	52.53%	86.96%	91.63%
Holding costs	€378,690	€34,392	€65,867	€40,887
#Spare parts on stock	4,518	348	362	175

Table 22 Key results of categorization case studies

The categorization leads to worse results than the optimization model of section 6.2.2. For roughly similar the same order fill rates, the outcomes of the experiments in e.g. section 6.3.3 show much lower holding costs. However, as stated earlier, this method is not to optimize the performance of the service central warehouse. The main goal of this option is to increase the level of acceptance and understanding at the users of our model.

7.2. Training

All functionalities that are described above have been clearly explained in manuals. Also, two training sessions have been scheduled with the future users of the model to give a live demonstration of all functionalities and to discuss the working of the model. Hereby the users have some on-hand experience of the service inventory tool and hence they are able to calculate the optimal stock levels without any knowledge of the model itself in the future.

8. Conclusions and recommendations

In this chapter we summarize the most important conclusions of our research, and we answer the research questions as stated in section 1.2. We make recommendations both to improve the performance of the current model and to extend the model.

8.1. Conclusions

The research problem was formulated as follows:

> *Develop an inventory policy for spare parts at the service central warehouse of Vanderlande Industries, taking into account the structure of the supply chain. Analyze how the stock levels resulting from the model influence the performance towards customers regarding on-time-delivery and customer order lead time.*

The research is roughly divided in two parts: 1) analysis of the supply chain of spare parts at Vanderlande and 2) development of an inventory model for the service central warehouse.

Analysis of the spare parts supply chain

Within the spare parts supply chain, we have recognized five different demand streams from customers:

- Emergency orders
- Preventive maintenance orders
- Replenishment orders
- Spare part packages
- Revisions, modifications and retrofits (RMR)

The only two demand streams that are of interest for the service central warehouse are preventive maintenance orders and replenishment orders. These orders are combined in one general demand stream ‘service orders’, characterized by customer orders of multiple order lines, and the quantity in which a spare part is ordered can be larger than one.

By assuming that spare parts are not correlated, i.e. the probability that a customer order contains a specific spare part does not depend on the other spare parts within that order, we are able to calculate the time-based order fill rate through the calculation of the performance per spare part. The resulting inventory model shows the following results.

Outcomes of the inventory model

Currently a time-based order fill rate of 60.2% is achieved with respect to a timeframe of 15 workdays. The total holding costs are around €37,000 per year and 632 spare parts are held on stock. Theoretically the current stock levels result in an order fill rate of around 53%; the difference is largely explained by the fact that in many cases spare parts can be picked from a second warehouse, normally intended for the project organization of Vanderlande. This reduces the replenishment lead times to a couple of days and hence the time-based order fill rate in practice exceeds the calculated performance.

For a target order fill rate of 60.2%, the inventory model results in 706 spare parts on stock and yearly holding costs of nearly €30,000. Hence, the costs are decreased with nearly 20% and the performance is at least equal to the current performance. Since in practice the time-based order fill rate will be higher than the calculated performance, the inventory model results in a better performance against less holding costs. From the current spare parts that are held on stock, 353 are held on stock in the proposed situation as well. The other spare parts are mainly characterized by low replenishment lead times; therefore these spare parts can be delivered within the timeframe by the supplier and it is not necessary to keep these spare parts on stock.

Analysis of different scenarios has shown that holding costs of more than €70,000 only lead to slightly higher order fill rates compared to the order fill rates for lower holding costs. From this point, increasing the holding costs with €100,000 leads in most cases to an increase in order fill rate of less than 1%.

8.2. Recommendations

In this section we formulate a number of recommendations to Vanderlande. We split the recommendations in two parts: 1) recommendations to improve the performance within the current model and 2) recommendations to extend the current model.

8.2.1. Recommendations for performance improvement

- Currently, the demand process of spare parts is purely based on historical data. Many spare parts have very low demand rates, especially for these spare parts it is difficult to find good parameters of the demand process since there is not much data to base the parameters on. We recommend that Vanderlande invests in the determination of a more accurate model to forecast spare parts demand, e.g. through a graduate research. A basis of this forecast model could be the installed base information that is available at Vanderlande.
- The holding costs per spare part are calculated as a percentage of the cost price in the current model. However, in reality the holding costs are calculated as a combination of a percentage of the cost price (indirect costs) and costs per used pallet or shelve in the warehouse (direct costs). The latter is not related to the cost price at all. For a better calculation of the holding costs, the costs per pallet or shelve need to be related to the spare part. The holding costs per spare part are best approximated by calculating them as a function of the volume of the spare part. Theretofore the dimensions of each spare part need to be known. We suggest that at least for the most frequently asked spare parts these dimensions are recorded in the database.
- The outcomes of our model analysis clearly show that the replenishment lead time heavily influences the time-based order fill rate, especially if the timeframe is set higher than the basic replenishment lead time of most spare parts. If these replenishment lead times can be reduced, the performance towards customers is increased without needing to invest in holding spare parts on stock at the service central warehouse.

8.2.2. Recommendations for model extensions

- In our research we have combined the two relevant demand streams – preventive maintenance orders and replenishment orders – into 'service orders'. Although we expect these two demand streams to be much alike, we suggest to record the type of demand for service orders in the future. Analysis of the demand for each of these demand streams can then be carried out to verify our assumption that the demand streams are alike. Besides, if the two demand streams are separately modeled, it is possible to adapt the desired performance per demand stream. More research needs to be conducted to find optimal stock values in case the two demand streams are separated and different desired performances are determined for each demand stream.

- Our model only observes the service central warehouse. A model that observes and optimizes the stock at system locations has been described in (van Sommeren, 2007). The next step is to combine the stock levels of the system locations and the service central warehouse into one model. The model of (van Sommeren, 2007) and our model can be used as a basis for this new model.

8.2.3. Recommendations for implementation of the model

- As discussed in chapter 7, the service department is able to apply the model through two interfaces. We recommend that preferably each month but at least once per three months new data is entered into the model and the resulting stock levels are analyzed. After execution of the model, the results of the model, i.e. the reorder points and replenishment order quantities for all spare parts that are suggest to be held on stock, should be compared to the previous results. There are three cases we distinguish:
    - A spare part is already held on stock, but the values of the reorder point and/or the replenishment order quantity have changed. In this case these changes should be copied.
    - A spare part that is not held on stock is present in the proposed situation. In this case there are two solutions: (1) immediately put the spare part on stock by ordering a sufficient amount – i.e. an amount between the R+1 and R+Q, with R and Q respectively the reorder point and the replenishment order quantity – or (2) wait until a customer order arrives and then order an amount such that the stock level lies between R+1 and R+Q. We suggest to choose option (1) for spare parts that have a relatively high demand rate since for these parts it is most likely that multiple customer orders arrive in the future. For lower demand rate spare parts, the possibility that no customer order will arrive in the (near) future is larger; therefore we suggest to choose option (2), where we wait until the next customer order for this spare part arrives before actually putting the spare part on stock.
    - In the current situation a spare part is held on stock, but in the proposed situation is it not held on stock anymore. In this case, we suggest to set the reorder point to the minimum value, implying that the stock is not replenished anymore. The remaining parts are not thrown away but they are used to fill customer orders that arrive.

  In the third case, it might occur that spare parts are left on stock for a long period since no customer orders arrive anymore for this spare part. The service department should register how many spare parts have not been demanded in e.g. the past year. If this is the case, possible solutions are selling the spare parts to a third party like the distribution center of Vanderlande, or selling the spare parts to a metal business to be recycled.

  By applying the procedure above, there is a smooth transition between the old stock levels and the suggested stock levels.

A. Definitions

Customer

A company that has one or more systems of Vanderlande installed, and requires *spare parts* to maintain this system.

Customer center

A division of Vanderlande, concerned with a part of the *customers*. Most customer centers are dedicated to a specific area of the world, but there is a customer center dedicated to all customers with distribution systems as well. The customer center forms the contact between the customer and the *service department*.

Customer order

If a customer needs *spare parts* to maintain his system, or to replenishment the stock of spare parts at the customer site, he places an order for spare parts at the service department. A customer order consists of *customer order lines*, each line having a certain *customer order line quantity*. The number of order lines determines the *customer order size*.

Customer order line

A customer order line is a part of a *customer order*. It contains one type of *spare part*, ordered in a certain amount. This amount is the *customer order line quantity*.

Customer order line quantity

The amount of a type of *spare part* that is ordered on a *customer order line*.

Customer order size

The number of *customer order lines* in a *customer order*.

Fixed order quantity

The *replenishment order quantity* needs to be an integer multiple of this quantity. The *supplier* determines this quantity, mostly based on the size of the package in which the *spare part* is sold.

Minimum order quantity

The *replenishment order quantity* needs to be at least equal to this quantity. The *supplier* determines this quantity, based on the size of the package in which the *spare part* is sold, or the total costs of the production of the spare part.

Order fill rate

The percentage of complete *customer orders* that can be filled from *stock*. Filling complete orders means that all *customer order lines* are filled completely from stock.

Order line fill rate

The percentage of complete *customer order lines* that can be filled from *stock*. Filling complete order lines means that the complete *customer order line quantity* is filled from stock.

Replenishment lead time

The number of days that it takes a *replenishment order* to arrive at the *service central warehouse* after the order has been placed at the *supplier*. Except for the basic model this is a stochastic variable in which the variability of the lead time, due to delays in the delivery of replenishment orders is taken into account.

Replenishment order

An order for *spare parts* made at the *supplier* in order to replenishment the stock at the *service central warehouse* or to fill a *customer order*. The ordered quantity is an integer multiple of the *replenishment order quantity*.

Replenishment order quantity

The fixed quantity of a *replenishment order*. In case a replenishment order needs to be placed, an integer multiple of this quantity is ordered. This quantity is subjected to a number of constraints; amongst others it needs to be at least equal to the *minimum order quantity* and an integer multiple of the *fixed order quantity*.

Service central warehouse

The warehouse where all *spare parts* are held on *stock*. Besides, the warehouse is used to consolidate *customer orders*. The daily activities, such as receiving *replenishment orders*, picking and packing, are outsourced to an external party. Throughout this thesis it is mostly denoted with “SCW”.

Service department

A department of Vanderlande, located at the headquarters in Veghel. Internally the department is called Supply Chain Management Services (SCMS). It consists of the *customer centers* International and Distribution, and the spanning Team Worldwide. The latter coordinates the *service central* warehouse and is responsible for the placement of *replenishment orders*.

Spare part

A part of the system of a *customer*, that is sold individually to maintain the system. At the *service central warehouse* a number of spare parts is held on *stock*. *Customer orders* consists of one or more *customer order lines*, where each line consists of a certain quantity of a single spare part.

Stock

Also called inventory; a number of *spare parts* are held on stock at the *service central warehouse*, in order to increase the *order (line) fill rate*.

Supplier

All *spare parts* are ordered at suppliers by the *service central warehouse*. This is either an external party or the factory of Vanderlande. The supplier may set restriction on the *replenishment order quantity* that is ordered at each *replenishment order*, through the *minimum and fixed order quantity*.

Time-based order fill rate

Percentage of all *customer orders* that is completely filled within a certain *timeframe*. This means that all *customer order lines* are filled completely within this timeframe.

Time-based order line fill rate

Percentage of all *customer order lines* that is completely filled within a certain *timeframe*. This means that the complete *customer order line quantity* is filled within this timeframe.

Timeframe

Period within which Vanderlande wants to deliver *customer orders*. The percentage of all orders that is delivered within this timeframe is denoted with *time-based order fill rate*.

B. Derivations

B.1. Order line fill rate calculation for the basic model

B.1.1. Using compound Poisson demand

We assume in this case that the customer demand is modeled by a compound Poisson process. This means that the inter arrival times for customers are exponentially distributed and the order quantity is determined from a (discrete) probability distribution. Because we assume exponential inter arrival times, the PASTA-property holds. This property states that in case customers arrive following a Poisson process – which implies exponential inter arrival times – the arriving customer observes the system (in our case, the inventory level) upon arrival in its steady state. This means that if in 50% of the time the inventory level is e.g. five, than 50% of all arriving customers will observe an inventory level of five upon arrival. The next counterexample shows that this does not hold in general.

Example 9

Suppose that a customer arrives every two days, at the start of the day, so the inter arrival time is constant and equal to two days. All customers order one spare part, for which the lead time is one day. The inventory level at the start is equal to one. When a customer arrives, he observes an inventory level of one. The takes the spare part and leaves, and the inventory level becomes zero. After one day the replenishment order arrives and the inventory level is one again, until the next customer arrives a day later. In total, the inventory level is equal to one 50% of the time, and equal to zero 50% of the time. However, 100% of the arriving customers observe an inventory level of one. Hence, arriving customers do not observe the inventory level in steady state, due to the deterministic arrival process of customers.

In our case each arriving customer observes the inventory level in steady state. By determining the steady state distribution of the inventory level we can calculate the probability that the inventory level is sufficiently large to fill the customer order. An inventory level of j is enough to fill all customer orders of at most j spare parts. Taking all possible values of the inventory level into account, the order line fill rate is defined as (Thorstenson, et al., 2008):

$$OLFR = \sum_{j=1}^{\infty} \Pr\left(IL = j\right) \Pr\left(F \leq j\right)$$

In section 4.1 we have seen that the inventory level is related to the inventory position. The inventory position equals the inventory level plus all outstanding replenishment orders that have not been received. Since the number of outstanding orders is not negative, the inventory level cannot exceed the maximum value of the inventory position.

The inventory position is used to determine when to order a replenishment order. Whenever the inventory position is equal to or lower than the reorder point R, the replenishment order quantity Q is ordered at the supplier. If the resulting inventory position (the ordered quantity is added to the outstanding orders and hence also to the inventory position) is still below the reorder point R, an additional integer multiple of Q is ordered until the resulting inventory position is larger than R. This policy ensures that the inventory position only takes values between R+1 and R+Q, so the inventory level cannot exceed R+Q as well. The calculation of the order line fill rate for given values of R and Q is hence given by:

$$OLFR(R,Q) = \sum_{j=1}^{R+Q} \Pr(IL = j)\Pr(F \leq j)$$

We assume that the probability distribution of the customer order quantity F is known. Next we determine the probability distribution of the inventory level by using the probability distribution of the inventory position.

In case the inventory position can obtain all values between its lower bound R+1 and upper bound R+Q, (Axsäter, 2006) has proven that the probability distribution of the inventory position is equal to the uniform distribution between R+1 and R+Q. This means that:

$$\Pr(IP = k) = \frac{1}{Q} \text{ for } k = R+1,\dots,R+Q$$

The next example shows that we can always use this uniform distribution of the inventory position:

Example 10

If the greatest common divisor (gcd) of all possible order quantities is equal to one, it is obvious that all values between R+1 and R+Q can be reached, so then the above is valid. In case the gcd is larger than one, we introduce a new spare part with a price that is equal to the 'old' price times the gcd, and divide all order quantities by the gcd. The gcd of the order quantities of this new spare part is equal to one again, and we are able to use the probability distribution of the inventory position as mentioned above.

Because in the basic model the replenishment lead time is fixed, there is another relationship between the inventory level and the inventory position. At an arbitrary time t, all replenishment orders that have been ordered at or before the time t-L (where L is the replenishment lead time) have arrived. All replenishment orders that have been ordered after the time t-L have not yet arrived. If no customers would arrive in the time between t-L and t, the inventory level at time t would be equal to the inventory position at time t-L, since the inventory position equals the inventory level plus the outstanding orders, and at time t all outstanding order from time t-L have been delivered. So without any arriving customers in the time between t-L and t we would have:

$$IL(t) = IP(t-L)$$

The actual inventory level at time t is hence equal to the inventory position at time t-L minus the total demand of all customers that have arrived between time t-L and t. We denote this by D(t-L,t), so the inventory level at time t is formulated as:

$$IL(t) = IP(t-L) - D(t-L,t)$$

The total demand during the period (t-L,t) is independent of the (arbitrary) time t, we therefore denote the demand during replenishment lead time L with D(L). The probability distribution of the inventory position is independent of the time t as well. We therefore rewrite the relation between the inventory level and inventory position as:

$$IL = IP - D(L)$$

The probability that the inventory level has a certain value – say j – hence equals the probability that the inventory position has a certain value – say k – times the probability that the customer demand during lead time equals the difference between j and k. The inventory level is at most R+Q, the inventory position can only take values between R+1 and R+Q. The probability distribution of the inventory level is hence given by (Axsäter, 2006):

$$\Pr(IL=j)=\sum_{k=\max(R+1,j)}^{R+Q}\Pr(IP=k)\Pr(D(L)=k-j) \qquad j\leq R+Q$$

$$\Pr(IL=j)=\frac{1}{Q}\sum_{k=\max(R+1,j)}^{R+Q}\Pr(D(L)=k-j) \qquad j\leq R+Q$$

The probability distribution of the demand during replenishment lead time is calculated by using the compound Poisson process that describes the customer orders. The probability that y customers order k-j spare parts is given by:

$$\Pr(\text{y customers order k-j spare parts})=f^{y}(k-j)$$

$f^{y}(k-j)$ is the y-fold convolution of the probability distribution of the customer order quantity. It is calculated recursively using the fact that the probability that y customers order k-j spare parts equals the probability that (y-1) customers order (k-j)-t spare parts times the probability that the y^th customer orders exactly t spare parts, for all possible values of t:

$$f^{y}(k-j)=\sum_{t=1}^{k-j}f^{(y-1)}(k-j-t)\Pr(F=t) \text{ for } y\geq 1 \text{ and}$$
$$f^{0}(0)=1$$
$$f^{0}(x)=0 \quad \forall x\neq 0$$

To calculate the probability that the demand during replenishment lead time equals k-j, we multiply the probability that y customers order k-j spare parts with the probability that during replenishment lead time y customers arrive. We therefore need the probability that y customers arrive in a period of length L. This is given by the Poisson distribution with rate λ, the arrival rate of customers:

$$\Pr(\text{y customers arrive in L})=\frac{(\lambda L)^{y}e^{-\lambda L}}{y!}$$

The probability of having a cumulative demand of k-j during the replenishment lead time L is now given by:

$$\Pr(D(L)=k-j)=\sum_{y}\Pr(\text{y customers arrive in L})\Pr(\text{y cust. order k-j spare parts})$$

Independent of the customer order quantity distribution the range of y is determined. We need at least one customer to have a positive value of the demand during replenishment lead time, so the minimum value of y is one. Since each customer orders at least one spare part, k-j spare parts cannot be ordered by more than k-j customers, so the maximum value of y equals k-j. This results in the following expression:

$$\Pr\left(D\left(L\right)=k-j\right)=\sum_{y=0}^{k-j}\frac{\left(\lambda L\right)^{y}e^{-\lambda L}}{y!}f^{y}\left(k-j\right)$$

To avoid calculating large factorials in the denominator as y increases, we use the following recursive relation on the probability of having a certain number of customers:

$$\Pr\left(\text{y customers arrive in L}\right)=\frac{\lambda L}{y}\Pr\left(\text{y-1 customers arrive in L}\right)$$

For large values of y, the probability that y customers arrive during the replenishment lead time is very small. We therefore define a maximum M on the number of customers we observe, so:

$$\Pr\left(D\left(L\right)=k-j\right)=\sum_{y=0}^{M}\frac{\left(\lambda L\right)^{y}e^{-\lambda L}}{y!}f^{y}\left(k-j\right)$$

The value of M is determined such that:

$$M=\min\left\{M\mid\Pr\left(\text{at most M customers arrive in L}\right)\geq 0.999999\right\}$$

$$M=\min\left\{M\mid\sum_{y=0}^{M}\frac{\left(\lambda L\right)^{y}e^{-\lambda L}}{y!}\geq 0.999999\right\}$$

Using the expression for the probability distributions of the inventory level and the demand during replenishment lead time we fill in the expression for the order line fill rate. It is given by:

$$OLFR\left(R,Q\right)=\sum_{j=1}^{R+Q}\Pr\left(IL=j\right)\Pr\left(F\leq j\right)$$

$$OLFR\left(R,Q\right)=\sum_{j=1}^{R+Q}\frac{1}{Q}\sum_{k=\max\left(R+1,j\right)}^{R+Q}\left[\Pr\left(D\left(L\right)=k-j\right)\right]\Pr\left(F\leq j\right)$$

$$OLFR\left(R,Q\right)=\frac{1}{Q}\sum_{j=1}^{R+Q}\Pr\left(F\leq j\right)\sum_{k=\max\left(R+1,j\right)}^{R+Q}\left[\sum_{y=0}^{M}\frac{\left(\lambda L\right)^{y}e^{-\lambda L}}{y!}f^{y}\left(k-j\right)\right]$$

We first calculate the probability distribution of D(L) for all possible values of the demand during replenishment lead time. In case the maximal value of F is 1,000 and M is 100 this takes around 10,000,000 operations. Calculating the order line fill rate then takes an additional (R+Q)*Q operations. For large values of λ and L the values of R and Q can increase up to over 10,000 and 1,000 respectively, implying a total number of operations for the order line fill rate calculation of more than 10,000,000.

B.1.2. Two moments method

The second method to approach the probability distribution of the demand during replenishment lead time makes use of the mean and variance of both the customer order quantity and the replenishment lead time. From these values we derive the mean and variance of the demand during lead time. These values are used to determine the parameters of a discrete probability distribution, the negative binomial distribution.

Throughout this section, the following notations are used:

- F: Stochastic variable that denotes the probability distribution of the customer order quantity. Its mean and variance can be calculated from the known discrete probability distribution.
- Λ: Stochastic variable of the daily arrival rate of the customer. This is a Poisson process with mean and variance equal to λ.
- L: Stochastic variable that denotes the probability distribution of the lead time. Its mean is EL, its variance varL. In the basic model, varL=0 since we have a fixed replenishment lead time. Since we know the probability distribution of the lead time we can calculate the mean and variance.
- N: Stochastic variable that denotes the number of customers during replenishment lead time. Its mean is EN, its variance varN.
- D: Stochastic variable that denotes the total demand during replenishment lead time. Its mean is ED, its variance varD.

We are interested in the mean and variance of D, the demand during replenishment lead time. This is derived in two steps. First we define the mean and variance of N, the number of customers during replenishment lead time, then we use this mean and variance to calculate ED and varD.

The number of customers during replenishment lead time is equal to the number of customers per day times the length of the replenishment lead time. Since these are both stochastic variables, the mean and variance of the product are given by:

$$\mathrm{E}N = \mathrm{E}L\mathrm{E}\Lambda$$

$$\operatorname{var} N = \mathrm{E}L \operatorname{var} \Lambda + (\mathrm{E}\Lambda)^2 \operatorname{var} L$$

$$\operatorname{var} N = \lambda \mathrm{E}L + \lambda^2 \operatorname{var} L$$

In the same way we construct the mean and variance of the demand during replenishment lead time, combining the number of customers during replenishment lead time and the probability distribution of the customer order quantity.

$$\mathrm{E}D = \mathrm{E}N\mathrm{E}F = \lambda \mathrm{E}L\mathrm{E}F$$

$$\operatorname{var} D = \mathrm{E}N \operatorname{var} F + (\mathrm{E}F)^2 \operatorname{var} N$$

$$\operatorname{var} D = \lambda \mathrm{E}L \operatorname{var} F + (\mathrm{E}F)^2 \left(\lambda \mathrm{E}L + \lambda^2 \operatorname{var} L\right)$$

$$\operatorname{var} D = \lambda\left(\mathrm{E}L \operatorname{var} F + (\mathrm{E}F)^2 (\mathrm{E}L + \lambda \operatorname{var} L)\right)$$

We want to fit this mean and variance to a suitable discrete probability distribution. There are three options, depending on the variance-to-mean ratio of the demand during replenishment lead time (i.e. varD/ED):

Binomial distribution

The probability mass function of the binomial distribution is given by:

$$\Pr\left(D(L)=k-j\right)=\left(\frac{t(t-1)(t-2)\ldots(t-(k-j)+1)}{(k-j)!}\right)p^{k-j}(1-p)^{t-(k-j)} \tag{9.1}$$

The input parameters are t and p, restricted by t>0 and 0<p<1. The mean of the binomial distribution is tp, its variance is tp(1-p). The values of t and p are found by solving the following two equations:

$$\begin{aligned} tp &= E\left(D(L)\right) \\ tp(1-p) &= \operatorname{var}\left(D(L)\right) \end{aligned} \tag{9.2}$$

Poisson distribution

The probability mass function of the Poisson distribution is given by:

$$\Pr\left(D(L)=k-j\right)=\frac{e^{-\lambda}\lambda^{k-j}}{(k-j)!} \tag{9.3}$$

The single input parameter is restricted by λ>0. The mean and variance of the Poisson distribution are both equal to λ. This value is found by:

$$\lambda = E\left(D(L)\right)\left(=\operatorname{var}\left(D(L)\right)\right) \tag{9.4}$$

In case the mean and variance of D(L) are not exactly equal to each other, we choose to use E(D(L)) to calculate the input parameter λ.

Negative binomial distribution

The probability mass function of the negative binomial distribution is given by:

$$\Pr\left(D(L)=k-j\right)=\frac{(r+k-j-1)(r+k-j-2)\ldots(r+1)r}{(k-j)!}p^{r}(1-p)^{k-j} \tag{9.5}$$

The two input parameters r and p are restricted by r>0 and 0<p<1. The mean of the negative binomial distribution is $r(1-p)/p$, its variance is $r(1-p)/p^2$. The values of the parameters r and p are found by solving:

$$\begin{aligned} \frac{r(1-p)}{p} &= E\left(D(L)\right) \\ \frac{r(1-p)}{p^2} &= \operatorname{var}\left(D(L)\right) \end{aligned} \tag{9.6}$$

The resulting discrete probability distribution of the demand during replenishment lead time is entered into the general order line fill rate calculation:

$$OLFR(R,Q)=\sum_{j=1}^{R+Q}\frac{1}{Q}\sum_{k=\max(R+1,j)}^{R+Q}\left[\Pr\left(D(L)=k-j\right)\right]\Pr\left(F\leq j\right)$$

Calculating the probability distribution of the demand during replenishment lead time now takes less than 1,000,000 operations instead of 10,000,000 in the compound Poisson method in the previous subsection.

B.2. Determining the reorder point R if the order quantity Q is given

The determination of the reorder point R and the order quantity Q is based on minimizing the total costs. In subsection 4.3.2 we have defined that the total costs include the holding costs and the order costs. We define them as HC(R,Q) and OC(Q) respectively, because the holding costs depend on both R and Q and the order costs only depend on Q. The total costs are:

$$C(R,Q)=HC(R,Q)+OC(Q)$$

$$C(R,Q)=\sum_{j=1}^{R+Q} jh\Pr(IL=j)+\frac{C\mu}{Q}$$

$$C(R,Q)=\sum_{j=1}^{R+Q}\frac{jh}{Q}\sum_{k=\max\{R+1.j\}}^{R+Q}\Pr\left(D(L)=k-j\right)+\frac{C\mu}{Q}$$

From the equations above follows that for a fixed value of Q, the total costs increase as R increases.

In our sequential method we first determine the replenishment order quantity Q before we determine the reorder point R. The objective is therefore to find the smallest reorder point R that achieves the target order line fill rate TFR. In mathematical notation, we are looking for:

$$R=\min\left\{R^{*}\mid OLFR\left(R^{*},Q\right)\geq TFR\right\}$$

$$R=\min\left\{R^{*}\mid\frac{1}{Q}\sum_{j=1}^{R^{*}+Q}\Pr\left(F\leq j\right)\left[\sum_{k=\max(R^{*}+1,j)}^{R^{*}+Q}\sum_{y=0}^{(k-j)}\frac{(\lambda L)^{y}e^{-\lambda L}}{y!}f^{y}(k-j)\right]\geq TFR\right\}$$

An easy way to find this reorder point is to start with its minimum value and increase by one until the target fill rate is obtained. This is suggested by (Axsäter, 2006). For R≤-Q, the inventory position (and hence the inventory level) is never positive since their maximum value is R+Q, so the order line fill rate is zero. We therefore start at R=-Q - for which OLFR(R,Q)=OLFR(-Q,Q)=0 – and increase R by one until the target order line fill rate is obtained. We have derived a recursive relation to calculate the order line fill rate:

$$OLFR(R,Q) = OLFR(R-1,Q) + \frac{1}{Q}\left[\sum_{j=1}^{R+Q} \Pr(F \le j)\Pr(D = R+Q-j) - \sum_{j=1}^{R} \Pr(F \le j)\Pr(D = R-j)\right] \qquad R > -Q$$

This calculation method works for all possible probability distributions of the demand during replenishment lead time.

B.3. Lagrangian relaxation

One of the methods to calculate the reorder point R and the replenishment order quantity Q in a simultaneous method is to minimize the total costs such that the target order line fill rate is achieved in a combinatorial optimization problem. This problem is formulated as:

$$\begin{aligned}
&\min C(R,Q) \\
&\text{s.t. } OLFR(R,Q) \ge TFR \\
&Q \in \mathbb{N}, R \in \{-Q, -Q+1, -Q+2, \ldots\}
\end{aligned}$$

The total costs consist of holding costs and replenishment order costs. The order line fill rate calculation is derived in Appendix B.1.1, the expression for the costs is given in Appendix B.2. Inserting these expressions for the costs and the order line fill rate we get:

$$\begin{aligned}
&\min \sum_{j=1}^{R+Q} jh \Pr(IL = j) + \frac{C\mu}{Q} \\
&\text{s.t. } \sum_{j=1}^{R+Q} \Pr(IL = j)\Pr(F \le j) \ge TFR \\
&Q \in \mathbb{N}, R \in \{-Q, -Q+1, \ldots\}
\end{aligned}$$

Due to the integrality constraints and R and Q and the large number of possible values for R and Q it is very hard to determine the optimal solution of this problem. We apply Lagrangian relaxation to simplify the solving of the problem. In Lagrangian relaxation we add the constraint to the objective function with a multiplier μ:

$$\begin{aligned}
&\min \sum_{j=1}^{R+Q} jh \Pr(F \le j) + \frac{C\mu}{Q} - \mu\left(\sum_{j=1}^{R+Q} \Pr(IL = j)\Pr(F \le j) - TFR\right) \\
&\text{s.t. } Q \in \mathbb{N}, R \in \{-Q, -Q+1, \ldots\}, \mu \ge 0
\end{aligned}$$

This simplifies to:

$$\begin{aligned}
&\min \sum_{j=1}^{R+Q} \left(jh - \mu \Pr(F \le j)\right)\Pr(IL = j) + \mu TFR + \frac{C\mu}{Q} \\
&\text{s.t. } Q \in \mathbb{N}, R \in \{-Q, -Q+1, \ldots\}, \mu \ge 0
\end{aligned}$$

(Fisher, 1981) discusses a methods to solve this problem. The solution of this problem is not necessarily the optimal solution of the original problem but (Fisher, 1981) indicates that the Lagrangian relaxation yields a good approach.

B.4. Periodic review model

We observe a time at the end of a period and denote this time with t. Any necessary replenishment order has just been placed and is hence included in the inventory position. The replenishment orders arrives at the time t+L. From The inventory level at this point equals:

$$IL(t+L) = IP(t) - D(t,t+L)$$

As for the basic model, the inventory position at time t is uniformly distributed between R+1 and R+Q. The demand during the period (t,t+L) is independent of the time t. We hence calculate the probability of having a certain inventory level x as:

$$\Pr\left(IL(t+L)=x\right) = \sum_{k=\max\{x,R+1\}}^{R+Q} \Pr\left(IP(t)=k\right)\Pr\left(D(L)=k-x\right)$$

$$\Pr\left(IL(t+L)=x\right) = \frac{1}{Q}\sum_{k=\max\{x,R+1\}}^{R+Q} \Pr\left(D(L)=k-x\right)$$

This inventory level 'has to last' until the next possible point of replenishment t+L+T.

We deduce the order line fill rate by looking at the number of fills for a fixed number of customers during this period, and sum over all possible numbers of customers. So:

$$OLFR = \sum_{y=0}^{\infty} \Pr\left(\text{y customers}\right)\sum_{x=0}^{y}\frac{x}{y}\Pr\left(\text{x out of y orders are filled}\right)$$

The probability that x out of y orders are filled also depends on the initial inventory level IL(t+L), so this has to be included in the equation as well.

The probability that x out of y orders are filled is calculated as follows. We assume that we know how much customers arrive during the period (t+L,t+L+T), so we know the number y. If x orders are filled, this means that the first x customers order a total of not more than the inventory level at time t+L. The (x+1)th customer orders more than the remaining inventory level, otherwise this order could be filled as well. At a given inventory level z, the probability of filling x out of y orders is hence:

$$\Pr\left(\text{x out of y orders filled}\right) = \sum_{j=0}^{z}\Pr\left(\text{x customers order j}\right)\Pr\left((x+1)^{th}\text{ customer orders more than z-j}\right)$$

$$\Pr\left(\text{x out of y orders filled}\right) = \sum_{j=0}^{z} f^{x*}(j)\Pr\left(F > z-j\right)$$

The corresponding order line fill rate is x/y. The calculation of the total order line fill rate hence becomes:

1) Set the inventory level z to 0.
2) Set the number of customers, denoted with y, equal to 1.
3) Calculate for y customers the probability that x out of y orders are filled for x from 0 to y-1 by the formula as given above. Multiply this probability with the resulting order line fill rate x/y.
4) Calculate the probability that all y customers are filled. This is equal to the probability that y customers order a number of spare parts at most equal to the inventory level (say, x):

$$\sum_{i=0}^{z} f^{y}(i)$$

5) Add the results of steps 2) en 3) to each other and multiply this with the probability that during the review period of length T, y customers arrive. We denote this with:

$$\Pr(C(T)=y)=\frac{(\lambda T)^{y} e^{-\lambda T}}{y!}$$

Add the result to a temporary variable.
6) If y is large enough (such that Pr(C(T)=y) is very small, i.e. <0.000001) go to step 7, else increase y with 1 and go to step 3.
7) Add the probability that no customers arrive during the review period to the temporary variable.
8) Multiply the temporary variable with the probability that the inventory level equals z and add this to the order line fill rate. If z equals R+Q, stop, else increase z by 1 and return to step 2.

The resulting expression for the order line fill rate is:

$$OLFR=\sum_{z=0}^{R+Q}\Pr\left(IL(t+L)=z\right)\left\{\sum_{y=1}^{\infty}\Pr\left(C(T)=y\right)*\ldots\right.$$

$$\left.\ldots\left[\sum_{x=0}^{y-1}\frac{x}{y}\sum_{j=0}^{z}f^{x}(j)\Pr(F>z-j)+\sum_{i=0}^{z}f^{y}(i)\right]+\Pr\left(C(T)=0\right)\right\}$$

C. Model verifications

In this Appendix detailed information is provided on probability distributions that are used in our verification studies.

C.1. Customer order quantity distributions

Option 1

x	Pr(F=x)	Pr(F≤x)
100	1	1

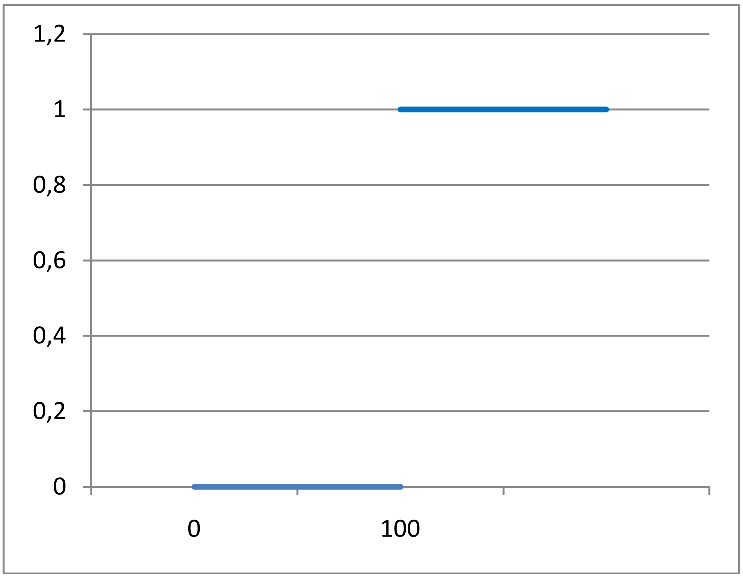

Option 2

x	Pr(F=x)	Pr(F≤x)
1	0.307937	0.307937
2	0.184127	0.492063
3	0.120635	0.612698
4	0.07619	0.688889
5	0.079365	0.768254
6	0.057143	0.825397
7	0.047619	0.873016
8	0.047619	0.920635
10	0.031746	0.952381
12	0.019048	0.971429
15	0.015873	0.987302
18	0.003175	0.990476
20	0.009524	1

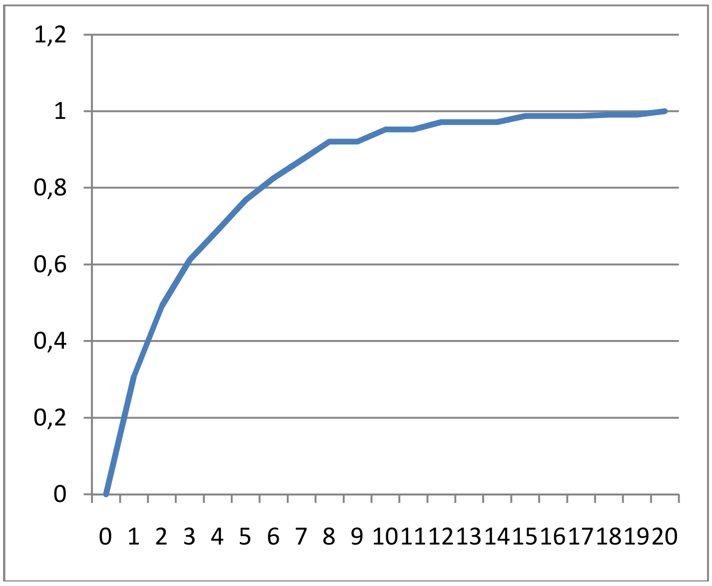

Option 3

x	Pr(F=x)	Pr(F≤x)
1	0.061135	0.061135
2	0.085153	0.146288
3	0.133188	0.279476
4	0.141921	0.421397
5	0.181223	0.60262
6	0.091703	0.694323
8	0.080786	0.775109
10	0.08952	0.864629
12	0.048035	0.912664
14	0.037118	0.949782
15	0.032751	0.982533
16	0.00655	0.989083
18	0.00655	0.995633
20	0.004367	1

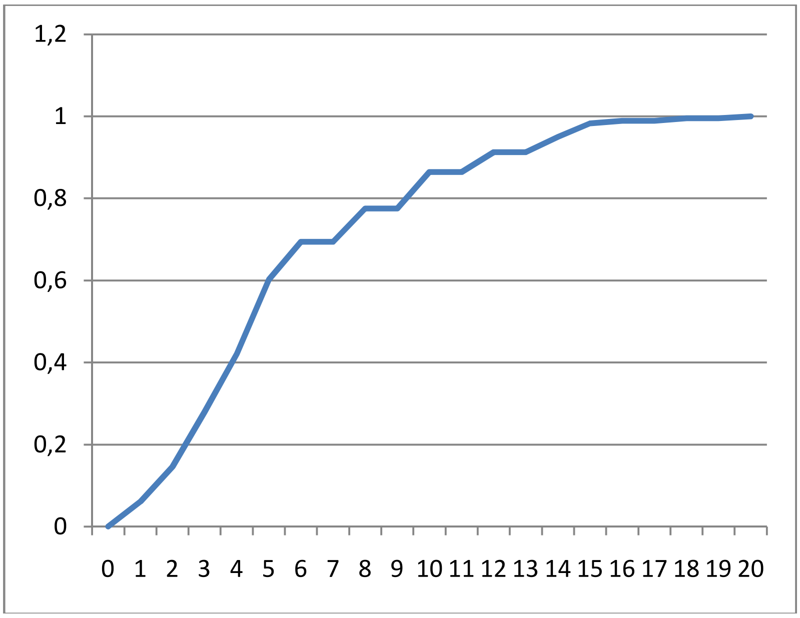

Option 4

x	Pr(F=x)	Pr(F≤x)
2	0.02849	0.02849
4	0.019943	0.048433
5	0.025641	0.0740741
6	0.02849	0.1025641
7	0.045584	0.1481481
8	0.051282	0.1994302
9	0.102564	0.3019943
10	0.222222	0.5242165
11	0.145299	0.6695157
12	0.116809	0.7863248
13	0.074074	0.8603989
14	0.037037	0.8974359
15	0.034188	0.9316239
16	0.017094	0.9487179
17	0.017094	0.965812
18	0.014245	0.980057
19	0.008547	0.988604
20	0.011396	1

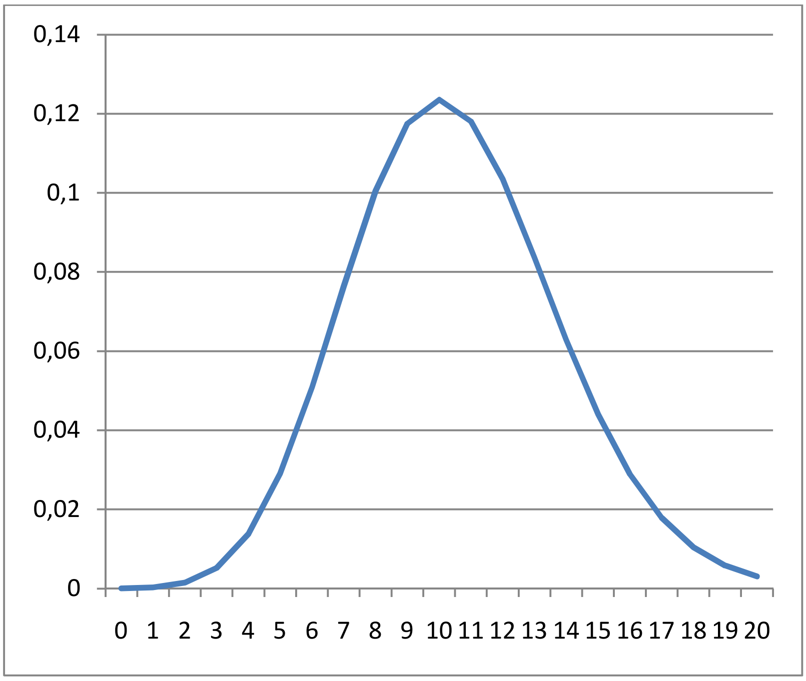

Option 5

x	Pr(F=x)	Pr(F≤x)
1	0.129237	0.129237
2	0.106992	0.236229
3	0.09428	0.330508
4	0.085805	0.416314
5	0.077331	0.493644
6	0.064619	0.558263
7	0.060381	0.618644
8	0.061441	0.680085
9	0.043432	0.723517
10	0.060381	0.783898
12	0.037076	0.820975
15	0.044492	0.865466
20	0.043432	0.908898
21	0.003178	0.912076
25	0.026483	0.938559
30	0.01589	0.954449
40	0.010593	0.965042
50	0.013771	0.978814
60	0.005297	0.98411
70	0.005297	0.989407
75	0.003178	0.992585
80	0.004237	0.996822
90	0.002119	0.998941
100	0.001059	1

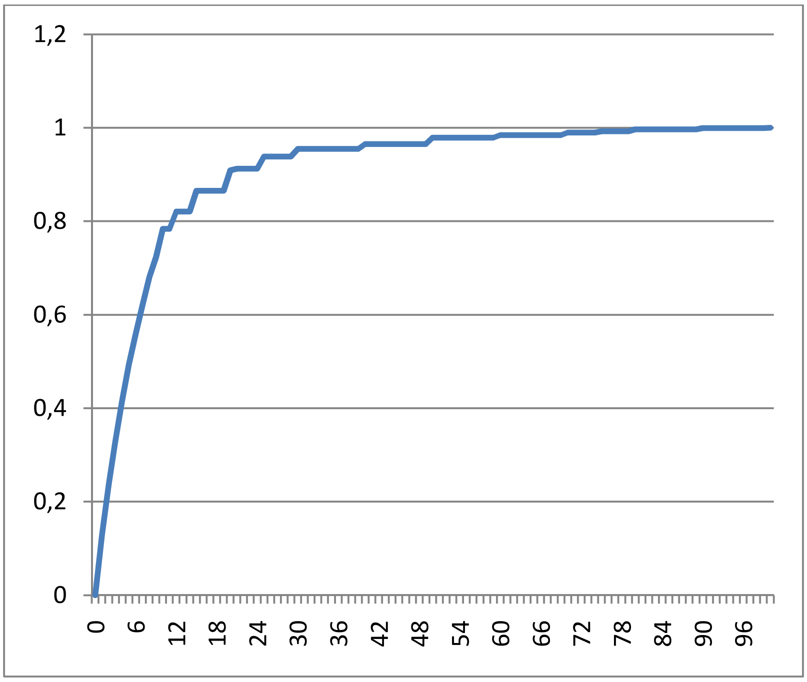

Option 6

x	Pr(F=x)	Pr(F≤x)
1	0.048232	0.048232
2	0.065916	0.114148
3	0.045016	0.159164
5	0.078778	0.237942
6	0.053055	0.290997
10	0.181672	0.472669
15	0.160772	0.633441
16	0.057878	0.691318
20	0.075563	0.766881
25	0.072347	0.839228
50	0.049839	0.889068
60	0.017685	0.906752
75	0.043408	0.950161
80	0.016077	0.966238
100	0.033762	1

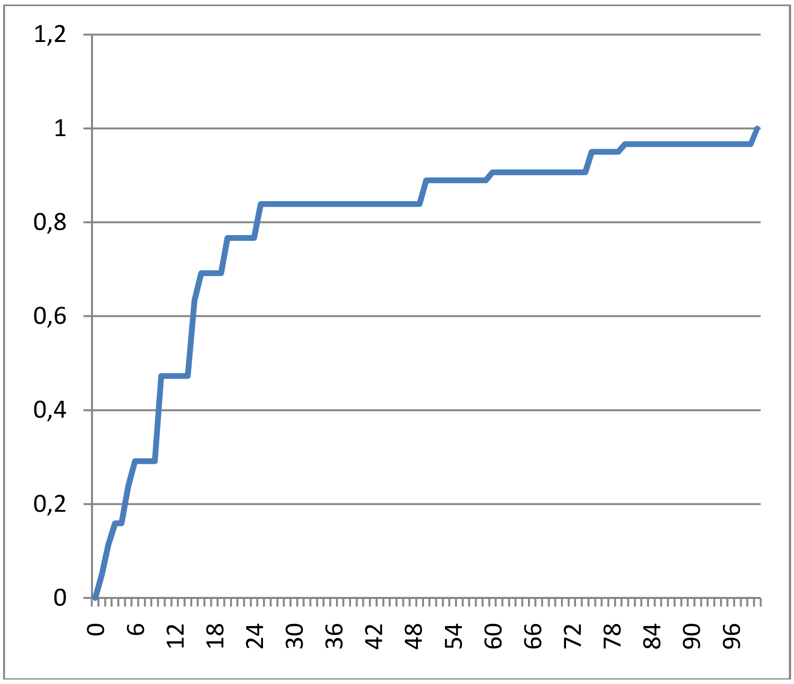

Option 7

x	Pr(F=x)	Pr(F≤x)
1	0.00656	0.00656
5	0.012184	0.018744
10	0.010309	0.029053
25	0.044049	0.073102
30	0.04686	0.119963
50	0.081537	0.2015
60	0.036551	0.238051
75	0.085286	0.323336
85	0.025305	0.348641
90	0.038425	0.387067
100	0.126523	0.51359
125	0.073102	0.586692
150	0.076851	0.663543
200	0.113402	0.776945
240	0.044986	0.821931
250	0.055295	0.877226
275	0.019681	0.896907
300	0.036551	0.933458
350	0.014058	0.947516
400	0.029991	0.977507
480	0.008435	0.985942
500	0.014058	1

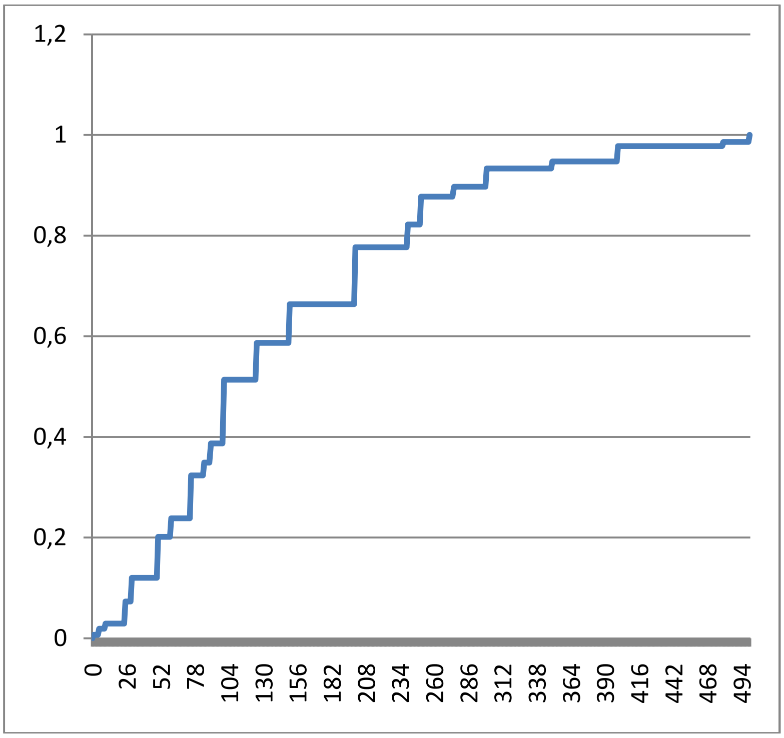

Option 8

x	Pr(F=x)	Pr(F≤x)
10	0.263889	0.263889
20	0.154762	0.418651
50	0.10119	0.519841
80	0.097222	0.617063
100	0.125	0.742063
200	0.095238	0.837302
250	0.063492	0.900794
300	0.049603	0.950397
500	0.03373	0.984127
1000	0.015873	1

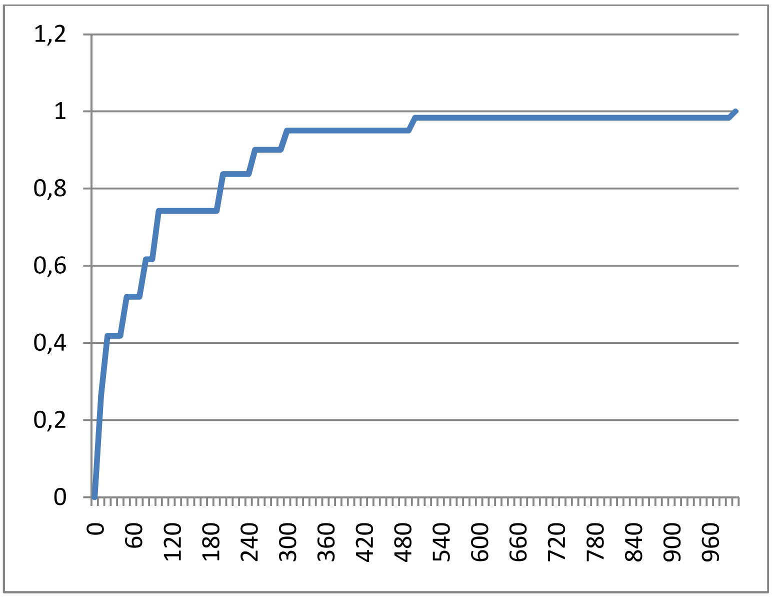

Option 9

x	Pr(F=x)	Pr(F≤x)
10	0.124108	0.124108
20	0.112696	0.236805
30	0.10699	0.343795
40	0.097004	0.440799
50	0.087019	0.527817
70	0.079886	0.607703
80	0.072753	0.680456
100	0.048502	0.728959
160	0.058488	0.787447
200	0.038516	0.825963
250	0.029957	0.85592
300	0.027104	0.883024
320	0.035663	0.918688
350	0.022825	0.941512
400	0.017118	0.958631
420	0.007133	0.965763
450	0.014265	0.980029
480	0.015692	0.99572
500	0.00428	1

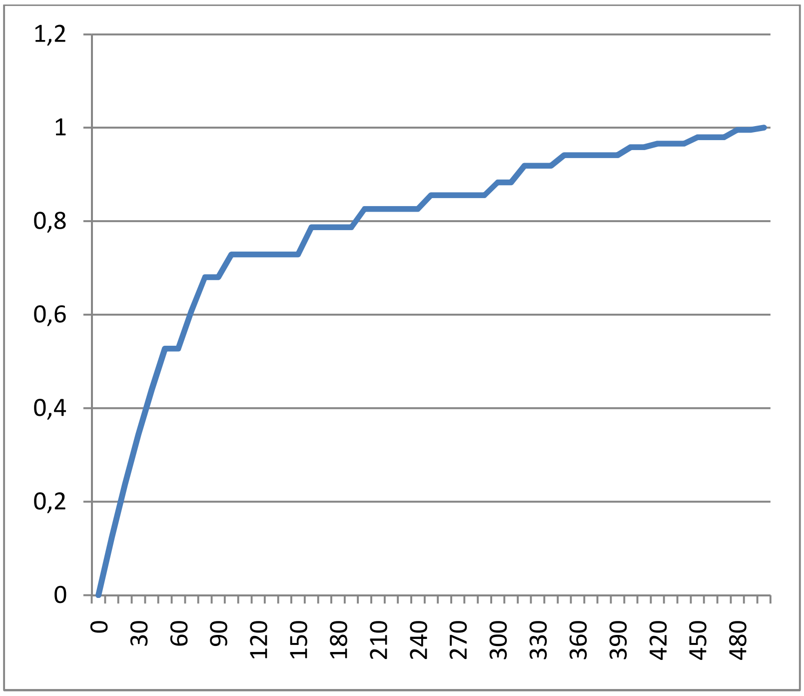

Option 10

x	Pr(F=x)	Pr(F≤x)
100	0.079452	0.079452
200	0.112329	0.191781
300	0.139726	0.331507
500	0.205479	0.536986
800	0.131507	0.668493
1000	0.109589	0.778082
1200	0.084932	0.863014
1500	0.068493	0.931507
1800	0.041096	0.972603
2000	0.027397	1

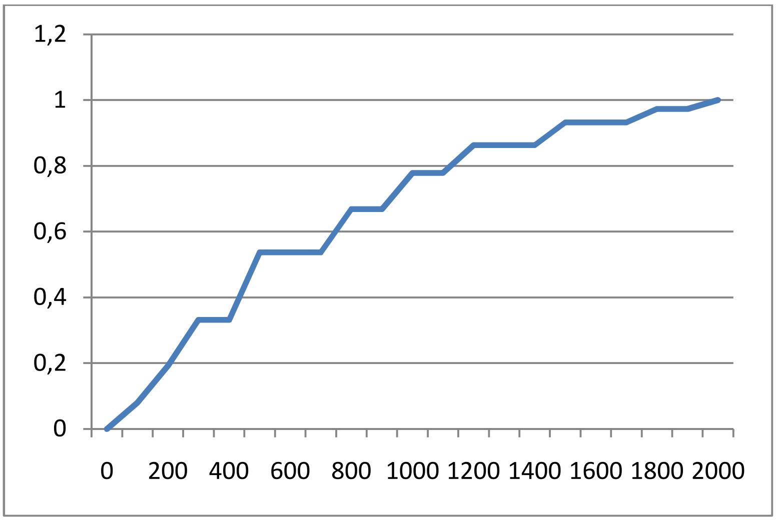

C.2. Supplier delay probability distributions

Per supplier we give the complete discrete delay probability distribution, plus the probability density plots of these distributions.

Supplier A

x	Pr(Delay=x)
0	0,941716136
1	0,016189962
2	0,009174312
3	0,005396654
4	0,006475985
5	0,004856989
6	0,002698327
7	0,001618996
8	0,001079331
9	0,001079331
10	0,000539665
11	0,000539665
12	0,001618996
15	0,001079331
16	0,002158662
20	0,000539665
21	0,000539665
35	0,002698327

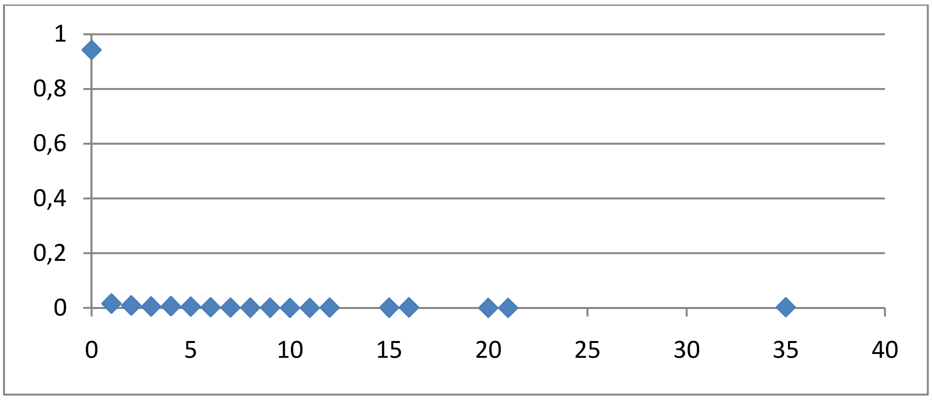

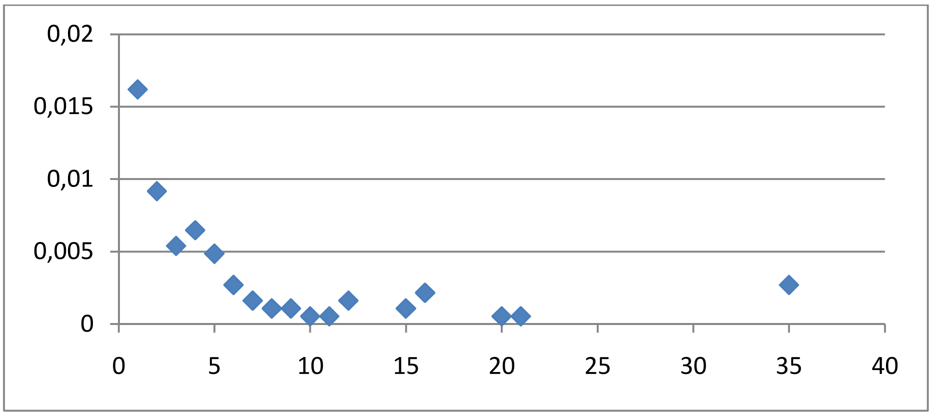

Supplier B

x	Pr(Delay=x)
0	0,976584022
1	0,015151515
2	0,004132231
3	0,00137741
8	0,00137741
13	0,00137741

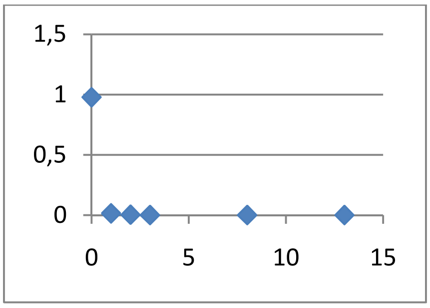

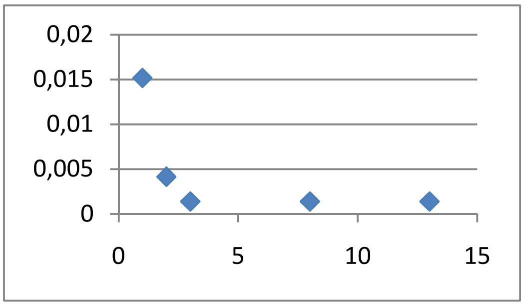

Supplier C

x	Pr(Delay=x)
0	0,046382189
1	0,141001855
2	0,14471243
3	0,161410019
4	0,085343228
5	0,077922078
6	0,090909091
7	0,109461967
8	0,037105751
9	0,009276438
10	0,024118738
11	0,003710575
12	0,027829314
14	0,001855288
15	0,003710575
16	0,003710575
17	0,009276438
19	0,001855288
20	0,001855288
22	0,003710575
23	0,001855288
25	0,001855288
35	0,011131725

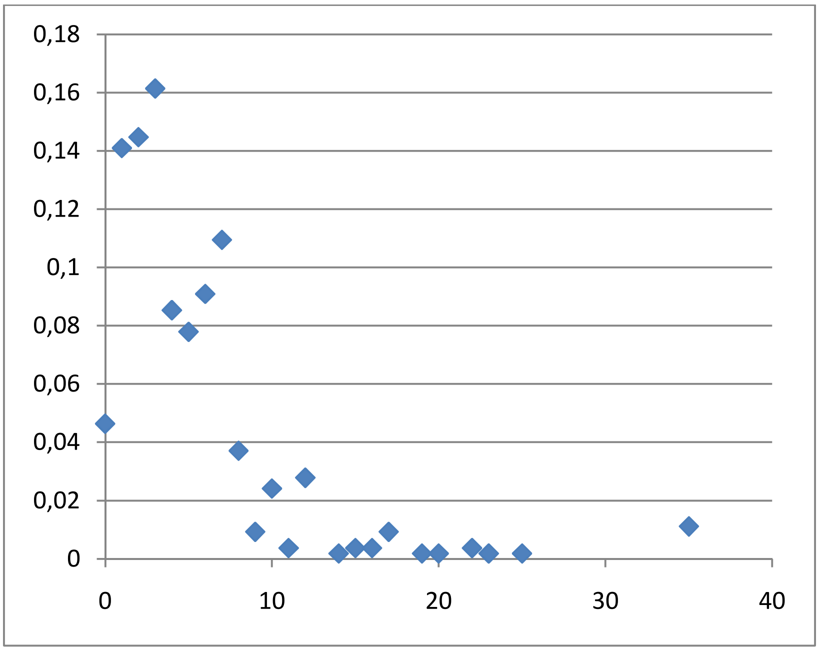

Supplier D

x	Pr(Delay=x)
0	0,883333333
1	0,026666667
3	0,01
5	0,023333333
6	0,003333333
7	0,01
10	0,006666667
16	0,003333333
20	0,003333333
21	0,003333333
24	0,003333333
29	0,003333333
35	0,02

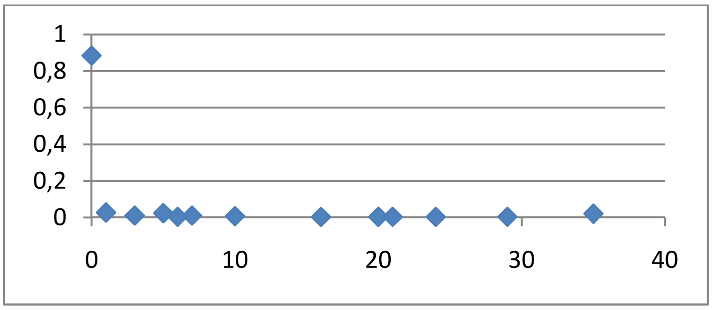

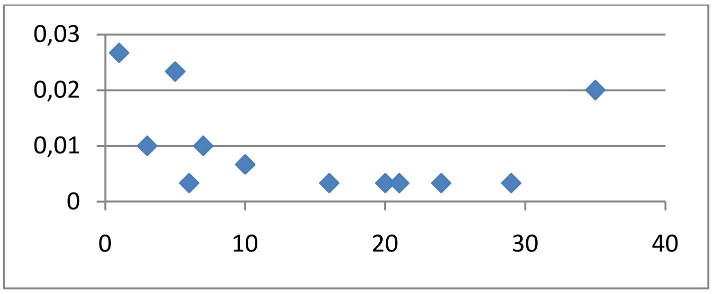

Supplier E

x	Pr(Delay=x)
0	0,469924812
1	0,063909774
2	0,063909774
3	0,045112782
4	0,071428571
5	0,041353383
6	0,037593985
7	0,026315789
8	0,033834586
9	0,003759398
10	0,015037594
12	0,003759398
13	0,015037594
14	0,007518797
15	0,011278195
20	0,003759398
21	0,007518797
22	0,026315789
23	0,018796992
25	0,007518797
28	0,007518797
29	0,003759398
35	0,015037594

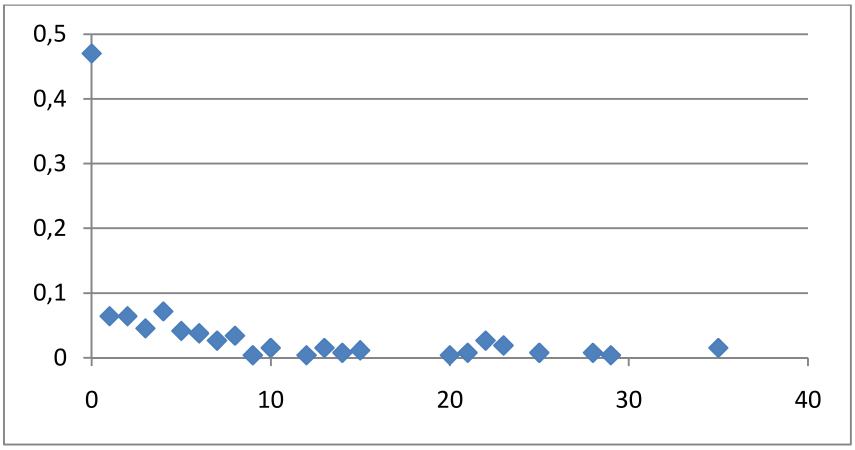

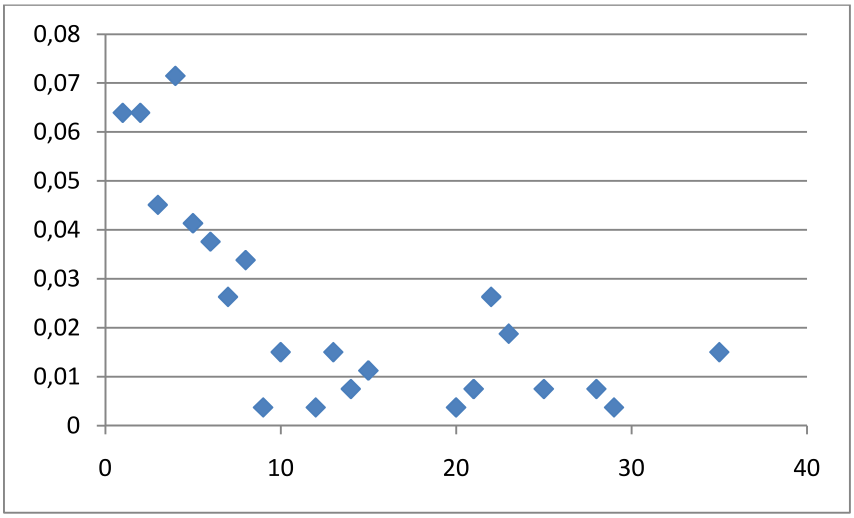

Supplier F

x	Pr(Delay=x)
0	0,94
1	0,04
2	0,01
5	0,01

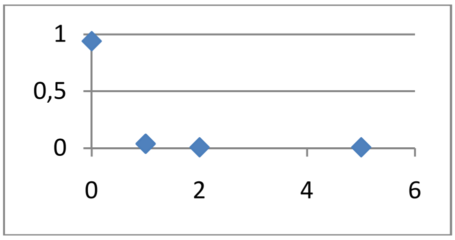

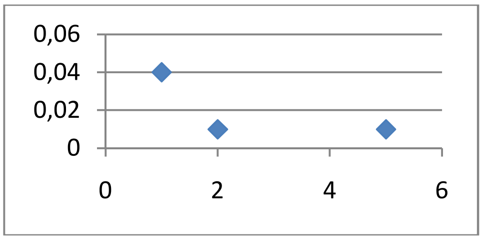

Supplier G

x	Pr(Delay=x)
0	0,929292929
1	0,04040404
2	0,01010101
3	0,01010101
12	0,01010101

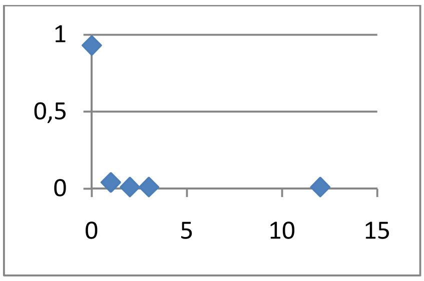

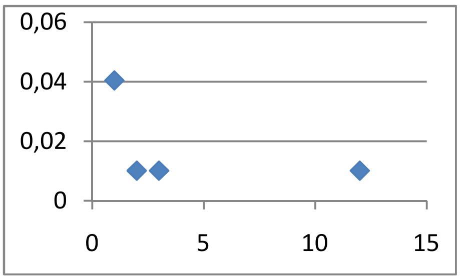

D. Model verification results

D.1. Basic model

Table 23 and Table 24 show the complete results for the compound Poisson method.

	Average difference	Std.dev difference	Average absolute difference	Std.dev absolute difference
Total:	-0,009%	0,082%	0,054%	0,063%
	Customer order quantity distributions			
1	0,003%	0,049%	0,033%	0,036%
2	-0,005%	0,084%	0,054%	0,065%
3	-0,015%	0,083%	0,051%	0,066%
4	-0,014%	0,083%	0,055%	0,064%
5	-0,004%	0,073%	0,052%	0,052%
6	0,000%	0,081%	0,057%	0,058%
7	-0,033%	0,113%	0,074%	0,091%
8	0,000%	0,070%	0,050%	0,049%
9	-0,005%	0,084%	0,058%	0,061%
10	-0,013%	0,084%	0,055%	0,065%
	Cost price			
1000	-0,016%	0,099%	0,068%	0,073%
5	-0,001%	0,061%	0,039%	0,046%
	Replenishment lead time			
65	-0,020%	0,121%	0,083%	0,090%
28	-0,012%	0,076%	0,053%	0,056%
15	0,002%	0,060%	0,042%	0,042%
5	-0,004%	0,053%	0,037%	0,038%
	Demand rate			
0.85	-0,034%	0,096%	0,063%	0,080%
0.75	-0,026%	0,096%	0,056%	0,082%
0.35	-0,005%	0,075%	0,049%	0,056%
0.25	-0,001%	0,079%	0,051%	0,060%
0.077	0,005%	0,077%	0,055%	0,055%
0.027	0,007%	0,070%	0,052%	0,047%
0.004	-0,006%	0,070%	0,051%	0,048%

Table 23 Means and standard deviations for basic model – compound Poisson method

	50%	75%	90%	95%	99%	Max
Total:	0,033%	0,071%	0,127%	0,175%	0,296%	0,634%
	Customer order quantity distributions					
1	0,020%	0,047%	0,071%	0,096%	0,175%	0,225%
2	0,032%	0,073%	0,126%	0,191%	0,291%	0,424%
3	0,030%	0,067%	0,111%	0,156%	0,385%	0,458%
4	0,032%	0,073%	0,137%	0,162%	0,280%	0,429%
5	0,035%	0,068%	0,114%	0,147%	0,244%	0,266%
6	0,040%	0,083%	0,135%	0,155%	0,259%	0,376%
7	0,046%	0,090%	0,190%	0,249%	0,417%	0,634%
8	0,036%	0,067%	0,105%	0,131%	0,259%	0,350%
9	0,034%	0,082%	0,143%	0,173%	0,251%	0,364%

10	0,030%	0,070%	0,137%	0,189%	0,303%	0,360%
	Cost price					
1000	0,045%	0,090%	0,156%	0,210%	0,356%	0,634%
5	0,024%	0,054%	0,089%	0,123%	0,215%	0,431%
	Replenishment lead time					
65	0,053%	0,106%	0,202%	0,275%	0,422%	0,634%
28	0,033%	0,072%	0,134%	0,164%	0,251%	0,364%
15	0,031%	0,058%	0,096%	0,130%	0,194%	0,248%
5	0,024%	0,053%	0,086%	0,112%	0,174%	0,215%
	Demand rate					
0.85	0,033%	0,084%	0,144%	0,253%	0,396%	0,431%
0.75	0,029%	0,063%	0,132%	0,210%	0,407%	0,634%
0.35	0,031%	0,069%	0,109%	0,171%	0,264%	0,350%
0.25	0,028%	0,066%	0,130%	0,190%	0,258%	0,412%
0.077	0,036%	0,077%	0,127%	0,173%	0,237%	0,301%
0.027	0,040%	0,069%	0,115%	0,159%	0,204%	0,249%
0.004	0,037%	0,075%	0,112%	0,139%	0,224%	0,279%

Table 24 Percentiles of absolute difference for basic model - compound Poisson method

Table 25 and Table 26 show the results of the two moments method for the basic model.

	Average difference	Std.dev difference	Average absolute difference	Std.dev absolute difference
Total:	0,035%	0,392%	0,237%	0,315%
	Customer order quantity distributions			
1	-0,006%	0,056%	0,037%	0,043%
2	0,013%	0,187%	0,144%	0,119%
3	0,012%	0,347%	0,259%	0,231%
4	0,150%	0,863%	0,521%	0,703%
5	0,044%	0,194%	0,138%	0,143%
6	0,069%	0,367%	0,275%	0,252%
7	-0,019%	0,408%	0,322%	0,250%
8	0,024%	0,170%	0,123%	0,120%
9	0,046%	0,332%	0,261%	0,209%
10	0,017%	0,381%	0,290%	0,247%
	Cost price			
1000	0,013%	0,441%	0,280%	0,342%
5	0,057%	0,335%	0,194%	0,279%
	Replenishment lead time			
65	-0,047%	0,422%	0,255%	0,340%
28	0,016%	0,397%	0,244%	0,313%
15	0,057%	0,365%	0,233%	0,287%
5	0,113%	0,366%	0,216%	0,317%
	Demand rate			
0.85	-0,096%	0,173%	0,157%	0,121%
0.75	-0,084%	0,182%	0,157%	0,125%
0.35	-0,051%	0,249%	0,198%	0,159%
0.25	-0,022%	0,278%	0,206%	0,187%
0.077	0,087%	0,425%	0,307%	0,305%
0.027	0,205%	0,547%	0,379%	0,445%

0.004	0,206%	0,544%	0,256%	0,523%

Table 25 Means and standard deviations for basic model - two moments method

	50%	75%	90%	95%	99%	Max
Total:	0,160%	0,312%	0,485%	0,706%	1,299%	4,395%
	Customer order quantity distributions					
1	0,023%	0,045%	0,085%	0,127%	0,194%	0,275%
2	0,107%	0,205%	0,337%	0,383%	0,471%	0,496%
3	0,205%	0,352%	0,556%	0,707%	1,035%	1,305%
4	0,294%	0,520%	1,129%	2,303%	3,435%	4,395%
5	0,083%	0,194%	0,332%	0,435%	0,640%	0,745%
6	0,219%	0,337%	0,515%	0,835%	1,246%	1,527%
7	0,280%	0,383%	0,611%	0,938%	1,139%	1,297%
8	0,086%	0,154%	0,249%	0,393%	0,563%	0,673%
9	0,236%	0,339%	0,433%	0,765%	0,960%	0,993%
10	0,238%	0,384%	0,535%	0,747%	1,150%	1,391%
	Cost price					
1000	0,215%	0,356%	0,528%	0,757%	1,489%	4,395%
5	0,118%	0,245%	0,395%	0,588%	1,135%	4,035%
	Replenishment lead time					
65	0,190%	0,344%	0,514%	0,653%	0,949%	4,395%
28	0,174%	0,312%	0,444%	0,715%	1,669%	2,592%
15	0,160%	0,302%	0,475%	0,799%	1,253%	2,498%
5	0,131%	0,278%	0,466%	0,704%	1,286%	3,139%
	Demand rate					
0.85	0,128%	0,243%	0,312%	0,374%	0,482%	0,586%
0.75	0,132%	0,233%	0,323%	0,401%	0,502%	0,699%
0.35	0,166%	0,282%	0,386%	0,486%	0,702%	0,921%
0.25	0,162%	0,305%	0,389%	0,514%	0,956%	1,237%
0.077	0,249%	0,410%	0,686%	0,855%	1,285%	2,585%
0.027	0,274%	0,473%	0,854%	1,149%	2,424%	3,139%
0.004	0,088%	0,259%	0,524%	0,939%	2,578%	4,395%

Table 26 Percentiles of absolute difference for basic model – two moments method

D.2. Lead time variability

Table 27 and Table 28 show the results of the first modeling option to include lead time variability, by using the mean of the replenishment lead time in the compound Poisson method.

	Average difference	Std.dev difference	Average absolute difference	Std.dev absolute difference
Total:	0,822%	1,827%	0,846%	1,816%
	Customer order quantity distributions			
1	0,910%	2,135%	0,923%	2,129%
2	0,701%	1,637%	0,724%	1,628%
3	0,857%	1,910%	0,884%	1,898%
4	1,056%	2,168%	1,076%	2,158%
5	0,567%	1,289%	0,593%	1,277%
6	0,745%	1,577%	0,774%	1,563%
7	1,262%	2,382%	1,293%	2,366%

8	0,569%	1,272%	0,593%	1,260%
9	0,698%	1,540%	0,724%	1,528%
10	0,855%	1,893%	0,881%	1,881%
	Cost price			
1000	1,205%	2,307%	1,234%	2,291%
5	0,439%	1,031%	0,459%	1,023%
	Replenishment lead time			
65	0,429%	0,804%	0,463%	0,785%
28	0,672%	1,175%	0,693%	1,163%
15	0,844%	1,444%	0,859%	1,435%
5	1,049%	1,787%	1,063%	1,779%
	Demand rate			
0.85	1,621%	1,955%	1,638%	1,941%
0.75	1,522%	1,851%	1,535%	1,840%
0.35	0,996%	1,344%	1,005%	1,337%
0.25	0,780%	1,113%	0,793%	1,104%
0.077	0,257%	0,384%	0,273%	0,372%
0.027	0,069%	0,120%	0,095%	0,100%
0.004	-0,003%	0,066%	0,048%	0,046%
	Supplier			
A	0,269%	0,350%	0,285%	0,337%
B	0,011%	0,100%	0,061%	0,079%
C	1,124%	1,430%	1,133%	1,423%
D	1,591%	1,926%	1,598%	1,920%
E	2,668%	3,346%	2,674%	3,341%
F	0,008%	0,088%	0,061%	0,064%
G	0,083%	0,153%	0,112%	0,132%

Table 27 Means and standard deviations for variable lead time verification – compound Poisson method 1

	50%	75%	90%	95%	99%	Max
Total:	0,115%	0,657%	2,705%	4,778%	9,280%	16,342%
	Customer order quantity distributions					
1	0,108%	0,676%	2,741%	5,421%	11,387%	16,342%
2	0,099%	0,560%	2,045%	4,186%	8,339%	13,038%
3	0,120%	0,689%	2,691%	5,053%	9,836%	14,282%
4	0,150%	0,937%	3,628%	5,511%	10,854%	15,169%
5	0,093%	0,450%	1,703%	3,209%	6,372%	10,608%
6	0,113%	0,604%	2,619%	4,359%	7,351%	11,756%
7	0,179%	1,153%	4,676%	6,216%	11,079%	13,789%
8	0,099%	0,452%	1,758%	3,190%	6,086%	10,514%
9	0,105%	0,550%	2,236%	3,947%	7,513%	12,013%
10	0,132%	0,646%	2,792%	4,896%	9,413%	14,109%
	Cost price					
1000	0,162%	1,125%	4,447%	6,135%	11,056%	16,342%
5	0,084%	0,397%	1,327%	2,168%	4,926%	13,564%
	Replenishment lead time					
65	0,107%	0,409%	1,546%	2,607%	4,957%	8,442%
28	0,110%	0,628%	2,531%	4,158%	7,747%	11,973%
15	0,118%	0,776%	3,316%	5,089%	9,668%	13,996%
5	0,138%	0,952%	4,478%	6,314%	12,337%	16,342%
	Demand rate					

0.85	0,491%	2,492%	5,598%	7,532%	12,813%	16,342%
0.75	0,480%	2,260%	5,267%	7,045%	11,908%	14,572%
0.35	0,298%	1,192%	3,610%	4,900%	8,438%	11,589%
0.25	0,236%	0,905%	2,706%	3,982%	6,612%	10,238%
0.077	0,107%	0,326%	0,810%	1,265%	2,325%	4,716%
0.027	0,063%	0,137%	0,247%	0,345%	0,629%	1,261%
0.004	0,033%	0,070%	0,110%	0,137%	0,199%	0,353%
	Supplier					
A	0,136%	0,404%	0,847%	1,041%	1,344%	1,823%
B	0,038%	0,083%	0,149%	0,186%	0,312%	1,949%
C	0,541%	1,612%	3,362%	4,397%	5,948%	7,622%
D	0,705%	2,538%	4,927%	5,659%	7,190%	8,485%
E	1,262%	3,757%	8,097%	10,401%	13,454%	16,342%
F	0,039%	0,083%	0,142%	0,188%	0,306%	0,451%
G	0,062%	0,138%	0,295%	0,411%	0,595%	0,776%

Table 28 Percentiles of absolute difference for variable lead time verification – compound Poisson method 1

Table 29 and Table 30 show the results of the second modeling option to include lead time variability, by using conditioning on the replenishment lead time in the compound Poisson method.

	Average difference	Std.dev difference	Average absolute difference	Std.dev absolute difference
Total:	-0,002%	0,070%	0,040%	0,057%
	Customer order quantity distributions			
1	-0,001%	0,043%	0,024%	0,036%
2	0,001%	0,070%	0,039%	0,058%
3	-0,007%	0,076%	0,041%	0,064%
4	-0,004%	0,069%	0,040%	0,057%
5	0,004%	0,065%	0,039%	0,052%
6	0,000%	0,071%	0,042%	0,057%
7	-0,007%	0,092%	0,057%	0,073%
8	0,004%	0,065%	0,038%	0,052%
9	-0,002%	0,071%	0,041%	0,058%
10	-0,005%	0,069%	0,039%	0,057%
	Cost price			
1000	-0,003%	0,086%	0,053%	0,067%
5	0,000%	0,050%	0,027%	0,042%
	Replenishment lead time			
65	-0,025%	0,104%	0,073%	0,079%
28	-0,003%	0,065%	0,044%	0,048%
15	0,012%	0,042%	0,026%	0,035%
5	0,010%	0,027%	0,015%	0,025%
	Demand rate			
0.85	-0,020%	0,080%	0,040%	0,072%
0.75	-0,014%	0,071%	0,038%	0,062%
0.35	0,002%	0,069%	0,041%	0,056%
0.25	0,006%	0,068%	0,041%	0,055%
0.077	0,007%	0,066%	0,043%	0,051%
0.027	0,008%	0,059%	0,041%	0,043%
0.004	0,001%	0,053%	0,033%	0,042%

	Supplier			
A	-0,001%	0,064%	0,037%	0,053%
B	-0,008%	0,068%	0,034%	0,059%
C	0,000%	0,071%	0,045%	0,055%
D	0,003%	0,068%	0,042%	0,053%
E	-0,005%	0,079%	0,050%	0,062%
F	-0,002%	0,072%	0,036%	0,062%
G	0,001%	0,066%	0,036%	0,055%

Table 29 Means and standard deviations for timeframe verification

	50%	75%	90%	95%	99%	Max
Total:	0,019%	0,051%	0,105%	0,151%	0,273%	0,836%
	Customer order quantity distributions					
1	0,011%	0,031%	0,066%	0,090%	0,152%	0,519%
2	0,018%	0,048%	0,101%	0,151%	0,264%	0,501%
3	0,018%	0,053%	0,110%	0,157%	0,300%	0,836%
4	0,019%	0,049%	0,102%	0,155%	0,268%	0,548%
5	0,020%	0,052%	0,101%	0,140%	0,243%	0,500%
6	0,023%	0,054%	0,110%	0,159%	0,267%	0,455%
7	0,031%	0,075%	0,139%	0,209%	0,344%	0,533%
8	0,018%	0,050%	0,107%	0,147%	0,227%	0,431%
9	0,020%	0,052%	0,110%	0,153%	0,251%	0,713%
10	0,017%	0,050%	0,106%	0,147%	0,265%	0,530%
	Cost price					
1000	0,029%	0,072%	0,135%	0,186%	0,323%	0,612%
5	0,012%	0,033%	0,071%	0,104%	0,181%	0,836%
	Replenishment lead time					
65	0,047%	0,100%	0,176%	0,248%	0,387%	0,836%
28	0,028%	0,060%	0,108%	0,141%	0,215%	0,713%
15	0,014%	0,035%	0,072%	0,099%	0,176%	0,298%
5	0,004%	0,015%	0,040%	0,070%	0,131%	0,229%
	Demand rate					
0.85	0,013%	0,041%	0,107%	0,172%	0,334%	0,836%
0.75	0,014%	0,043%	0,100%	0,153%	0,326%	0,470%
0.35	0,019%	0,051%	0,107%	0,155%	0,264%	0,519%
0.25	0,020%	0,056%	0,111%	0,159%	0,266%	0,500%
0.077	0,022%	0,061%	0,119%	0,163%	0,240%	0,440%
0.027	0,026%	0,058%	0,104%	0,131%	0,207%	0,360%
0.004	0,019%	0,047%	0,093%	0,124%	0,191%	0,329%
	Supplier					
A	0,016%	0,046%	0,101%	0,142%	0,247%	0,548%
B	0,009%	0,040%	0,096%	0,142%	0,285%	0,713%
C	0,026%	0,058%	0,109%	0,151%	0,269%	0,428%
D	0,024%	0,056%	0,101%	0,149%	0,253%	0,452%
E	0,029%	0,065%	0,119%	0,168%	0,262%	0,836%
F	0,012%	0,043%	0,104%	0,159%	0,300%	0,533%
G	0,014%	0,044%	0,104%	0,144%	0,266%	0,530%

Table 30 Percentiles of absolute difference for variable lead time verification – compound Poisson method 2

Table 31 and Table 32 show the results of the third modeling option to include lead time variability, by using the two moments method.

	Average difference	Std.dev difference	Average absolute difference	Std.dev absolute difference
Total:	0,259%	0,709%	0,421%	0,627%
	Customer order quantity distributions			
1	0,200%	0,645%	0,296%	0,607%
2	0,210%	0,593%	0,316%	0,543%
3	0,237%	0,707%	0,436%	0,605%
4	0,385%	1,059%	0,701%	0,882%
5	0,218%	0,547%	0,299%	0,508%
6	0,278%	0,631%	0,415%	0,550%
7	0,386%	0,839%	0,600%	0,701%
8	0,192%	0,520%	0,270%	0,484%
9	0,242%	0,643%	0,413%	0,549%
10	0,246%	0,717%	0,460%	0,603%
	Cost price			
1000	0,431%	0,864%	0,585%	0,768%
5	0,088%	0,449%	0,256%	0,379%
	Replenishment lead time			
65	0,149%	0,534%	0,336%	0,441%
28	0,242%	0,598%	0,400%	0,506%
15	0,276%	0,631%	0,431%	0,537%
5	0,287%	0,695%	0,451%	0,601%
	Demand rate			
0.85	0,250%	0,767%	0,524%	0,613%
0.75	0,242%	0,712%	0,487%	0,572%
0.35	0,262%	0,586%	0,391%	0,508%
0.25	0,251%	0,568%	0,376%	0,494%
0.077	0,196%	0,513%	0,362%	0,413%
0.027	0,254%	0,579%	0,420%	0,473%
0.004	0,214%	0,570%	0,271%	0,545%
	Supplier			
A	0,070%	0,426%	0,257%	0,347%
B	0,039%	0,392%	0,234%	0,317%
C	0,304%	0,536%	0,415%	0,455%
D	0,540%	0,866%	0,711%	0,731%
E	0,778%	1,148%	0,859%	1,088%
F	0,036%	0,397%	0,238%	0,320%
G	0,048%	0,382%	0,229%	0,310%

Table 31 Means and standard deviations for variable lead time verification – two moments method

	50%	75%	90%	95%	99%	Max
Total:	0,203%	0,443%	1,082%	1,782%	3,184%	5,018%
	Customer order quantity distributions					
1	0,076%	0,242%	0,738%	1,590%	3,215%	4,664%
2	0,136%	0,305%	0,740%	1,466%	2,903%	4,536%
3	0,237%	0,481%	1,068%	1,668%	3,190%	4,667%
4	0,363%	0,748%	2,076%	2,614%	4,274%	5,018%
5	0,115%	0,311%	0,710%	1,415%	2,446%	4,503%
6	0,230%	0,450%	1,006%	1,505%	2,750%	4,248%
7	0,349%	0,751%	1,421%	2,135%	3,275%	4,569%

8	0,109%	0,244%	0,593%	1,329%	2,403%	4,375%
9	0,241%	0,429%	0,967%	1,484%	2,817%	4,718%
10	0,261%	0,499%	1,178%	1,637%	3,076%	4,729%
	Cost price					
1000	0,288%	0,694%	1,634%	2,253%	3,487%	5,018%
5	0,141%	0,307%	0,566%	0,907%	2,005%	4,450%
	Replenishment lead time					
65	0,202%	0,416%	0,794%	1,243%	2,308%	4,679%
28	0,213%	0,427%	1,174%	1,755%	2,842%	3,891%
15	0,203%	0,484%	1,228%	1,919%	3,221%	4,490%
5	0,188%	0,483%	1,205%	2,159%	4,173%	5,018%
	Demand rate					
0.85	0,230%	0,604%	1,582%	2,374%	3,732%	4,664%
0.75	0,224%	0,539%	1,516%	2,222%	3,407%	4,384%
0.35	0,179%	0,403%	1,257%	1,929%	3,031%	4,729%
0.25	0,187%	0,386%	1,120%	1,756%	2,845%	4,713%
0.077	0,255%	0,429%	0,852%	1,170%	2,312%	4,217%
0.027	0,304%	0,512%	1,026%	1,282%	2,272%	5,018%
0.004	0,095%	0,260%	0,601%	0,973%	2,827%	4,679%
	Supplier					
A	0,153%	0,318%	0,567%	0,864%	1,597%	4,413%
B	0,154%	0,305%	0,486%	0,681%	1,341%	4,450%
C	0,249%	0,625%	0,996%	1,271%	1,936%	4,785%
D	0,412%	1,198%	1,866%	2,129%	2,711%	4,455%
E	0,367%	1,143%	2,705%	3,278%	4,377%	5,018%
F	0,158%	0,314%	0,494%	0,740%	1,395%	4,387%
G	0,146%	0,294%	0,473%	0,701%	1,307%	4,134%

Table 32 Percentiles of absolute difference for variable lead time verification – two moments method

D.3. Timeframe

Table 33 and Table 34 show the results of the verification of the model including a timeframe of 15 days, by using conditioning on the replenishment lead time in the compound Poisson method.

	Average difference	Std.dev difference	Average absolute difference	Std.dev absolute difference
Total:	-0,011%	0,085%	0,056%	0,065%
	Customer order quantity distributions			
1	-0,005%	0,060%	0,036%	0,048%
2	-0,011%	0,087%	0,054%	0,069%
3	-0,015%	0,084%	0,054%	0,065%
4	-0,011%	0,081%	0,053%	0,062%
5	-0,002%	0,085%	0,057%	0,063%
6	-0,006%	0,085%	0,058%	0,063%
7	-0,025%	0,109%	0,078%	0,080%
8	-0,003%	0,083%	0,057%	0,060%
9	-0,007%	0,082%	0,056%	0,060%
10	-0,021%	0,083%	0,055%	0,065%
	Cost price			
1000	-0,018%	0,103%	0,072%	0,076%
5	-0,003%	0,061%	0,040%	0,046%

	Replenishment lead time			
65	-0,023%	0,117%	0,081%	0,088%
28	-0,013%	0,083%	0,058%	0,060%
15	-0,004%	0,067%	0,047%	0,049%
5	-0,001%	0,054%	0,036%	0,039%
	Demand rate			
0.85	-0,038%	0,097%	0,063%	0,083%
0.75	-0,026%	0,089%	0,057%	0,073%
0.35	0,001%	0,084%	0,055%	0,064%
0.25	-0,002%	0,082%	0,054%	0,063%
0.077	-0,002%	0,079%	0,055%	0,057%
0.027	-0,001%	0,079%	0,056%	0,055%
0.004	-0,003%	0,067%	0,048%	0,046%
	Supplier			
A	-0,012%	0,086%	0,055%	0,067%
B	-0,010%	0,084%	0,055%	0,064%
C	-0,011%	0,086%	0,058%	0,065%
D	-0,008%	0,080%	0,054%	0,060%
E	-0,014%	0,087%	0,059%	0,066%
F	-0,009%	0,085%	0,056%	0,065%
G	-0,010%	0,085%	0,055%	0,066%

Table 33 Means and standard deviations for timeframe verification

	50%	**75%**	**90%**	**95%**	**99%**	**Max**
Total:	0,034%	0,073%	0,132%	0,183%	0,311%	0,669%
	Customer order quantity distributions					
1	0,020%	0,046%	0,085%	0,116%	0,237%	0,515%
2	0,029%	0,069%	0,130%	0,188%	0,335%	0,561%
3	0,033%	0,069%	0,130%	0,183%	0,317%	0,519%
4	0,032%	0,070%	0,127%	0,173%	0,295%	0,635%
5	0,036%	0,074%	0,133%	0,180%	0,289%	0,563%
6	0,037%	0,076%	0,131%	0,179%	0,319%	0,472%
7	0,051%	0,107%	0,173%	0,243%	0,365%	0,601%
8	0,039%	0,076%	0,132%	0,180%	0,293%	0,400%
9	0,036%	0,073%	0,127%	0,184%	0,280%	0,493%
10	0,035%	0,070%	0,136%	0,181%	0,308%	0,669%
	Cost price					
1000	0,047%	0,096%	0,165%	0,230%	0,357%	0,669%
5	0,025%	0,052%	0,097%	0,129%	0,214%	0,635%
	Replenishment lead time					
65	0,051%	0,108%	0,199%	0,267%	0,423%	0,669%
28	0,037%	0,077%	0,140%	0,183%	0,281%	0,416%
15	0,032%	0,064%	0,109%	0,144%	0,223%	0,635%
5	0,024%	0,051%	0,086%	0,120%	0,183%	0,333%
	Demand rate					
0.85	0,034%	0,078%	0,153%	0,226%	0,408%	0,601%
0.75	0,032%	0,072%	0,137%	0,219%	0,352%	0,669%
0.35	0,034%	0,073%	0,132%	0,184%	0,331%	0,515%
0.25	0,032%	0,072%	0,130%	0,193%	0,309%	0,561%
0.077	0,035%	0,071%	0,132%	0,174%	0,271%	0,502%
0.027	0,039%	0,075%	0,129%	0,160%	0,250%	0,394%

0.004	0,034%	0,071%	0,114%	0,146%	0,209%	0,317%
	Supplier					
A	0,034%	0,069%	0,124%	0,178%	0,342%	0,563%
B	0,034%	0,071%	0,135%	0,183%	0,304%	0,669%
C	0,036%	0,077%	0,137%	0,186%	0,301%	0,635%
D	0,033%	0,072%	0,127%	0,173%	0,268%	0,561%
E	0,037%	0,079%	0,141%	0,189%	0,330%	0,477%
F	0,034%	0,070%	0,129%	0,187%	0,322%	0,493%
G	0,031%	0,072%	0,130%	0,183%	0,327%	0,601%

Table 34 Percentiles of absolute difference for timeframe verification

D.4. Fixed delivery days

D.4.1. Continuous review approximation

Table 35 and Table 36 show the results of the verification of the model including the fixed delivery days of suppliers, without a timeframe. The scenarios refer to the supplier number and the number of delivery days per week, as explained in section 5.3.2.

	Average difference	Std.dev difference	Average absolute difference	Std.dev absolute difference
Total:	-0,017%	0,086%	0,059%	0,065%
	Customer order quantity distributions			
1	-0,015%	0,063%	0,042%	0,049%
2	-0,021%	0,087%	0,059%	0,068%
3	-0,020%	0,090%	0,059%	0,070%
4	-0,021%	0,086%	0,060%	0,065%
5	-0,010%	0,081%	0,057%	0,058%
6	-0,012%	0,090%	0,061%	0,066%
7	-0,031%	0,106%	0,079%	0,077%
8	-0,005%	0,085%	0,057%	0,064%
9	-0,014%	0,085%	0,060%	0,062%
10	-0,020%	0,079%	0,057%	0,059%
	Cost price			
1000	-0,024%	0,102%	0,073%	0,075%
5	-0,010%	0,066%	0,045%	0,049%
	Replenishment lead time			
65	-0,024%	0,108%	0,076%	0,081%
28	-0,016%	0,084%	0,059%	0,062%
15	-0,013%	0,069%	0,050%	0,050%
5	-0,011%	0,066%	0,047%	0,047%
	Demand rate			
0.85	-0,044%	0,092%	0,069%	0,074%
0.75	-0,038%	0,091%	0,065%	0,074%
0.35	-0,013%	0,086%	0,059%	0,065%
0.25	-0,007%	0,080%	0,055%	0,059%
0.077	-0,004%	0,083%	0,057%	0,061%
0.027	-0,004%	0,074%	0,053%	0,052%
0.004	-0,004%	0,066%	0,048%	0,046%
	Scenario			
B-1	-0,006%	0,079%	0,053%	0,059%

B-2	-0,042%	0,085%	0,069%	0,065%
B-5	-0,002%	0,079%	0,052%	0,060%
C-1	-0,014%	0,085%	0,058%	0,065%
C-2	-0,032%	0,094%	0,066%	0,073%
C-5	-0,005%	0,085%	0,056%	0,063%

Table 35 Means and standard deviations for fixed delivery days verification – without timeframe

	50%	75%	90%	95%	99%	Max
Total:	0,038%	0,079%	0,139%	0,185%	0,304%	0,614%
	Customer order quantity distributions					
1	0,025%	0,055%	0,103%	0,146%	0,231%	0,369%
2	0,037%	0,077%	0,138%	0,191%	0,315%	0,550%
3	0,037%	0,076%	0,145%	0,185%	0,310%	0,571%
4	0,040%	0,081%	0,141%	0,188%	0,306%	0,405%
5	0,039%	0,080%	0,131%	0,171%	0,251%	0,561%
6	0,041%	0,084%	0,139%	0,189%	0,319%	0,614%
7	0,053%	0,113%	0,176%	0,229%	0,356%	0,594%
8	0,037%	0,076%	0,130%	0,171%	0,314%	0,517%
9	0,039%	0,086%	0,143%	0,181%	0,274%	0,363%
10	0,038%	0,075%	0,130%	0,184%	0,284%	0,387%
	Cost price					
1000	0,049%	0,101%	0,171%	0,222%	0,348%	0,614%
5	0,030%	0,061%	0,106%	0,138%	0,227%	0,594%
	Replenishment lead time					
65	0,050%	0,105%	0,185%	0,251%	0,402%	0,614%
28	0,039%	0,082%	0,142%	0,183%	0,299%	0,474%
15	0,034%	0,071%	0,124%	0,160%	0,240%	0,448%
5	0,033%	0,066%	0,111%	0,143%	0,210%	0,319%
	Demand rate					
0.85	0,043%	0,097%	0,171%	0,234%	0,355%	0,550%
0.75	0,042%	0,089%	0,157%	0,212%	0,360%	0,594%
0.35	0,038%	0,078%	0,141%	0,192%	0,305%	0,614%
0.25	0,035%	0,075%	0,134%	0,185%	0,301%	0,465%
0.077	0,040%	0,078%	0,131%	0,182%	0,278%	0,477%
0.027	0,037%	0,073%	0,125%	0,158%	0,256%	0,375%
0.004	0,034%	0,065%	0,114%	0,143%	0,207%	0,315%
	Scenario					
B-1	0,033%	0,070%	0,126%	0,170%	0,275%	0,517%
B-2	0,049%	0,098%	0,154%	0,192%	0,283%	0,614%
B-5	0,033%	0,068%	0,123%	0,165%	0,291%	0,504%
C-1	0,036%	0,077%	0,137%	0,178%	0,308%	0,594%
C-2	0,042%	0,090%	0,160%	0,211%	0,355%	0,571%
C-5	0,037%	0,074%	0,129%	0,179%	0,299%	0,570%

Table 36 Percentiles of absolute difference for fixed delivery days verification – without timeframe

Table 37 and Table 38 show the results of the verification of the model including the fixed delivery days of suppliers, with a timeframe of 15 days.

	Average difference	Std.dev difference	Average absolute difference	Std.dev absolute difference

Total:	-0,015%	0,074%	0,047%	0,059%
	Customer order quantity distributions			
1	-0,010%	0,056%	0,034%	0,046%
2	-0,016%	0,070%	0,044%	0,056%
3	-0,018%	0,073%	0,047%	0,059%
4	-0,020%	0,083%	0,052%	0,068%
5	-0,009%	0,071%	0,046%	0,054%
6	-0,016%	0,076%	0,050%	0,059%
7	-0,021%	0,089%	0,059%	0,070%
8	-0,011%	0,072%	0,045%	0,057%
9	-0,013%	0,072%	0,048%	0,056%
10	-0,018%	0,074%	0,046%	0,060%
	Cost price			
1000	-0,021%	0,088%	0,058%	0,068%
5	-0,010%	0,057%	0,036%	0,045%
	Replenishment lead time			
65	-0,022%	0,105%	0,073%	0,078%
28	-0,011%	0,066%	0,048%	0,046%
15	-0,018%	0,076%	0,050%	0,060%
5	-0,002%	0,023%	0,010%	0,021%
	Demand rate			
0.85	-0,030%	0,085%	0,053%	0,073%
0.75	-0,031%	0,090%	0,055%	0,077%
0.35	-0,009%	0,075%	0,044%	0,061%
0.25	-0,007%	0,071%	0,043%	0,057%
0.077	-0,009%	0,066%	0,042%	0,052%
0.027	-0,004%	0,062%	0,040%	0,047%
0.004	-0,002%	0,059%	0,039%	0,044%
	Scenario			
B-1	0,000%	0,063%	0,036%	0,051%
B-2	-0,043%	0,085%	0,059%	0,076%
B-5	-0,001%	0,070%	0,039%	0,057%
C-1	-0,009%	0,068%	0,047%	0,050%
C-2	-0,035%	0,079%	0,059%	0,063%
C-5	-0,004%	0,064%	0,042%	0,049%

Table 37 Means and standard deviations for fixed delivery days verification – timeframe of 15 days

	50%	75%	90%	95%	99%	Max
Total:	0,027%	0,064%	0,119%	0,160%	0,280%	0,669%
	Customer order quantity distributions					
1	0,017%	0,045%	0,083%	0,127%	0,215%	0,324%
2	0,023%	0,060%	0,116%	0,151%	0,256%	0,441%
3	0,026%	0,063%	0,116%	0,159%	0,295%	0,443%
4	0,029%	0,071%	0,132%	0,189%	0,312%	0,613%
5	0,028%	0,062%	0,117%	0,150%	0,253%	0,410%
6	0,032%	0,070%	0,121%	0,160%	0,291%	0,550%
7	0,037%	0,077%	0,146%	0,209%	0,298%	0,561%
8	0,025%	0,061%	0,110%	0,151%	0,251%	0,551%
9	0,029%	0,066%	0,123%	0,157%	0,248%	0,391%
10	0,026%	0,061%	0,117%	0,159%	0,272%	0,669%
	Cost price					

1000	0,035%	0,080%	0,146%	0,195%	0,319%	0,669%
5	0,020%	0,049%	0,090%	0,120%	0,206%	0,561%
	Replenishment lead time					
65	0,046%	0,098%	0,172%	0,235%	0,359%	0,669%
28	0,034%	0,069%	0,116%	0,147%	0,220%	0,390%
15	0,031%	0,066%	0,118%	0,155%	0,275%	0,416%
5	0,004%	0,020%	0,049%	0,074%	0,141%	0,210%
	Demand rate					
0.85	0,027%	0,075%	0,147%	0,205%	0,327%	0,408%
0.75	0,031%	0,071%	0,138%	0,205%	0,343%	0,669%
0.35	0,026%	0,063%	0,122%	0,163%	0,277%	0,561%
0.25	0,025%	0,062%	0,115%	0,157%	0,279%	0,425%
0.077	0,025%	0,060%	0,113%	0,146%	0,236%	0,461%
0.027	0,028%	0,059%	0,105%	0,134%	0,234%	0,390%
0.004	0,025%	0,060%	0,102%	0,134%	0,186%	0,289%
	Scenario					
B-1	0,019%	0,048%	0,095%	0,135%	0,224%	0,561%
B-2	0,032%	0,080%	0,159%	0,223%	0,335%	0,550%
B-5	0,018%	0,055%	0,107%	0,145%	0,259%	0,613%
C-1	0,030%	0,064%	0,109%	0,149%	0,230%	0,370%
C-2	0,039%	0,082%	0,137%	0,179%	0,284%	0,669%
C-5	0,025%	0,055%	0,098%	0,138%	0,243%	0,390%

Table 38 Percentiles of absolute difference for fixed delivery days verification – timeframe of 15 days

D.4.2. Order line fill rate calculation in periodic review method

Below the results of the verification of the calculation method of equation (5.9) for an arbitrary set of 200 fictional spare parts. As in all other verification studies, we have compared the outcomes of the calculation to the average of five simulation runs. The results are shown in Table 39.

	Difference	Absolute difference
Mean	3.652%	3.708%
Standard Devation	2.677%	2.598%
Minimum	-2.992%	0.088%
Maximum	10.500%	10.500%

Table 39 Results of periodic review verification

It is clear that these results are unacceptable with respect to the calculated order line fill rate.

We have looked at the calculation times in this verification as well. For spare parts with a high demand rate, replenishment lead time and average customer order quantity, the calculation of the order line fill rate took nearly 15 minutes. Since there are thousands of spare parts in real life, a calculation time of multiple minutes is not desirable.

D.5. Order fill rate

#Parts	Order size Distr.	Target Fill rate	Number of parts on stock	Average Order line fill rate
100	1	60%	54	75,206%
100	2	60%	67	87,603%

100	3	60%	79	84,421%
100	4	60%	66	83,636%
100	1	80%	74	83,929%
100	2	80%	86	88,402%
100	3	80%	87	91,129%
100	4	80%	84	87,462%
150	1	60%	78	81,021%
150	2	60%	107	88,338%
150	3	60%	125	87,588%
150	4	60%	107	82,781%
150	1	80%	118	86,177%
150	2	80%	132	91,082%
150	3	80%	134	93,803%
150	4	80%	131	91,196%
300	1	60%	163	81,940%
300	2	60%	221	86,744%
300	3	60%	240	89,470%
300	4	60%	206	84,683%
300	1	80%	224	89,388%
300	2	80%	267	90,036%
300	3	80%	273	91,796%
300	4	80%	260	90,278%

Table 40 Key results of order fill rate verification

#Parts	Order size Distr.	Target Fill rate	Price 150	Price 5	Price 0,05	LT 65	LT 28	LT 15	LT 5
100	1	60%	15%	63%	100%	65%	60%	37%	53%
100	2	60%	33%	80%	100%	77%	80%	48%	59%
100	3	60%	56%	89%	100%	88%	87%	70%	65%
100	4	60%	31%	80%	100%	73%	80%	48%	59%
100	1	80%	49%	83%	100%	85%	87%	59%	59%
100	2	80%	69%	94%	100%	88%	90%	85%	76%
100	3	80%	72%	94%	100%	88%	90%	89%	76%
100	4	80%	69%	89%	100%	88%	87%	81%	76%
150	1	60%	11%	55%	100%	61%	57%	32%	54%
150	2	60%	35%	86%	100%	80%	81%	50%	68%
150	3	60%	62%	92%	100%	90%	87%	71%	82%
150	4	60%	35%	86%	100%	80%	81%	50%	68%
150	1	80%	53%	88%	100%	88%	87%	59%	75%
150	2	80%	71%	96%	100%	90%	91%	85%	82%
150	3	80%	75%	96%	100%	93%	91%	88%	82%
150	4	80%	69%	96%	100%	90%	91%	82%	82%
300	1	60%	11%	61%	99%	59%	55%	47%	58%
300	2	60%	39%	87%	100%	77%	74%	69%	75%
300	3	60%	55%	89%	100%	83%	82%	75%	79%
300	4	60%	31%	81%	100%	73%	70%	62%	70%
300	1	80%	41%	88%	100%	79%	75%	69%	75%
300	2	80%	73%	96%	100%	90%	88%	90%	88%
300	3	80%	79%	96%	100%	91%	90%	94%	88%
300	4	80%	69%	93%	100%	87%	83%	89%	88%

Table 41 Detailed results of order fill rate verification - I

#Parts	Order size Distr.	Target Fill rate	Dem. rate 221	Dem. rate 195	Dem. rate 91	Dem. rate 65	Dem. rate 20	Dem. rate 7	Dem. rate 1
100	1	60%	67%	68%	77%	53%	59%	13%	15%
100	2	60%	80%	74%	92%	67%	82%	38%	15%
100	3	60%	93%	84%	92%	73%	100%	75%	23%
100	4	60%	100%	100%	100%	73%	100%	75%	23%
100	5	60%	80%	74%	92%	67%	76%	38%	15%
100	1	80%	87%	79%	92%	67%	88%	75%	23%
100	2	80%	100%	100%	100%	73%	100%	75%	38%
100	3	80%	100%	100%	100%	80%	100%	75%	38%
100	4	80%	100%	100%	100%	80%	100%	75%	38%
100	5	80%	100%	100%	100%	73%	100%	75%	23%
150	1	60%	68%	72%	71%	45%	39%	43%	24%
150	2	60%	80%	76%	94%	65%	82%	64%	33%
150	3	60%	96%	88%	100%	70%	100%	86%	38%
150	4	60%	96%	96%	100%	75%	100%	86%	52%
150	5	60%	80%	76%	94%	65%	82%	64%	33%
150	1	80%	88%	80%	94%	70%	93%	86%	38%
150	2	80%	96%	96%	100%	80%	100%	86%	52%
150	3	80%	100%	96%	100%	85%	100%	86%	52%
150	4	80%	100%	100%	100%	85%	100%	86%	52%
150	5	80%	96%	96%	100%	75%	100%	86%	52%
300	1	60%	69%	72%	75%	45%	45%	48%	15%
300	2	60%	83%	83%	93%	71%	80%	59%	33%
300	3	60%	91%	87%	93%	76%	90%	74%	38%
300	4	60%	98%	96%	95%	79%	96%	74%	41%
300	5	60%	83%	83%	88%	71%	69%	48%	21%
300	1	80%	87%	83%	93%	71%	80%	59%	36%
300	2	80%	98%	98%	98%	87%	96%	78%	56%
300	3	80%	100%	98%	100%	92%	100%	78%	56%
300	4	80%	100%	100%	100%	92%	100%	89%	56%
300	5	80%	98%	98%	98%	87%	96%	74%	41%

Table 42 Detailed results of order fill rate verification - II

#Parts	Order size Distr.	Target Fill rate	Scen. B-1	Scen. B-2	Scen. B-5	Scen. C-1	Scen. C-2	Scen. C-5
100	1	60%	60%	57%	62%	39%	50%	54%
100	2	60%	75%	79%	71%	50%	71%	54%
100	3	60%	95%	86%	71%	61%	79%	85%
100	4	60%	70%	79%	71%	50%	71%	54%
100	1	80%	80%	86%	71%	56%	79%	77%
100	2	80%	95%	93%	86%	72%	86%	85%
100	3	80%	95%	93%	86%	72%	86%	92%
100	4	80%	95%	93%	81%	67%	86%	85%
150	1	60%	62%	48%	54%	48%	39%	57%
150	2	60%	73%	83%	69%	63%	78%	67%
150	3	60%	96%	87%	74%	74%	89%	86%

150	4	60%	73%	83%	69%	63%	78%	67%
150	1	80%	88%	87%	69%	70%	83%	81%
150	2	80%	96%	91%	83%	81%	89%	90%
150	3	80%	96%	91%	89%	81%	89%	90%
150	4	80%	96%	91%	83%	81%	89%	86%
300	1	60%	55%	52%	52%	57%	46%	62%
300	2	60%	73%	83%	68%	73%	71%	76%
300	3	60%	82%	86%	71%	79%	76%	88%
300	4	60%	62%	79%	63%	71%	68%	72%
300	1	80%	73%	83%	70%	73%	71%	80%
300	2	80%	91%	90%	84%	89%	90%	90%
300	3	80%	93%	93%	88%	91%	90%	92%
300	4	80%	89%	90%	80%	86%	85%	90%

Table 43 Detailed results of order fill rate verification - III

#Parts	Order size Distr.	Target Fill rate	Cust. Ord. Qt. 1	Cust. Ord. Qt. 2	Cust. Ord. Qt. 3	Cust. Ord. Qt. 4	Cust. Ord. Qt. 5	Cust. Ord. Qt. 6	Cust. Ord. Qt. 7	Cust. Ord. Qt. 8	Cust. Ord. Qt. 9	Cust. Ord. Qt. 10
100	1	60%	50%	67%	78%	78%	50%	44%	50%	56%	63%	29%
100	2	60%	57%	100%	78%	78%	83%	89%	60%	61%	75%	43%
100	3	60%	79%	100%	100%	89%	100%	89%	60%	72%	88%	57%
100	4	60%	57%	100%	78%	78%	83%	78%	60%	61%	75%	43%
100	1	80%	64%	100%	100%	89%	100%	89%	60%	72%	75%	43%
100	2	80%	100%	100%	100%	89%	100%	89%	80%	83%	88%	57%
100	3	80%	100%	100%	100%	89%	100%	89%	90%	83%	88%	57%
100	4	80%	86%	100%	100%	89%	100%	89%	80%	83%	88%	57%
150	1	60%	40%	78%	67%	86%	40%	50%	46%	50%	58%	30%
150	2	60%	60%	89%	83%	86%	93%	75%	62%	67%	75%	45%
150	3	60%	87%	89%	100%	93%	100%	88%	62%	88%	83%	55%
150	4	60%	60%	89%	83%	86%	93%	75%	62%	67%	75%	45%
150	1	80%	73%	89%	100%	93%	100%	88%	62%	79%	75%	45%
150	2	80%	100%	100%	100%	93%	100%	88%	85%	88%	92%	55%
150	3	80%	100%	100%	100%	93%	100%	88%	92%	88%	92%	60%
150	4	80%	100%	100%	100%	93%	100%	88%	77%	88%	92%	55%
300	1	60%	55%	83%	59%	64%	59%	50%	28%	55%	61%	30%
300	2	60%	65%	88%	94%	89%	93%	70%	55%	68%	68%	43%
300	3	60%	81%	88%	100%	93%	100%	87%	55%	75%	71%	43%
300	4	60%	61%	83%	74%	89%	83%	63%	55%	66%	68%	43%
300	1	80%	68%	88%	94%	89%	93%	73%	55%	68%	71%	43%
300	2	80%	94%	100%	100%	93%	100%	90%	76%	89%	86%	57%
300	3	80%	94%	100%	100%	93%	100%	90%	86%	89%	93%	61%
300	4	80%	84%	88%	100%	93%	100%	87%	76%	89%	86%	57%

Table 44 Detailed results of order fill rate verification - IV

E. Results from real life case study

E.1. Number of orders for spare parts on stock

Table 45 shows the number of orders in the past eight years for the spare parts that are put on stock in the basic scenario.

#of past orders	#of spare parts
1	401
2	168
3	106
4	75
5	42
6-10	170
11-20	166
21-30	85
31-40	45
41-50	28
>50	89

Table 45 Number of past orders for spare parts on stock – basic scenario

Table 46 shows the number of orders in the past eight years for the spare parts that are put on stock:

#of past orders	#of spare parts
1	428
2	153
3	84
4	72
5	41
6-10	164
11-20	153
21-30	82
31-40	45
41-50	24
>50	86

Table 46 Number of past orders for spare parts on stock – alternative model

Applying the adjustments mentioned in section 6.2, the following results are obtained.

#of past orders	#of spare parts
1	0
2	0
3	19
4	39
5	35
6-10	161
11-20	186
21-30	94
31-40	50
41-50	30
>50	92

Table 47 Number of past orders for spare parts on stock – alternative model + excluding low demand spare parts

E.2. Scenarios

Scenario A

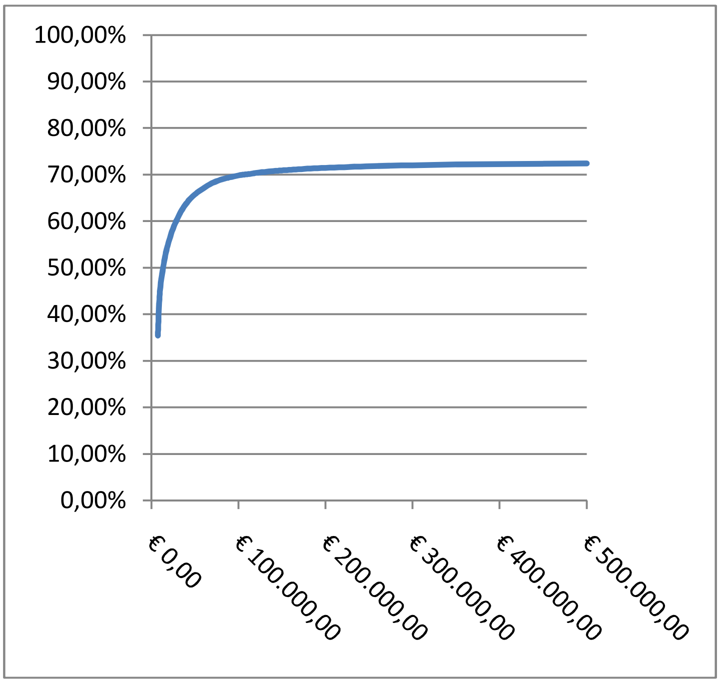

Timeframe	(1)	(2)
0	26.03%	49.15%
10	44.08%	60.97%
15	60.21%	70.00%
20	82.30%	90.05%
30	92.86%	95.73%
KOSTEN	€29,066	€104,333

Scenario B

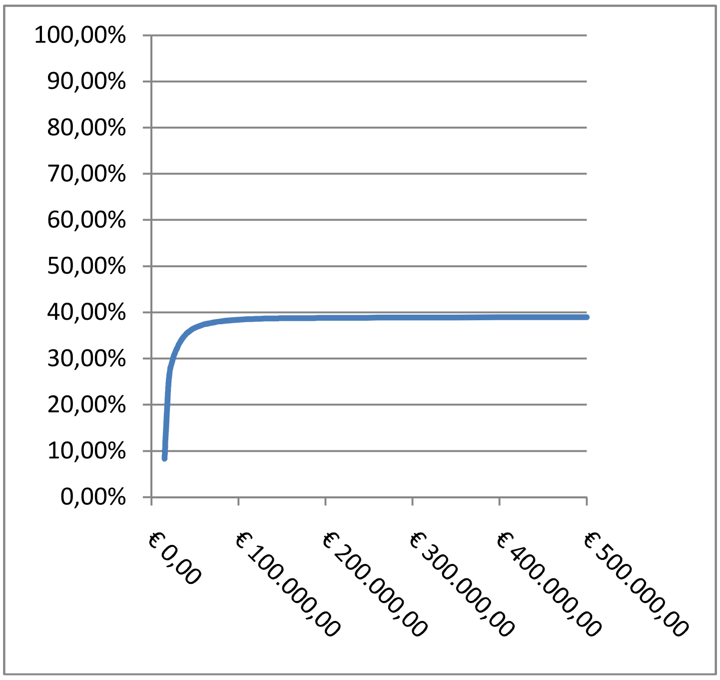

Timeframe	(1)	(2)
0	35.00%	38.00%
10	45.87%	48.00%
15	57.06%	58.82%
20	81.40%	83.02%
30	91.91%	92.63%
KOSTEN	€38,402	€77,663

Scenario C

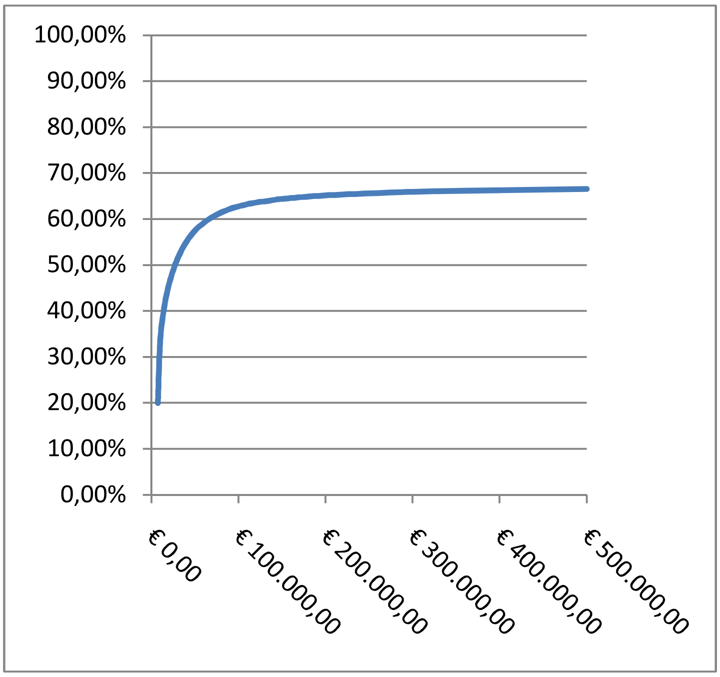

Timeframe	(1)	(2)
0	37.85%	52.76%
10	55.00%	63.50%
15	63.21%	70.39%
20	84.53%	90.15%
30	93.39%	95.75%
KOSTEN	€39,938	€117,633

Scenario D

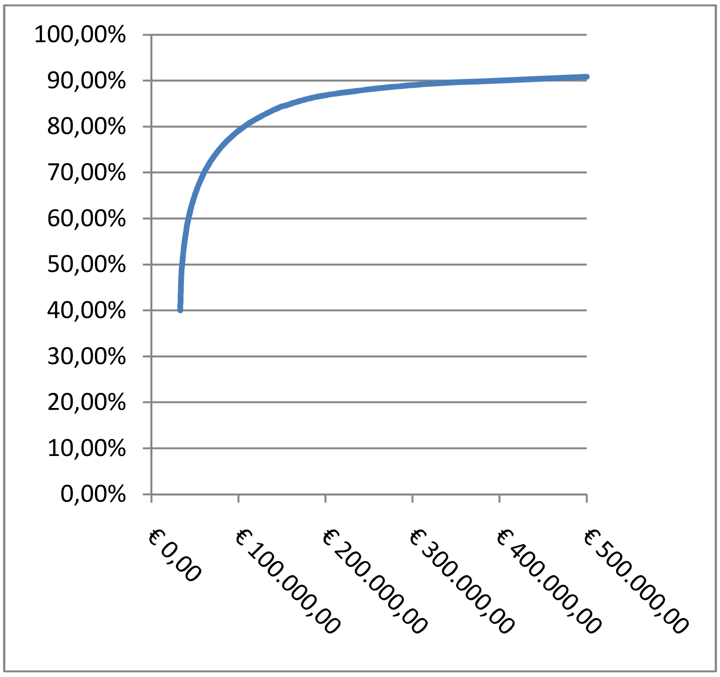

Timeframe	(1)	(2)
0	34.79%	74.57%
10	53.93%	83.53%
15	70.00%	90.00%
20	87.05%	96.81%
30	94.02%	98.52%
KOSTEN	€60,574	±€400,000

Scenario E

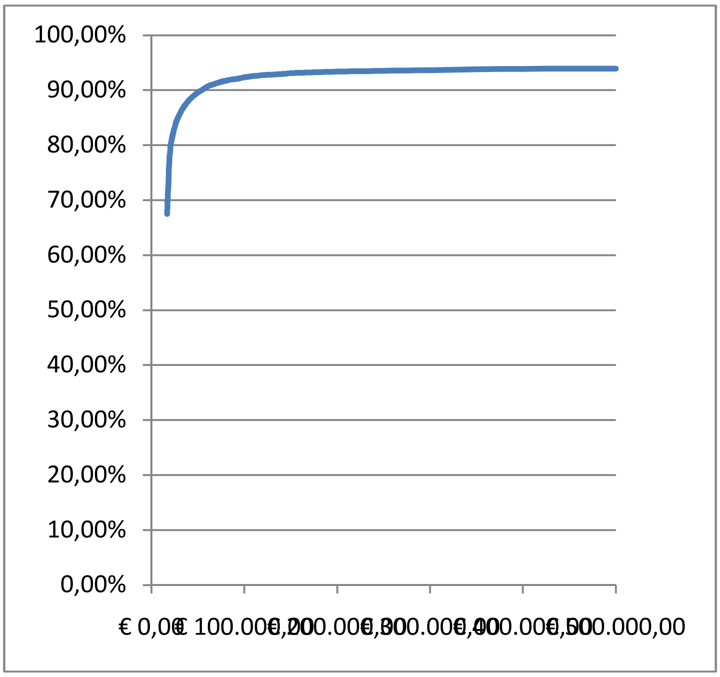

Timeframe	(1)	(2)
0	20.06%	42.30%
10	36.34%	51.99%
15	51.15%	58.84%
20	85.00%	92.50%
30	93.51%	96.90%
KOSTEN	€28,545	€106,580

E.3. Separation of demand streams

Scenario	Order fill rate service orders	Order fill rate prev.maint.	Start with prev.maint.	Order line fill rate replenishment
A-(1)	60.21%	61.28%	60.22%	64.29%
A-(2)	70.00%	70.23%	70.01%	52.54%
B-(1)	35.00%	36.54%	35.00%	40.84%
B-(2)	38.00%	38.31%	38.00%	18.38%
C-(1)	55.00%	56.43%	55.00%	52.15%
C-(2)	63.50%	63.84%	63.51%	33.91%
D-(1)	70.00%	71.16%	70.00%	71.10%
D-(2)	Execution failed			
E-(1)	85.00%	85.82%	85.00%	88.99%
E-(2)	92.50%	92.66%	92.50%	82.21%

Table 48 Order fill rates for separation of demand streams experiments

F. Model characteristics

This is the complete list of model characteristics as is used by (Hoving, 2008) for the spare parts supply chain at Vanderlande.

- *Number of spare parts.* There are around 27,000 different items that are labeled as spare part within Vanderlande.
- *Cost price of spare parts.* Excluding the most expensive 1% of all spare parts, the cost price of spare parts lies between €0.01 and €700. Including all spare parts, the most expensive part costs €3,000.-.
- *Lead time of spare parts.* Spare parts are ordered either at the factory of Vanderlande or at external suppliers; lead time is measured in weeks. Spare parts from the former category have a lead time between 1 and 7 weeks, excluding the 5% spare parts with largest lead time. The absolute maximum is 20 weeks. For spare parts from external suppliers the lead time lies between 1 and 8 weeks (excluding 5% largest lead times), with an absolute maximum of 31 weeks. This standard lead time is not fixed for all suppliers; some suppliers are more reliable in meeting the standard lead time than others.
- *Spare parts supply.* Spare parts are ordered at the external supplier by the piece or in batches, a.o. depending on the spare part characteristics (price, average order quantity).
- *Demand per year.* In the period we reviewed, spare parts were demanded on average between 0 and 8 times per year, excluding the 5% most frequently ordered parts. One part was ordered 213 times per year, which is the maximum. Over 22,000 spare parts were not ordered during the review period.
- *Order quantity.* The order quantity represents the amount of one specific spare part in one order. Excluding the 5% largest amounts, spare parts are ordered in quantities from 0.1 to 150. The highest order quantity was 50,000, this concerned a large order for washers (which are always ordered in very large amounts).
- *Costs related to spare parts.* The costs for holding spare parts on stock and transport to the customer are both dependent on the size and weight of the spare part.
- *Size and shape of spare parts.* The volume, height and weight of all spare parts are known. They are mainly used to determine holding and transportation costs. The shape of spare parts is not important in our research.
- *Order characteristics.* Spare parts are not ordered one-by-one, but they are grouped in customer orders. Per year, around 4,500 orders are placed, with an average of 2.6 order lines (i.e. different spare parts) per order. The number of different spare parts per order lies between 1 and 20, excluding the 1% orders with the most order lines. The maximum number of order lines was 126. Most customers agree on the lead time based on the spare parts lead time from the factory or external supplier.
- *Order handling.* All orders (and possible backorders) are handled following a First Come, First Served principle.
- *Service differentiation.* Within all service orders, we want to differentiate in desired performance between preventive maintenance orders (more important, higher performance) and replenishment orders (less important, lower performance). For the basic model we do not apply service differentiation.

- *Obsolescence.* We assume that spare parts do not become obsolete. In practice this is not completely true; some parts may be replaced in time by other, better parts. Mostly these new parts are completely substitutable, so we do not expect major problems by assuming no obsolescence.
- *Criticality.* We assume that all spare parts are critical, i.e. for all spare parts, failure of the part causes failure of the system. In practice there are spare parts that are not critical to the customers' system, but to achieve the highest performance in the supply of spare part orders we assume all parts to be critical to the system.
- *Number of echelons in supply chain.* Formally, the supply chain consists of multiple echelons; the central service warehouse, local warehouses, and system location stock. In our research, we will assume a single echelon; we will observe the demand at the central service warehouse as if it origins from one demand location.
- *Single- or multi-indenture.* We do not look at the relation between different spare parts within a bill of material. Hence, we apply a single-indenture model.
- *Lateral transshipments.* We do not include the possibility of lateral transshipments in our model. This means that all demand at the central warehouse is fulfilled via the central warehouse; no other (local) warehouses can supply in case the central warehouse has a stock-out.
- *Commonality.* Items can occur in multiple types of orders.
- *Static environment.* We assume that the supply- and demand processes are similar throughout the year; there are no seasonal or other influences to these processes.
- *Transportation modes.* We are not interested in the number of transportation modes. We try to optimize the availability of the orders at the central service warehouse, so before the shipment to the customer.
- *Backordering.* We assume that all orders that cannot be fulfilled immediately are fully backlogged.
- *Reparability.* We do not take the reparability of spare parts into account.

G. Literature study (in Dutch)

Achtergrond onderzoek

Het onderzoek dient ter afronding van twee masterstudies: Applied Mathematics en Industrial Engineering & Management. Het onderzoek wordt uitgevoerd bij Vanderlande Industries BV in Veghel, een internationale leverancier van 'materials handling systems' voor voornamelijk vliegvelden en distributiecentra, specifiek bij de serviceafdeling van Vanderlande Industries. Deze afdeling is verantwoordelijk voor de levering van reserveonderdelen om het systeem draaiende te houden, en maakt hierbij gebruik van een centraal magazijn. Het doel van het onderzoek is om een voorraadmodel voor dit magazijn op te stellen dat de leverprestatie richting de klanten verbeterd.

Aanleiding literatuuronderzoek

In de meeste voorraadmodellen die op reserveonderdelen gebaseerd zijn, wordt aangenomen dat onderdelen één voor één worden besteld. Bij Vanderlande is het echter zo dat klanten vaak meerdere (verschillende) onderdelen tegelijk bestellen. De geijkte prestatiemaat item fill rate, i.e. het percentage onderdelen dat uit voorraad kan worden geleverd, is daarom niet de gewenste prestatiemaat. Tevens is Vanderlande niet zozeer geïnteresseerd in het meteen uitleveren van onderdelen, maar meer in het uitleveren binnen een gestelde termijn. Deze prestatiemaat wordt time-based fill rate genoemd, i.e. het percentage complete orders dat binnen een gestelde levertermijn wordt uitgeleverd. Het literatuuronderzoek is er op gericht om geschikte modellen te vinden die deze prestatiemaat hanteren. Op basis van de overige eigenschappen die de gevonden modellen hebben kan dan een keuze worden gemaakt om één of meerdere modellen uit te werken voor de specifieke situatie bij Vanderlande.

Zoekcriteria

Het belangrijkste zoekcriterium is order fill rate, dit is de prestatiemaat waar we modellen bij zoeken. De overige criteria zijn opgesteld op basis van de eigenschappen van de situatie:

- Spare parts / service parts
- Low demand
- Batch demand
- Single-echelon
- Single-indenture
- Backordering

De laatste drie criteria is niet expliciet op gezocht, uit de abstracts van de eerder gevonden artikelen bleek dat alle gevonden artikelen aan deze eigenschappen voldeden.

Literatuuronderzoek

Ik heb mijn onderzoek uitgevoerd via twee zoekmachines. Allereerst heb ik mijn zoekcriteria ingevoerd in Web of Science. Later heb ik dezelfde combinaties ingevoerd in Scopus. In vrijwel alle gevallen gaf Scopus meer resultaten. Per combinatie zal ik hieronder aangeven hoeveel zoekresultaten ik heb gevonden op iedere website en welke artikelen ik geselecteerd heb. Het selecteren van artikelen heb ik voornamelijk gedaan op basis van aantal keren geciteerd, bron (e.g. Management Science en Operations Research staan hoog aangeschreven) en uiteraard inhoud op basis van het abstract. Tevens heb ik in sommige gevallen doorgezocht naar artikelen die een interessant artikel geciteerd hebben.

Zoekterm: "order fill rate"

Web of Science: 10 resultaten
Scopus: 26 resultaten

Gevonden artikelen:

- (Thorstenson, et al., 2008) (WoS)
- (Lu, 2008) (WoS)
- (Lu, et al., 2003) (WoS)
- (Song, 1998) (WoS)
- (Vliegen, et al., 2008) (Scopus)
- (Song, 2000) (Scopus)

Zoekterm: Alle artikelen die (Song, 1998) geciteerd hebben

Gevonden artikelen:

- (Thorstenson, et al., 2008) (al eerder gevonden)
- (Lu, et al., 2003) (al eerder gevonden)
- (Song, et al., 2002)

Zoekterm: Alle artikelen die (Lu, et al., 2003) geciteerd hebben

Gevonden artikelen:

- (Lu, 2008) (al eerder gevonden)

Zoekterm: "spare parts" & onderwerp: OR & Management Science
Web of Science: 148 resultaten
Scopus: 340 resultaten

Toegevoegd: Inventory
Web of Science: 99 resultaten
Scopus: 125 resultaten

Gevonden artikelen:

- (Sani, et al., 1997) (WoS)
- (Eaves, et al., 2004) (WoS)
- (Strijbosch, et al., 2000) (Scopus)

Zoekterm: time based order fill rate (N.B.: zonder aanhalingstekens)
Web of Science: 30 resultaten
Scopus: 82 resultaten (exclusief niet relevante onderwerpen: 42 resultaten)

Gevonden artikelen:

- (Yao, et al., 2006) (WoS)
- (Wang, et al., 2005) (WoS)
- (Zhao, 2008) (Scopus)

Zoekterm: "service parts" en inventory
Web of Science: 35 resultaten
Scopus: 73 resultaten

Gevonden artikelen:

- (Cohen, et al., 1997) (WoS)
- (Caggiano, et al., 2007) (WoS)
- (Caggiano, et al., 2008) (Scopus)

- (Fortuin, et al., 1999) (Scopus)

Zoekterm: order fill rate en service parts
Web of Science: 2 resultaten
Scopus: 8 resultaten

Gevonden artikelen:

- (Kranenburg, et al., 2007) (WoS)

Zoekterm: “low demand” en order
Web of Science: 33 resultaten
Scopus: 99 resultaten

Gevonden artikelen:

- (Pujawan, 2004) (WoS)
- (Sox, et al., 1997) (Scopus)

Tweede selectieronde

De artikelen uit deze eerste selectie heb ik uitvoeriger bestudeerd. Door de artikelen door te lezen heb ik een tweede selectie gemaakt van artikelen die de situatie bij Vanderlande het beste benaderen. In sectie 3.2 is aangegeven welke artikelen hier zijn geselecteerd en welke aannames ze maken.

Vervolg

In eerste instantie heb ik het artikel van (Lu, 2008) geselecteerd om uit te werken. Zoals echter al in de laatste alinea aangegeven bleek de situatie bij Vanderlande na nader (data-)onderzoek niet geschikt voor het soort modellen dat ik in de literatuurstudie heb gezocht. In overleg met mijn begeleiders heb ik dan ook gekozen om een andere benadering te kiezen, die gebaseerd is op (eenvoudigere) item fill rate modellen. Middels een aantal studieboeken aangeleverd door mijn begeleiders heb ik deze benadering verder uit kunnen werken.

Evaluatie

Dit literatuuronderzoek was bedoeld als startpunt voor het modelleringgedeelte van mijn afstudeeronderzoek. Op basis van de situatie zoals in kaart gebracht in het eerste gedeelte van mijn onderzoek leek het de juiste benadering om order fill rate modellen te observeren. Door de literatuurstudie heb ik een beter beeld gekregen de mogelijkheden binnen dit soort modellen. Hiermee is echter boven water gekomen dat deze benadering niet geschikt is voor de situatie bij Vanderlande. Het literatuuronderzoek heeft dus zeker resultaat gehad, maar niet het resultaat – i.e. een geschikt model – waar ik in eerste instantie op hoopte.

Bibliography

Axsäter, S. 2006. *Inventory Control.* Second edition. Lund : Springer, 2006.

Caggiano, K.E., et al. 2008. Efficient computation of time-based customer service levels in a multi-item, multi-echelon supply chain: A practical approach for inventory optimization. *European Journal of Operational Research.* 2008, Artical in Press.

Caggiano, K.E., et al. 2007. Optimizing service parts inventory in a multiechelon, multi-item supply chain with time-based customer service-level agreements. *Operations Reserach.* 2007, Vol. 55, 2, pp. 303-318.

Chang, P., Chou, Y. and Huang, M. 2005. A (r,r,Q) inventory model for spare parts involving equipment criticality. *International Journal of Production Economics.* 2005, Vol. 97, pp. 66-74.

Cohen, M.A., Zheng, Y-S. and Agrawal, V. 1997. Service parts logistics: a benchmark analysis. *IIE Transactions.* 1997, Vol. 29, 8, pp. 627-639.

Eaves, A.H.C. and Kingsman, B.G. 2004. Forecasting for the ordering and stock-holding of spare parts. *Journal of the Operations Research Society.* 2004, Vol. 55, 4, pp. 431-437.

Federgruen, A. and Zheng, Y-S. 1992. An efficient algorithm for computing an optimal (r,Q) policy in continuous review stochastic inventory systems. *Operations Research.* 1992, Vol. 40, pp. 808-813.

Feeney, G.J. and Sherbrooke, C.C. 1966. The (s-1,s) inventory policy under compound poisson demand. *Management Science.* 1966, Vol. 12, 5, pp. 391-411.

Fisher, M.L. 1981. The Lagrangian relaxation method for solving integer programming problems. *Management Science.* 1981, Vol. 27, 1, pp. 1-18.

Fortuin, L. and Martin, H. 1999. Control of service parts. *International Journal of Operations and Production Management.* 1999, Vol. 19, 9, pp. 950-971.

Harris, F.W. 1913. How many parts to make at once. *Factory, the magazine of management.* 1913, Vol. 10, 2, pp. 135-136.

Hoving, E.J. 2008. *Spare Parts Inventory Control at Central Warehouses.* University of Groningen : MSc. Thesis, 2008.

Kranenburg, A.A. and Van Houtum, G.J. 2007. Effect of commonality on spare parts provisioning costs for capital goods. *International Journal of Production Economics.* 2007, Vol. 108, 1-2, pp. 221-227.

Law, A.M. and Kelton, W.D. 2000. *Simulation modeling and analysis.* Third edition. : McGraw-Hill, 2000.

Lu, Y. 2008. Performance analysis for assemble-to-order system with general renewal arrivals and random batch demands. *European Journal of Operational Research.* 2008, Vol. 185, pp. 635-647.

Lu, Y., Song, J-S. and Yao, D.D. 2003. Order fill rate, leadtime variability, and advance demand information in an assemble-to-order system. *Operations Research.* 2003, Vol. 51, 2, pp. 292-308.

Pujawan, I.N. 2004. The effect of lot sizing rules on order variability. *European Journal of Operational Research.* 2004, Vol. 159, 3, pp. 617-635.

Sani, B. and Kingsman, B.G. 1997. Selecting the best periodic inventory control and demand forecasting methods for low demand items. *Journal of the Operational Research Society.* 1997, Vol. 48, 7, pp. 700-713.

Sherbrooke, C.C. 2004. *Optimal inventory modeling of systems: multi-echelon techniques.* Second edition. Boston : Kluwer academic publishers, 2004.

Song, J-S. 2000. A note on assemble-to-order systems with batch ordering. *Management science.* 2000, Vol. 46, 5, pp. 739-743.

Song, J-S. and Yao, D.D. 2002. Performance analysis and optimization of assemble-to-order systems with random lead times. *Operations Research.* 2002, Vol. 50, 5, pp. 889-903.

Song, J-S. 1998. On the order fill rate in a multi-item, base-stock inventory system. *Operations Research.* 1998, Vol. 46, 6, pp. 831-845.

Sox, C.R., Thomas, L.J. and McClain, J.O. 1997. Coordinating production and inventory to improve service. *Management Science.* 1997, Vol. 43, 9, pp. 1189-1197.

Strijbosch, L.W.G., Heuts, R.M.J. and Van der Schoot, E.H.M. 2000. Combined forecast - inventory control procedure for spare parts. 2000, Vol. 51, 10, pp. 1184-1192.

Thorstenson, A. and Larsen, C. 2008. A comparison between the order and the volume fill rate for a base-stock inventory control system under a compound renewal demand process. *Journal of the Operational Research Society.* 2008, Vol. 59, 6, pp. 798-804.

van Sommeren, R.J.M. 2007. *Spare parts inventory control for a single-echelon, multi-location model with lateral transhipments.* Eindhoven University of Technology : MSc. Thesis, 2007.

Vliegen, I.M.H. and Van Houtum, G.J. 2008. Approximate evaluation of order fill rates for an inventory system of service tools. *International Journal of Production Economics.* 2008, Article in Press.

Wang, T., Chen, Y.H. and Feng, Y.Y. 2005. On the time-window fulfillment rate in a single-item min-max inventory control system. *IIE Transactions.* 2005, Vol. 37, 7, pp. 667-680.

Wilson, R.H. 1934. A scientific routine for stock control. *Harvard Business Review.* 1934, Vol. 13, pp. 116-128.

Yao, Z. and Simchi-Levi, D. 2006. Performance analysis and evaluation of assemble-to-order systems with stochastic sequential lead times. *Operations Research.* 2006, Vol. 54, 4, pp. 706-724.

Zhao, Y. 2008. Analysis and evaluation of an Assemble-to-Order system with batch ordering policy and compound Poisson demand. *European Journal of Operational Research.* 2008.

Zwillinger, D. 2003. *CRC standard mathematical tables and formulae.* 31st edition. Boca Raton : Chapman & Hall, 2003.