The Bigger Picture

Optimizing material availability in a contract assembly setting

Remy Gankema
Author

Dr. Engin Topan
Supervisor

Maaike Groot Dengerink, MSc
Company supervisor

Dr. Ipek Seyran Topan
Co-supervisor

Master's thesis Industrial Engineering and Management
September 3, 2020

THALES

UNIVERSITY OF TWENTE.

Preface

This thesis presents the findings of the graduation assignment for my master Industrial Engineering and Management. I am pleased to show you what I have learned and achieved during over half a year of research. The company I performed my research at, Thales Nederland, gave me the opportunity to define my own research as long as it helped improving service levels by optimising material availability. While this was definitely challenging at times, it also greatly aided in developing research skills and finding the right focus. I could not have achieved this by myself, and therefore I would like to express a few words of gratitude.

First of all, I would like to thank Engin for our useful discussions, extensive feedback, and most importantly for keeping me on track. I also would like to thank Ipek for her advices and contributions to improve my thesis. I would like to thank Maaike for her daily supervision, proofreading, having answers to most of my questions and else showing me directions. Also to my other colleagues at Thales, thank you for all your help and for making it an enjoyable period. Additionally, I would like to thank my friends and family for providing the sometimes much needed distraction during my thesis, but especially for making my student's life as good as it has been. Lastly, I would like to thank my girlfriend, for her unconditional support and understanding during all those times I pondered on my research amidst our conversations. For now, I am looking forward to the next step in my life!

Remy Gankema
September 3, 2020

Management summary
The product chain Aerospace of Thales Netherlands faced an *on-time delivery rate* (OTD) of 30-40% in the past few years. After an internal audit, an end-to-end improvement team executed many revision projects of the process, and this increased the OTD towards 80-90%. Despite these positive changes, there is still room for improvement. In order to become more competitive, Thales wants to increase OTD towards its goal of 98%, while reducing customer lead times and costs related to purchasing, holding items and production delay. One approach towards reaching that goal is increasing material availability. This thesis describes several tactical decisions in material management that help to achieve this higher material availability.

Finding research opportunities
At the start, there was no clear-cut solution-path. Uncertainty is generally the hardest thing to deal with in a supply chain. However, as there are numerous options to consider which require thorough analysis to assess their feasibility, we first enter explorative research. Initially, we consider both uncertainty in demand and supply. We observe that the current forecasts are not accurate enough to base purchasing decisions on. At the same time, since not Thales but the customer is the product owner of Aerospace products, there are more risks than usual with uncertain demand. Examples are uncontrolled revision of materials and completely obsolete stock if demand stops. We identify these risks, but decide to leave further dealing with uncertain demand for future work.

We also see that there is much more to be improved regarding supply uncertainty. We observe that safety lead times are roughly the same for most materials. As each supplier performs differently and its price (and therefore holding costs) also differ, these rather constant safety lead times can never be optimal. We decide that optimising these safety lead times is the core problem to tackle in this research. It aligns very well with Thales' purchase-to-order policy, and we expect to improve much on this issue. Furthermore, we see the need to quickly identify materials on their importance and risks. For that, we choose to design a classification method. It aids in assessing which materials need more attention, or possibly safety stocks. The combination of this classification method, optimising safety lead times, and suggesting process changes together help in dealing with supply uncertainty. Finally, we also observe that quantity discounts play a large role in costs of materials, but that holding costs are not always considered at purchasing decisions. We therefore decide to build a *Total Cost of Ownership* (TCO) model to make decisions given those quantities and other factors. We consider this a separate improvement, that allows to reduce costs in various settings.

Designing solutions
After studying relevant literature on safety lead time models and classification methods, we build a mathematical model to determine the optimal safety lead times. It optimises the sum of holding costs of materials being earlier than needed, holding costs of materials being longer on-stock while waiting on late materials, and costs of production delay (tardiness costs). In this model we use delivery lateness data, which we define as delivery delay caused by suppliers. The model does not fully capture the setting at Aerospace, but it still offers good insight in optimising safety lead times. Additionally, to more reliably measure the results of the model, we do a simulation study which incorporates more important details of reality compared to the mathematical model. The model can be built and solved in Excel, making it easily integrable at Thales. Implementing safety lead times require changes in the ERP system with a three-months test period.

We also suggest additions to the communication process of sales and procurement. The model only captures material lateness caused by suppliers, but delays may also occur on behalf of Thales or the customer. The process changes aim to improve that issue. These process changes can be summarised as earlier and more communication between sales and procurement, and also regular updates on procurement lead times that the suppliers can actualise. Furthermore, we design the TCO model. It is kept simple, and allows insight in the total costs given current and future demand and an order quantity to purchase now.

In addition, we design a material classification method. Together with colleagues, we decide that important parameters to classify materials on are slack to the critical path (a scale to place procurement lead time in relation to those of other materials in a final assembly), material costs per system, revision risk and forecasted demand (whether we expect *any* demand in the next year). We suggest five policies regarding what to do with materials in a specific category, which mainly focus on qualitative advices about safety stocks and using the TCO model. It allows for easy decision making, but we also note that it is only a start. We propose to further expand the policies or even quantify them, and to consider other parameters.

Results
The key result of the safety lead times optimisation model is that we can expect a total cost reduction between 15% and 50% for three evaluated Aerospace products (not named for confidentiality reasons). In absolute numbers, this would already be hundreds of thousands euros for a year's production of those products alone. We find these results by the mathematical model, but also by the more realistic simulation study. In addition to a reduction in costs, we also observe much less production delay (decreases of 5-60% of expected lateness per batch). The actual savings do depend a lot on the tardiness costs that we expect to make, which is very difficult to measure. A study on quantifying these costs may therefore be worthwhile.

Initially, we expected the lack of delivery lateness data for 75% of the suppliers a major issue. Fortunately, it turns out that the material's cost price has much more impact on the optimal safety lead times, allowing us to even find fairly good safety lead times in case we lack delivery lateness data for a supplier. Furthermore, we see that the model finds a solution quickly for Aerospace products, but requires a long run-time for radar systems (up to more than a full day when considering 2,000 materials). It may be necessary to decouple such a large final assembly into smaller subassemblies to analyse individually. Also, when considering to use this model outside Aerospace, the impact of commonality (usage of materials in multiple final assemblies) on optimal safety lead times must be considered.

Recommendations
Except that we recommend to implement the proposed solutions, we gain key insights for managerial implications from this research. First, we advise to shift the mindset for most procurement and inventory decisions from a material view to a final assembly view. Only then we can properly assess each material and make the right decisions. For example, a material with 20 weeks procurement lead time may not seem critical on its own, but if all other required materials for the final assembly have lower lead times, it is. At the same time, a material with 40 weeks procurement lead time is much less relevant if many other materials for the same final assembly exceed 60 weeks lead time. By adopting this mindset, overviewing materials and making the right decisions becomes more simple and effective.

Furthermore, we recommend to intensify decision making based on delivery lateness data and improve recording these data. Currently, delivery lateness data are mainly used to assess whether a supplier is performing well and if actions are required. However, supply uncertainty is the major factor of uncertainty in a contract assembly setting, so quantitative decisions on purchasing and inventory should include that as factor. Additionally, these data should not be polluted, as now can be the case due to e.g. not wanting to make a supplier look bad. Keeping data separate from such actions is essential. Lastly, we urge to consider the type of uncertainty for which a safety stock is held. At Thales, we generally hold safety stocks to deal with procurement lead time variance. For that, we often only need enough stock to work for a few weeks. If it is held for expected defects, we should have just enough (or a bit more) for those expected defects. Together with the safety lead times, when keeping this in mind we should be able to both increase material availability and reduce holding costs.

The most promising directions of future research are to optimise quality inspections from the same quantitative approach, optimising *Fixed Days Supply* (FDS) settings, and a study to quantify tardiness costs such that the model gives better solutions.

Contents

List of abbreviations

AHP	Analytic Hierarchy Process
ASL	Approved Supplier List
BOM	Bill Of Materials
BPA	Blanket Purchase Agreements
CAT	Customer Account Team
CW	Central Warehouse
DEA	Data Envelopment Analysis
DFARS	Defense Federal Acquisition Regulation Supplement
EOQ	Economic Order Quantity
ERP	Enterprise Resource Planning
FAI	First Article Inspection
FAR	Federal Acquisition Regulation
FDS	Fixed Days Supply
ITP	Instruction-To-Proceed
ITR	Inventory Turnover Rate
LSC	Letter Subcontract
MAPE	Mean Absolute Percentage Error
MHC	Material Handling Charge
MM	Mathematical Model
MOQ	Minimum Order Quantity
MPS	Master Production Schedule
MRP	Material Requirements Planning
NRE	Non-Recurring Expenditure
OTD	On-Time Delivery rate
PLM	Product Lifecycle Management
PTO	Purchase-To-Order
RFP	Request For Proposal
S&OP	Sales and Operations Planning
SM	Simulation Model
SMAPE	Symmetric Mean Absolute Percentage Error
TCO	Total Cost of Ownership
TINA	Truth In Negotiations Act
WIP	Work-In-Process

1. Introduction

This thesis describes a new structure of material management to be implemented for a specific product chain of Thales Netherlands, called Aerospace. This product chain faces an outgoing *on-time delivery rate* (OTD) that is below its goal of 98%. Thales wonders what changes are required to improve the outgoing OTD, and meanwhile decrease costs associated with purchasing, holding items and production delay. The specific focus is on material availability. We research several tactical purchasing and inventory decisions which can help in achieving this.

Section 1.1 describes Thales and the specific product chain. The problem statement is given in Section 1.2. Following on that, Section 1.3 gives definitions used in this report related to the problem context. Section 1.4 describes the research opportunities and our research goals. Finally, Section 1.5 introduces the research objective.

1.1. Thales

Thales Group is a multinational high-tech company that creates electrical systems and provides services for the aerospace, defence, transportation and security markets. Thales Group finds its origin in France, where its predecessor was established in 1893. Since then, Thales has grown to a company employing over 80,000 people worldwide and reporting sales of €18.4 billion in 2019. Customers of the global Thales Group are large and well-known, such as armed forces and big cities. The corporate head office of Thales Nederland is situated in Hengelo. Around 1,700 of the 2,100 employees of Thales Nederland work in Hengelo, at either production or the office. It is the worldwide leader in innovative radar technologies and radar systems for naval vessels.

A more specific so-called product chain for Thales Netherlands (hereafter simply referred to as Thales) is named Aerospace. Orders from Aerospace often follow from offset agreements, meaning that when the Dutch government issues orders at foreign military industry firms, the firm should place counter-orders at Dutch firms. Aerospace products incorporate mechanical and electrical parts that are used in attack helicopters, fighter jets and missiles. For those products, Thales typically acts as a contract assembler.

Aerospace significantly differs with radar systems in that the customers are the product owners. They provide a blueprint of an assembly which Thales has to follow, and are the only customer for that product. With radar systems, Thales engineers to order or even builds full radar systems it expects to sell in the future. Additionally, whereas radar systems are hardly sold over five times a year, we find throughputs of dozens up to hundreds each year for Aerospace products. Next to that, the costs and lead time of raw materials are a much larger share of total costs and lead time to the customer at Aerospace products, compared to radar systems. In the remainder of this chapter, we further address named and other differences to explain the impact on supply chain activities. The bottom line is that it requires quite a different way of working and brings more uncertainty.

This research is performed for Aerospace products within the team Master Planning, which is part of the department Industrial Supply Chain (Figure 1.1 shows the organisation structure). It is responsible for both capacity planning and item availability for all manufacturing activities of Thales. The direct supervisor for this project is the Inventory Controller, who is responsible for controlling all inventories at Thales. The subjects of this thesis have much in common with the function Material Manager, which was filled only near the end of this research.

1.2. Problem statement

Thales acknowledges the differences of Aerospace activities compared to those associated with manufacturing radar systems. To determine whether to continue with Aerospace products or to focus on its core business, an internal audit took place last year. Contrary to expectations, the conclusions from the audit were that Aerospace is a profitable part of Thales with many opportunities, and activities should be intensified rather than dropped. Since then, Aerospace has higher priority within the company. At the same time, the manufacturing process and supply chain process within Aerospace were not in control. Aerospace orders faced an outgoing OTD of 30-40%, where orders are considered on-time if delivered at most two weeks earlier or later than agreed to. To improve the process, an end-to-end improvement team was established. Through many revision projects of the process, the OTD has increased towards 80-90%. These projects included among others investigating workload inconsistency, improving the bid process, and adapting the promised customer lead times. Next to that, Thales is currently implementing a *Sales and Operations Planning* (S&OP) cycle, in which (forecasted) demand, supply and operations over a 6 to 36 months period ahead should be discussed and aligned in monthly meetings.

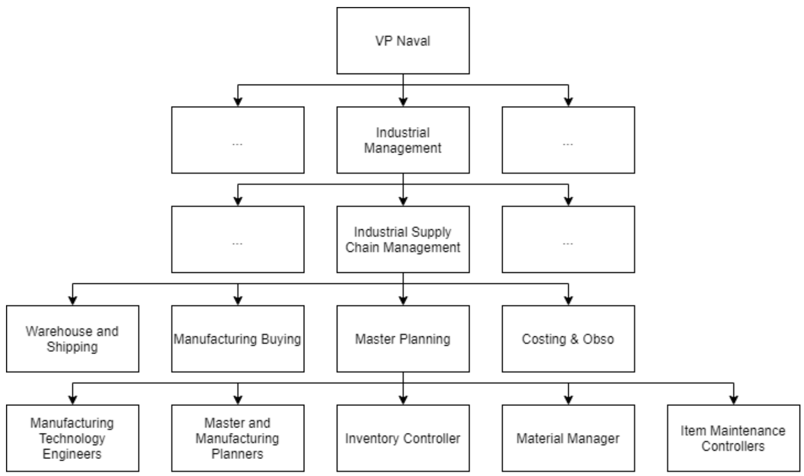

Figure 1.1. Organogram of Industrial Supply Chain.

Despite the positive changes within Aerospace, there is still room for improvement. In order to become more competitive, Thales wants to increase OTD towards its goal of 98%, reduce customer lead times, and lower costs associated with purchasing, holding items and production delays. One approach towards reaching that goal is increasing raw material and component availability. Currently, being a typically order-driven company, Thales wants to keep stock at a minimum and only purchase raw materials and components if there is a sales order. However, as lead times for these items can reach up to two years and show high variation in length, a relatively small delay can already result in production delay. Thales sees an opportunity in researching these and other types of uncertainty. The focus lies on material management. In particular, Thales wonders what tactical decisions in inventory and possibly procurement management should be taken, and what benefits and risks result from a new material management approach. We further consider the research opportunities in Section 1.4.

1.3. Problem context

For the context of the reader, we first have to introduce some definitions that apply to the Thales case. The final assemblies manufactured for the customer are referred to as *products* or *systems*. More specifically, a system consists of multiple final assemblies (products) that are often sold in a certain fixed ratio (e.g. a system that includes 2/1/1 of three different products). As mentioned, many raw materials and components (hereafter named *materials*) needed to manufacture the products are only requested or purchased once an order is accepted. Within materials, we for now consider two types: *buy* and *make* materials. By far, most materials in Aerospace are buy materials, which means they have to be purchased from a supplier. A small part is make materials, which requires Thales to manufacture the material. This is the case for subassemblies created from the buy materials. Note that products may be considered a special case of make materials, made for the customer. Both products and materials are also referred to as *items*.

Another important definition to describe is lead time, as we have different types of lead time to discuss. In the Thales *Enterprise Resource Planning* (ERP) system, both buy and make items have *processing lead times*. This is the time that elapses since consecutively a purchase order (work order) is placed until the items arrive (are manufactured) at Thales. *Pre-processing lead time* is the time it takes to make a purchase order (work order) from the moment the need for the item is stated. Subsequently, the time between material arrival (manufacturing completion), and the item being readily available for manufacturing or the customer, is *post-processing lead time*. As we want to distinct between buy and make items, we consider the sum of these three lead times for buy materials as *procurement lead time*, and the sum for make items as *manufacturing lead time*.

Additionally, manufacturing processing lead time has a fixed and variable component based on the manufacturing batch size. Next to that, we consider *cumulative total lead time* as the total time that elapses between the moment the need for the item is stated, until that item is readily available. For make materials, it then also includes procurement lead times for the needed materials and possibly other manufacturing lead times, while for buy materials this lead time is of course the same as procurement lead time. *Cumulative manufacturing lead time* simply is the sum of all manufacturing processing lead times needed for a make item. Finally, the time it takes to fulfil customer demand from the moment a sales order is placed is *customer lead time*, which in turn is equal to the cumulative total lead time of the product.

Figure 1.2 provides an overview of the types of lead time. Note that the order from procurement lead time to manufacturing lead time is not always strict: Certain manufacturing activities for subassemblies may already start while still waiting for other buy materials.

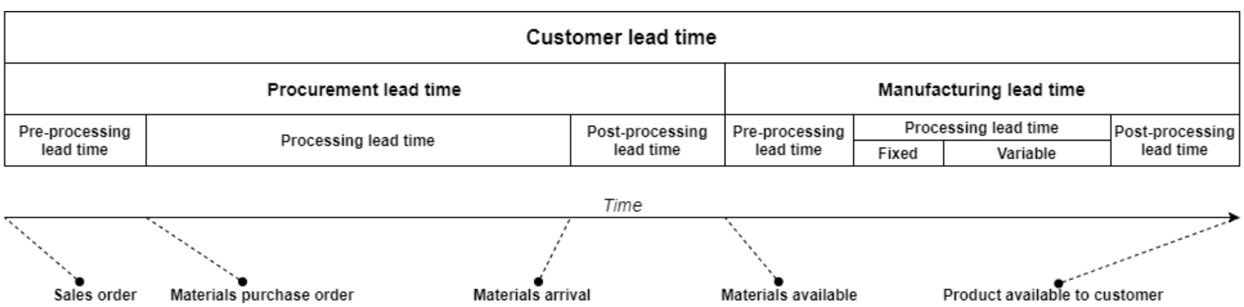

Figure 1.2. Different types of lead time.

1.4 Research opportunities

Dealing with uncertainty is generally the hardest part of material management. We distinguish between demand uncertainty and supply uncertainty. To understand what research opportunities we observe for Aerospace, we first have to elaborate on these two types of uncertainty.

Thales has a *purchase-to-order* (PTO) policy for most materials of Aerospace products. The materials are only purchased once a sales contract is signed or if the customer does an advance payment – in both cases, the quantity and moment of demand are known. For Aerospace specifically, there are several good reasons why this policy is adopted. First of all, Thales is not the product owner in Aerospace. In particular, each product is only sold to one customer who fully controls demand. Our product is a part in the customer's final assembly, and the customer also controls the design of our product. The materials used in Aerospace products are typically not used for other products of Thales. When the customer no longer demands our product (e.g. if they no longer sell their own product, or if they switch to another supplier), all our inventory of that product and its materials become obsolete. If they decide to revise the design, the same applies for the materials that are no longer needed.

A full stop of demand is not the only risk in demand uncertainty. For some materials, high set-up costs and *minimum order quantities* (MOQs) may apply. If we have purchased slightly less than necessary for a next sales order, the total costs may increase when a small but expensive extra order is needed. Next to that, there is uncertainty in the moment of demand. In the case that the customer requires products later than expected, we would have a worse cash position for a longer period if we had already purchased the materials. Also, warranty periods of materials may end before we use them in assemblies. By the PTO policy, we eliminate all of the named risks.

Of course, there are also reasons why we *could* anticipate with purchasing materials on uncertain demand. By placing less orders but in larger order sizes , we may receive higher quantity discounts from suppliers. The customer lead time may drastically decrease, and OTD can increase. Fixed costs of ordering, handling and shipping could decrease with the fewer orders as well. Finally, if materials are expected to increase in price over time, purchasing more now can reduce total costs. Gaining these benefits heavily depends on how well we can predict actual demand and revisions. Unless we can forecast demand with great precision or if we observe very high quantity discounts, the benefits unlikely outweigh the risks.

There are two unknown aspects in supply uncertainty. Materials may be delivered later than expected (uncertainty or variance in procurement lead times), and the quality of materials may be insufficient. In order to cope with variance in procurement lead times, one can do several things. Examples are to keep a safety stock, to order materials earlier (create a safety lead time) for materials, or to build an extra buffer in the customer lead time (such that we have less chance of delivering later than promised). With regards to imperfect quality, we could for instance perform inspections at an order arrival, or have such inspections already performed at the supplier. Generally speaking, we could also switch to another supplier in both cases. However, at Aerospace products we often have a limitation in supplier selection by orders of the customers, so that option is mostly excluded.

1.5. Research objective

Given the request from Thales and the research opportunities, we can capture the objective of this research in the following main research question:

> *For the Aerospace product chain of Thales, how can we use material management to cope with uncertainty in supply and demand, such that on-time deliveries increase, customer lead times decrease, while costs related to purchasing, holding inventory and production delays are minimised?*

At this point, we do not have sufficient information to decide on the solution areas. We require a thorough analysis of the current situation to understand which of the research opportunities (the sketched ones or new) add most value. Only once we have gained that knowledge, we can select the best options and further specify research questions. Therefore, in the following chapter, we clearly define the current situation. We first describe the relevant activities of bid and procurement management. To see how well Thales can predict demand, we evaluate the current S&OP forecasts. We consider which material characteristics play an important role in inventory and purchasing decisions. Furthermore, we list the current inventory policies and observe how well they perform. Finally, we identify which core problems are the focus points for the remainder of this research.

2. Current situation

The analysis of the current situation is the exploratory research on which we build. We first give background information on the sales and procurement processes in Section 2.1. Subsequently, we head into more detail concerning multiple procurement decisions in Section 2.2. Then, in Section 2.3, we describe how current forecasts are made and how they perform. Section 2.4 serves to highlight different material characteristics that may be important in procurement and inventory decisions. After that, we evaluate current inventory policies at Aerospace in Section 2.5. Finally, we decide how to continue the remainder of this research in Section 2.6.

2.1. Process background

In this section, we provide a description of the relevant processes in our research context. These processes include a combination of bid management and procurement management. As it is mainly background information for the reader, we do not head into the details but rather give an overview. Figure 2.1 shows the cross-functional flowchart, which was validated through interviews with the concerned parties. As exact functions and persons are less important, we include two broader areas: the sales and bid functions (the *Customer Account Team*, CAT) and procurement functions (Procurement).

The starting point of bid management is a *request for proposal* (RFP). This means that a customer wants to purchase a certain number of products with all sorts of requirements. Thales is asked to indicate their price, but also the capability of delivering, customer service, involved subcontractors, and so forth. Obviously, preparing such a proposal – or bid – requires a lot of time and manpower. For that reason, senior management evaluates at Gate 1 whether to bid or not. For that evaluation, a capture team is formed who have to develop a bid strategy, profile the customer, prepare potential partnerships and identify potential risks. It is also the stage where the capture team quantifies the probability of winning the bid with an opportunity and risk analysis quantification tool, which is explained more thoroughly in Section 2.3.1. If Thales decides that it should compete for the order, the bid is framed and prepared. That means that the bid organisation is planned, and all information that the customer requires must be gathered.

For Procurement, this is the moment to start. The Procurement Project Manager records the need and initialises the procurement process. Together with the buyers, an acquisition strategy is proposed. It may occur that new suppliers have to be identified. Each customer has different rules for Thales' suppliers to be approved. One customer may have general requirements that suppliers must meet, while another will have an *approved supplier list* (ASL) with specific suppliers that may provide materials. Therefore, this process varies heavily in terms of time and parties involved. Once the acquisition strategy is clear, buyers send their RFP files to suppliers and receive proposals. This is mostly done for about 20% of the suppliers that deliver the materials worth 80% of the costs. For cheaper materials, an estimate for the price is given by the buyers based on prices from the past. This often includes a risk mark-up for an expected price increase since then. Before proposals of suppliers are recorded, negotiations with suppliers should take place. In practice this step is often skipped at this point, due to a lack of time.

Information from Procurement is used as input for the price and cash sheet, which basically presents all costs and the price for the customer, and is part of the bid. This bid may be reviewed, after which the bid is sent to the customer and negotiations may start. Naturally, the order may be denied by the customer during this time. Also, an outcome of the negotiations could be that material costs are considered too high, which means that Procurement must (re)negotiate with the suppliers. Additionally, while an order will be placed but negotiations on the price are still ongoing, the customer may already give an *instruction-to-proceed* (ITP) or *letter subcontract* (LSC). This means that the customer does an advance payment such that Thales can already purchase crucial materials, or do an advance payment in their turn to allow suppliers to purchase their own crucial materials – depending on the sum of the payment. Eventually, both parties may come to an agreement, after which materials are purchased, monitored and controlled, and finally received.

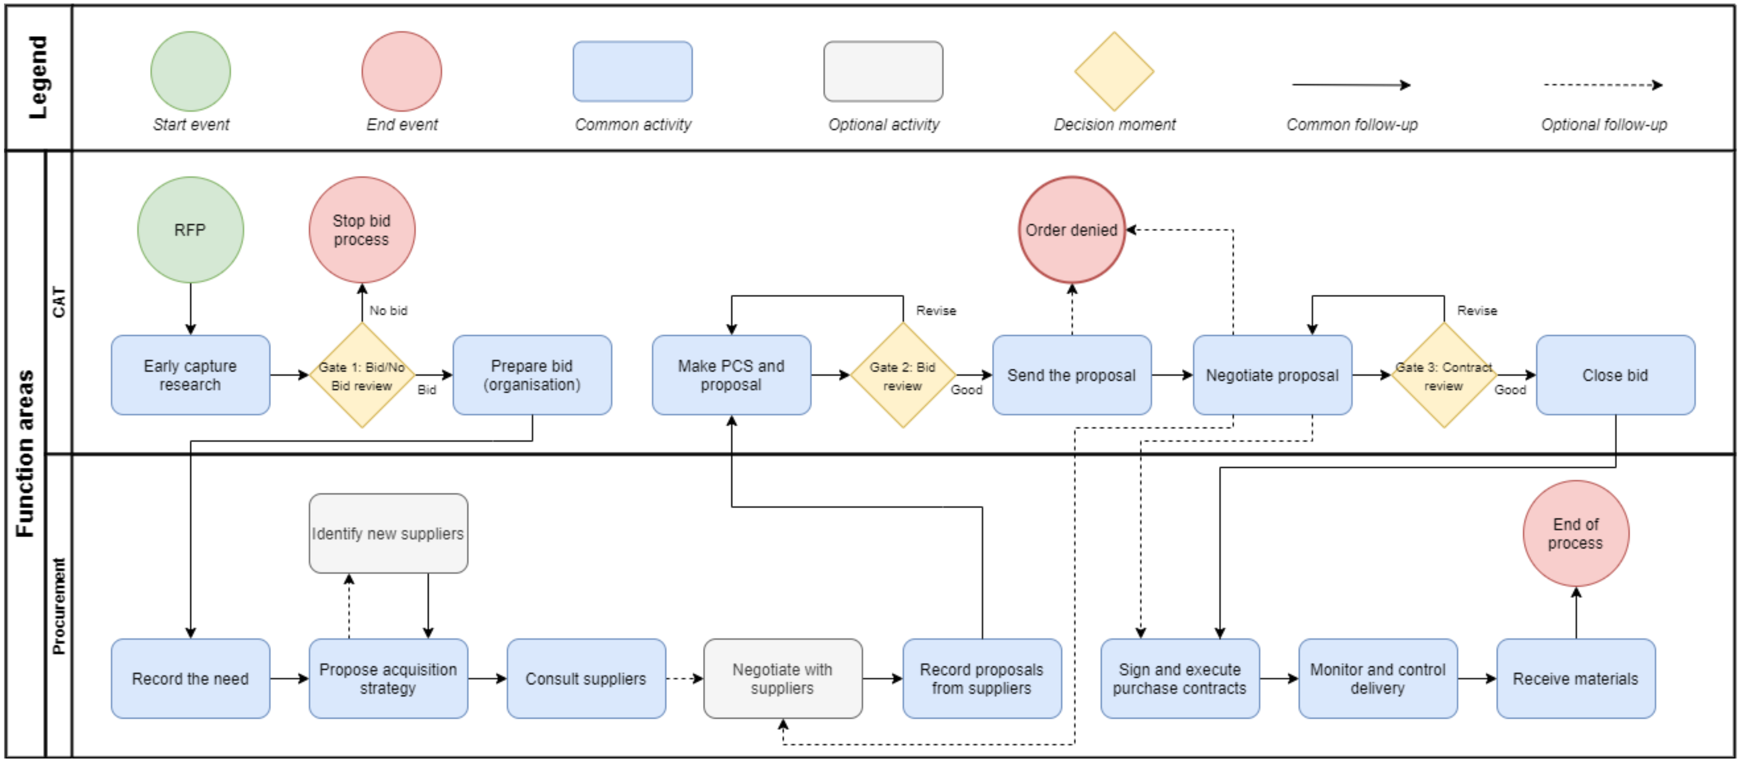

Figure 2.1. Cross-functional flowchart of bid and procurement management.

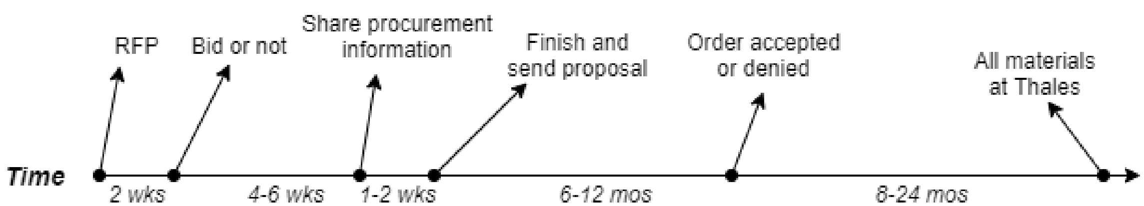

Figure 2.2. Timeline of the bid and procurement processes.

For known products and customers, the CAT aims to send a proposal within 8 weeks after receiving an RFP. In 2 weeks, the decision to bid or not should be made. From that point, Procurement generally needs about 4 to 6 weeks to give relevant information to the CAT. In the remaining time, the proposal should be finished and sent. Once the customer receives the proposal, it often takes about 6 to 12 months until an order is accepted or denied, and a contract signed. At that point, it takes days to at most a few weeks to embed the bid in the ERP system and execute purchase orders at suppliers. Finally, material lead times may take up to 2 years. Figure 2.2 summarises these periods.

It is important to note that the process durations may vary heavily depending on the type of product and customer. There are exceptions which we do not show in the figure. First, the *Federal Acquisition Regulation* (FAR) and *Defense Federal Acquisition Regulation Supplement* (DFARS) of the United States obligates so-called fair pricing from single source suppliers. These regulations include the Truthful Cost or Pricing Data Act threshold, still often referred to as TINA threshold from its former meaning of *Truth in Negotiations Act*. If an applicable order is over the TINA threshold of $2 million, the supplier must give full insight in all costs. Underlying purchase orders of Thales may also surpass that threshold. Gathering all the necessary information may add months to the process.

Next to that, identifying new suppliers for a customer that has an ASL could add weeks or months to the procurement process. On the other hand, repeated orders for more simple and often mechanical products could be fully handled in mere months. Another important issue is the validity of proposals. The proposals of suppliers are only valid for a certain period, ranging from a few months to multiple years. Therefore Thales generally sets the validity of its own proposals for a shorter period ahead than those of the suppliers. If a customer has not yet accepted the proposal within this validity period, Thales might have to renew its proposal. The long period between sending a proposal and having a signed contract, in combination with long material lead times, underlines the value of an ITP. However, in practice it even occurred that most of the products were already delivered to the customer based on advance payments, while price negotiations were still ongoing hence a contract not yet signed. Naturally, that is also undesirable. To conclude, we observe activities with many parties involved, long and varying periods of waiting with uncertainty, which emphasises the urge for a more efficient process.

2.2. Procurement management

We describe in Section 2.1 that the purchase contracts may be executed once demand is known. However, that does not mean that all materials are actually ordered at that moment. To explain how the order moment is currently determined, we have to introduce several concepts. We first explain how assembly activities are structured to see when materials are exactly required. Then, we describe when purchase orders are placed, and how it influences actual delivery. Subsequently, we analyse how Thales deals with lead time uncertainty from a procurement perspective. Finally, we take a first look at delivery delay data.

2.2.1. Manufacturing activities

A sales order of Aerospace products is typically delivered to the customer in multiple batches. Depending on the size of the sales order, the batches may be delivered over up to multiple years. These batches generally follow a certain batch cycle, e.g. 10 units per week or 12 units per 4 weeks. Within these batches, a product may again be produced in batches or one-by-one. Underlying make materials are mostly produced in batches. These make materials are then either stored until they are needed in the next assembly stage or manufactured just-in-time for the next stage. It also means that some materials are required earlier than others, depending on when the underlying make materials are planned to be manufactured. Figure 2.3 shows an example of the assembly system, where red blocks are make items and green blocks buy materials. Aerospace products have 2 to 4 assembly levels, with the product included. That is disregarding further levels within purchased components. Thales plans to infinite capacity, which means that planning ignores resource constraints (e.g. people and tools).

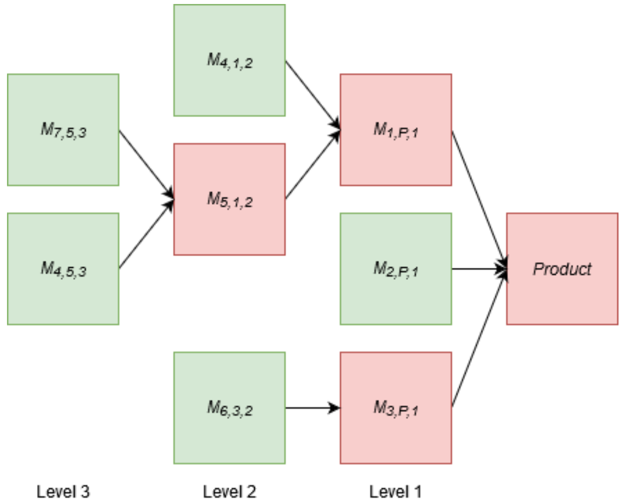

Figure 2.3. Assembly system ($M_{i,j,k}$ means material i, used in item j, at assembly level k).

Make materials in Aerospace are always used in only one other make item. Buy materials may be used in multiple make items, as is shown for material 4 at both assembly levels 2 and 3. As manufacturing lead times are rather short, the total need of a buy material at all stages together is often requested for the first stage in which it is required already. For the production of any make item, all different materials and mostly their full batch quantity must be available before production may start (the *production start date*). Only if there are at least some of all different materials and starting production already can significantly reduce the total delay, the production may start already. Delay in the production of one make material needs not necessarily delay next stages depending on how it is scheduled, but it often does.

Obviously, what we describe is how processes should go in general. In practice, all kind of disruptions may change it. There may be capacity issues, changing delivery dates from the customer, material quality problems with possible rework, and so forth. For any of those reasons we may have to expedite or postpone the production start date. Unless production postponing was known far in advance, the communicated buy material need with suppliers is generally not delayed as well. As a result, we also have those materials on stock for a longer time.

2.2.2. Purchase order management

Once an order is accepted, the production start date of the first batch depends on the longest procurement processing lead time of the underlying buy materials. When there is a go for purchase orders (*purchase kick-off moment*), which is typically marked with signing a contract or receiving an ITP with advance payment, we again check the longest procurement processing lead time. The reason for this is that named lead time in the system is what the supplier has indicated perhaps a long time ago, and it may not be what they can actually promise anymore. Based on the newly given lead time, Thales determines the actual production start date.

Materials are planned to arrive just-in-time. Bearing the production start date and the procurement post-processing lead time in mind, the moment a material should arrive (the *need-by date*) is determined. It is generally also the date for which we request the material (the *request date*). At the moment an order is placed at the supplier, they give a *promised arrival date*, which will not necessarily match the request date. Finally, the moment the shipment arrives is the *arrival date*. Figure 2.4 illustrates a timeline with the different types of dates. As can be seen, in case 1 the supplier promised to deliver the materials just before they were needed, but actually delivered after the production start date. In the second case, the supplier delivered earlier than promised, but still much later than needed. We observe that in practice the promised arrival date ranges between slightly before and much later than the request date, and the arrival date is often a bit earlier than the promised arrival date but with much higher variance when later. We further discuss this in Section 2.2.4.

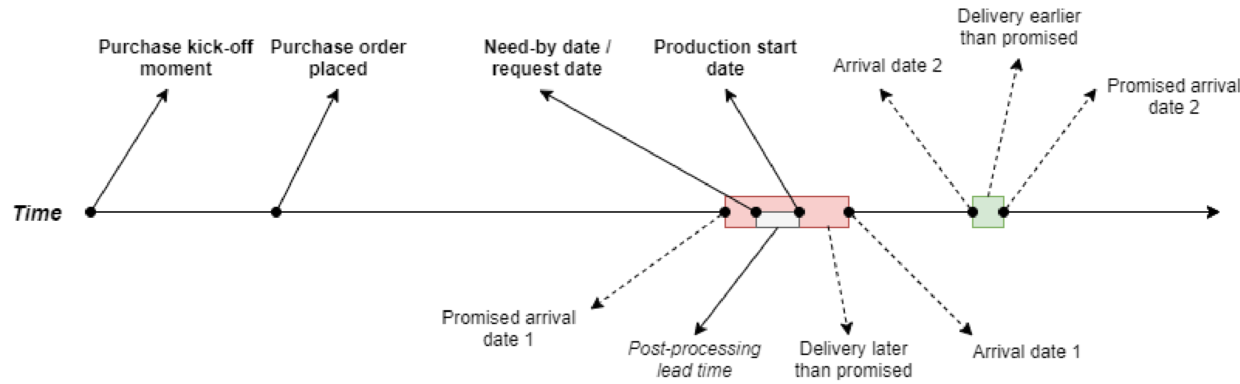

Figure 2.4. Example timeline from purchase order until receival.

Having materials either early or late is both undesirable. Materials are generally paid for after delivery, often 1-2 months later. Only in some cases of expensive materials, Thales may do an advance payment such that the supplier can cover part of his expenses already. In both cases, an earlier arrival lowers the cash position and results in inventory holding costs. On the other side, if materials are late then production cannot start obviously, resulting in tardiness costs and holding costs for materials that cannot yet be used. The results are less severe if the lateness is in the promised arrival date, because we can then anticipate on earlier notice. Still, even then, it hardly occurs that all purchase orders are rescheduled if one of them is known to be delayed.

A typical reason why the promised arrival date is later than requested is again that the registered procurement processing lead time in the ERP is outdated. The supplier cannot promise to deliver in that timeframe anymore, for instance if he received a large order from another customer. We may also have ordered too late but still give an early request date, in order to try and speed up the delivery. Or we simply place rush orders, with higher costs included. An order being placed too late may for example be caused by specific information being not yet clear or parameters that are incorrect (due to which the ERP system does not create the purchase order need). The arrival date being later than the promised arrival date is generally caused by the supplier, although transportation may also play its part (see next section).

The moment at which materials are actually ordered depends on the type of scheduling. We distinguish between *forward scheduling* and *backward scheduling*. With backward scheduling, the ERP system checks when the material is needed, what the total procurement lead time is, and then gives a notification to the concerned buyer that the material should be purchased. With forward scheduling, we do not wait with placing an order. Instead we order it once demand is known, and will also receive it earlier than truly necessary (except for the materials with the longest procurement lead time). This reduces the risk of materials being late, but it also increases holding costs. For some specific materials at Thales this forward planning approach is adopted, but we more often still see backward scheduling with possibly outdated procurement processing lead times. When a purchase order is placed for a first batch of materials, the later batches are already ordered as well.

Another very important factor in purchase orders is the *Fixed Days Supply* (FDS). All materials at Thales currently have 30 workdays FDS. This means that if a purchase order is requested in the ERP system, it automatically includes demand for the next 30 workdays as well. By combining this into one order, Thales cuts on fixed costs of ordering and has a larger chance of receiving the materials for the later demand moment on-time. However, by combining demand of multiple batches it also increases holding costs. In a specific extreme case, we see that the required number of a material for one batch costs over €500K, and due to FDS it is always ordered per 2 batches making it more than €1M per purchase order. This vastly increases holding costs. Furthermore, after total demand is evaluated in the FDS horizon, new demand in that period does not automatically create purchase order need. It therefore is more difficult to work with higher FDS settings in an unstable planning, as Thales should regularly check if new demand was added. A final reason why materials may be ordered in higher quantities is due to MOQs given by the supplier. This of course cannot be influenced.

What we describe in this section does not apply to all materials. A material can be included in the *Material Requirements Planning* (MRP). This means that its total demand is evaluated based on the projects (sales orders) in the system minus current stocks. Given that demand, it shows when and how much to purchase. If a material is excluded from the MRP, this evaluation does not take place. Instead, materials excluded from the MRP have a min/max policy. We further explain this difference in Section 2.5.

2.2.3. Dealing with lead time uncertainty

In order to cope with uncertainty in procurement lead time from a procurement perspective, a safety lead time is added to the actual post-processing lead time. By increasing this lead time, the ERP system will notify the buyer earlier to execute the purchase. Appendix A shows agreements that were made on these lead times. Recall from Section 1.3 that the procurement processing lead time is the actual time between placing the purchase order and arrival at Thales. Depending on the agreement with the supplier, this can also include transportation from the supplier to Thales. The procurement post-processing lead time is supposed to be the time needed between arrival of materials, picking the materials and possibly inspection. But as can be seen, the safety lead time is also added to that value.

In practice, we indeed see post-processing lead times of 17 and 27 workdays, but also some differing values. An on-going project is to change procurement post-processing lead times of certain materials that were often delivered too late, although this has only been done to a small part of the materials. Also, changes may have been made in the past for reasons that are no longer valid. We furthermore observe more extreme cases, in particular a material with 5 weeks actual processing lead time and half a year post-processing lead time. Next to that, quite a few materials have post-processing lead times of 7 workdays, so these have at most 5 workdays safety lead time. A major issue of including the safety lead time in the post-processing lead time is that it is no longer clear whether the material actually is required earlier, for instance to be inspected. It can also lead to the problem where a buyer may speed up delivery at an expected delivery delay – possibly at higher expenses – while the material would have still been on-time considering the actual material need. Preferably, the request date is expedited while the real need-by date still remains clear.

Although these structure issues deserve attention, a larger problem is that the safety lead times are still rather fixed. Each supplier is different, different materials within suppliers may be more difficult to deliver on-time, and transportation lateness may also play a role. At the same time, as material costs range from less than a cent to over €20,000 per produced product, holding materials longer than needed results in very different costs for each material. With some basic knowledge of statistics, we can already say that having quite fixed values of safety lead times for all materials will be far from optimal regarding the sum of holding costs and production delay costs. This is a key problem, especially regarding supply uncertainty with a PTO policy. We therefore conclude that this would be very interesting to investigate.

2.2.4. Delivery lateness data

Thales evaluates all shipments in its so-called *vendor ratings*. By comparing when shipments arrived on the one hand, with either when the suppliers promised the shipment or a revised date from Thales' side, we can check the OTD of each supplier. Next to that, it is also used to see how many materials were rejected originating from that supplier. All of these data are mainly used to assess how well a supplier is performing and to check if follow-ups are necessary. However, the source data of vendor ratings could also be very useful when considering safety lead times. By using the same data that are used to determine whether a supplier delivered on-time, we can determine exactly how many workdays the delivery was early or late. We consider this *delivery lateness*, where a positive number equals a late delivery and a negative number is an early delivery. With these data, we can try to fit a probability distribution on the delivery lateness later.

As the vendor ratings have changed in structure from 2019 compared to earlier years, we use the data from 2019 until now. Some data have to be removed .These are order lines that are still open and duplicate lines. In addition, some data may have been filled in incorrectly, or there were other issues for which the delay is not correct. These lines are marked manually by the buyers. As we cannot be certain which of these are actually correct or not, we also remove these lines from our analysis. After these removals, there are 6,872 shipments lines left to analyse. Figure 2.5 shows the histogram of delivery lateness, cropped from 30 workdays early to 40 workdays late (containing 98% of the records).

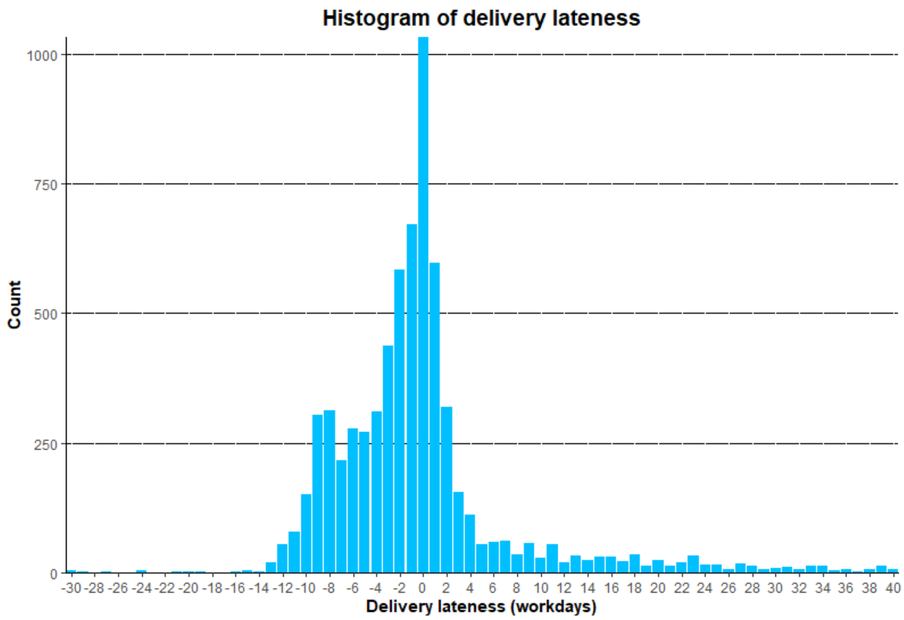

Figure 2.5. Cropped histogram of delivery lateness.

Although nearly 7,000 records is a good amount for possible further analysis, we do have some remarks. Only 244 of these records concerned materials that are used in Aerospace products which we include in our research (we elaborate on these products in Section 2.3.2). More specifically, these are 122 unique buy materials (~25% of unique materials we now analyse). Out of the 6,872 records, a total of 2,494 records concerns suppliers that also supply Aerospace materials. However, these are 22 unique suppliers (~54 of unique suppliers we now analyse), of which only 7 suppliers have more than 25 order lines. In other words, for most of the materials and suppliers we have very limited date. With further analysis, we should consider how to deal with that.

2.3. Forecasting

We elaborate on the role of bid management in Section 2.1. Within the process, just before Gate 1 (see Figure 2.1), Thales quantifies the probability of winning the bid with an opportunity and risk quantification tool. This is input for the S&OP forecasts. This section describes named tool, the S&OP forecasts that are based on the probabilities resulting from the tool, and how these forecasts perform.

2.3.1. Opportunity and risk analysis

In order to make an informed decision at Gate 1, Thales has an opportunity and risk quantification tool. It consists of twenty questions or areas, where at each area the potential order is scored on a scale from 1 to 5.

Examples of these areas are to what extent additional orders might follow or the degree to which the product has been made before. Based on the scores on each area, the order is placed on an opportunity and risk matrix as shown in Figure 2.6. This aids the capture team to decide whether to bid or not, and is used in the discussion on the decision. Generally speaking, an order above the upper orange line should be bid on, one below the lower orange line should not, and more discussion is necessary for an order between both orange lines.

In addition to depicting the risk matrix for the bid decision, the tool also indicates the probability of receiving the order. The probability is constructed as follows: the probability of the customer placing the order *at all* (P1), the probability of Thales receiving the order *given it is actualised* (P2), and the product of both (P3) being the actual probability of receiving the order. The first is determined on two questions: will the customer be able to finance the order, and how strong is their need for the order? For the second probability, Thales evaluates the customer's perception towards Thales and the extent to which Thales stands out to competitors. The probabilities for P1 and P2 may vary between 10% and 90%, resulting in a P3 value between 1% and 81%. The three probabilities are also adjusted later on at Gate 2. Next to that, Thales might already register a prospective order prior to receiving an RFP, if it believes a customer would send an RFP in the future. This opportunity will then usually have very low P-values, as there is more uncertainty over a longer period.

Figure 2.6. Opportunity and risk matrix.

Each month, Thales' management team discusses the *scenario list*, which consists of all sales orders and promising prospective sales orders for the coming years. Every three months, a *scenario freeze list* is created, based on the same orders of the next four years. This is input for the S&OP forecasts, in which the Demand Manager for each quarter ahead provides an indication of the demand per product (in terms of products shipped) and the expected workload. In order to get realistic demand numbers, he validates and adjusts these numbers based on information from Supply Chain Managers and Project Managers. In these forecasts, we have different forecast levels. Accepted sales orders are either *work-in-process* (WIP) or *accepted* (but not yet started). Prospective sales orders with an estimated P3 value higher than 70% are labelled as *commit*. Finally, prospective sales orders with a P3 value between 30% and 70% are labelled as *risk*. These numbers of P3 values are guidelines for the labels though, not strict separators. Figure 2.7 illustrates the format of a possible forecast made in Q3 of 2015.

2.3.2. S&OP forecast performance

As we noted earlier, demand forecasts have to be very precise for the benefits of ordering on uncertain demand to outweigh the risks. That is why the PTO policy is used in general. In Section 2.1, we already see that there is much uncertainty in the moment of demand. The exact quantity is also only known in a later stage. In this section, we quantify how well we can actually predict demand. We also have to decide which products to consider in the analysis. For these products, sufficient data should be available. At least some forecasts must have been made and the product must have been sold over the last five years. Based on those criteria, we evaluate six systems, with fourteen underlying products.

	15Q4	16Q1	16Q2	16Q3	16Q4	17Q1	17Q2	17Q3	17Q4	18Q1	18Q2	18Q3	18Q4	19Q1	19Q2	19Q3
WIP	12	16	9	10												
Accepted																
Commit						24	24	24	16							
Risk												50	50			

Figure 2.7. Possible forecast made in Q3 2015 (numbers are units).

The S&OP forecasts predict the number of products shipped for each quarter four years ahead. This means that we have to compare them with actual sales order lines. We do note that this does not align with the moment materials are needed, but the actual material need could be traced back through expert opinion. While forecasts are made at a system level, we have sales order lines of the products for that system. Next to that, we aggregate predictions per year. We want to know whether forecasts perform well in the bigger picture, but are less interested in fluctuations from quarter to quarter. This means that for each quarter where a forecast is made, we look at the total forecasted numbers per year instead of forecasts per quarter. For example, in the case of Figure 2.7, we have the forecast made in Q3 of 2015, and evaluate the forecasted demand of 35 in 2016, 88 in 2017, 100 in 2018 and 0 in 2019. The forecasts we take into accounts are those made in each quarter of 2015 to 2018, for the predicted total sales per year of the years 2016 to 2019. Forecasts of the fourth quarter of 2015 and 2017 are missing, so we interpolated these values.

To measure performance, we have to consider two factors. First, we look at the different forecast levels. We separate between risk, commit, both combined, and a total of both including already accepted orders (WIP and accepted). The latter is important as risk and commit forecasts will at some point either become accepted or cancelled orders, and therefore on themselves may be misleading on the shorter term. Second, we must consider how well forecasts perform based on the forecast period. These periods range from 1 quarter ahead (e.g. forecast of Q4 2017 for 2018) to 16 periods ahead (forecast of Q1 2015 for 2019). As a result, we only have one observation per product of a 13-16 quarters ahead forecast period: Those made in each quarter of 2015 for the year 2019. On the other hand, we have 4 observations per product of a 1-4 quarters ahead forecast period: All forecasts made in any quarter for the next year. Obviously, we have 3 observations for 5-8 quarters ahead, and 2 observations for 9-12 quarters ahead. For that reason, we consider performance evaluation on the short term more reliable.

Currently, Thales determines the *Mean Absolute Percentage Error* (MAPE) and the *Symmetric Mean Absolute Percentage Error* (SMAPE) to measure forecast accuracy. Equations 2.1 and 2.2 describe these measures, where F_t and A_t are respectively the forecasted and actual amount for period t. However, these measures take the absolute value of errors. We are also interested to know whether we typically under- or overpredict, as they result in different risks (see Section 1.4). In order to measure under- or overprediction, we use the relative bias. Mostly, relative bias is measured as the mean of all errors divided by the actual amounts. In our case, we often see that actual amounts are zero as no products were shipped in that period. As we cannot divide by zero, we instead use the sum of errors divided by the sum of absolute errors as relative bias, as shown in Equation 2.3. This measure scores between -1 and 1, where a positive value indicates underprediction and a negative value overprediction.

As an example, Figure 2.8 shows the forecast and actual amount (black line), for all products, all forecast periods and the actual year 2019. Although it is not very clear, the forecasted amount is often below the actual amount for 2019. However, that is certainly not the case for the other years (shown in Appendix B.1). For the grand total of all materials, forecasting levels and forecasting periods, we find a relative bias of 0.348. This means that the forecasts in general are underpredicting. Appendix B.2 shows the relative bias per forecasted period, for the different products and systems.

$$\text{MAPE} = \frac{1}{n}\sum_{t=1}^{n}\left|\frac{A_t - F_t}{A_t}\right|, \qquad (2.1) \qquad \text{SMAPE} = \frac{1}{n}\sum_{t=1}^{n}\frac{|F_t - A_t|}{|A_t| + |F_t|}, \qquad (2.2)$$

$$\text{Relative bias} = \frac{\sum_{t=1}^{n} A_t - F_t}{\sum_{t=1}^{n}|A_t - F_t|}, \qquad (2.3)$$

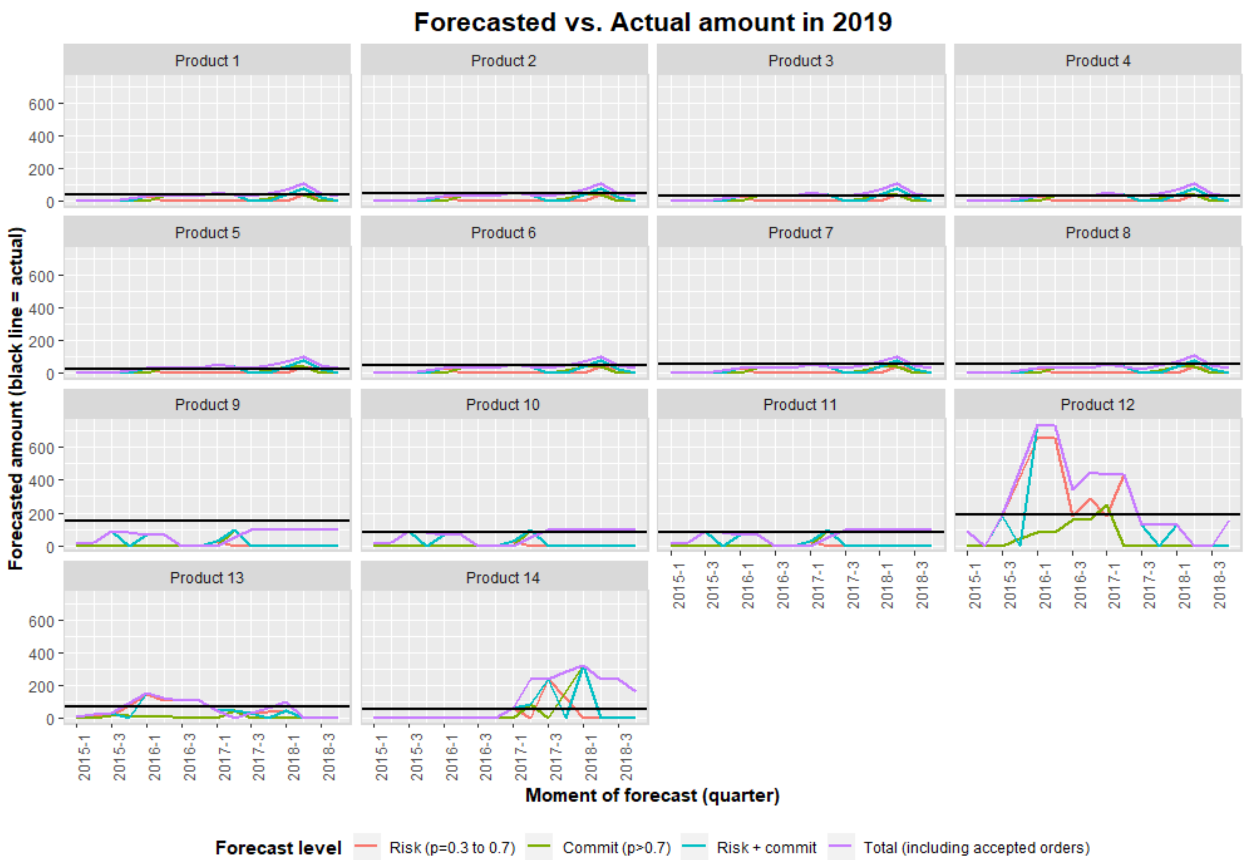

Figure 2.8. Forecasted vs. actual amount in 2019, per product and forecast level. Except for Product 12, the forecasts are underpredicting for 2019.

We do not need further quantifications to observe that the forecasts have too high errors. From neither a certain forecast level or a forecast period, we can draw reliable information. At times we may predict zero sales while actually having a lot, and also the other way around. In addition, we see that there are differences between actual sales of products within a system. That is a big issue, as we then need materials for different products in a different proportion. What makes these forecasts even more difficult to work with, is that they predict when products will be shipped (to determine when manufacturing activities have to take place). We cannot trace back the total (prospective) sales order quantity or expected order acceptance moment, especially because forecasted amounts are often shifting over time due to rescheduling. For example, if it turns out a forecast for 2017 was underpredicting, it could be that much was rescheduled to be shipped in 2018. The total sales order quantity may still have been correct, but we cannot measure it through these forecasts. For the same reason, quantitative forecasting with these data seems also less useful as it will also show expected sales order lines over a period instead of an expected prospective sales order size.

Altogether, the forecasts may be useful from a broader financial perspective or with regards to total capacity planning, but on a material level they are not precise enough. It is not our main focus to improve forecasts, nor is it desirable to work with different forecasts throughout Thales. We conclude that improvements to the forecast must be made before using them in purchasing decisions with uncertain demand.

2.4. Material characteristics

This section describes multiple material characteristics. We aim to do the following by that. First, we list material characteristics and assess their impact on purchasing and inventory decisions. Next, we determine whether they should possibly play a role in the remainder of this research. Furthermore, we take a next step by analysing how these characteristics should then be parametrised. In other words, we already look at data transformations or calculations that may be interesting.

To analyse the characteristics, we have to gather data from multiple sources. These sources include the ERP system, the *Product Lifecycle Management* (PLM) system, and separate Excel files and documents. Some characteristics can be found as is in a source, while others have to be calculated based on other numbers. The chosen characteristics to analyse are based on own insight and information from colleagues, and are as follows:

- Slack to critical path (explained in Section 2.4.1)
- Material costs
- Revision risk
- Shrinkage
- Warranty periods
- Commonality
- Salvage value

2.4.1. Slack to the critical path

Recall from Section 1.3 the different types of lead time. The customer lead time depends on manufacturing lead time and procurement lead time. These procurement lead times are essential in our research. However, it is hard to compare the procurement lead times of materials amongst different assemblies. A procurement lead time of 20 weeks may be the longest procurement lead time in one assembly, while a procurement lead time of 40 weeks is not in another. Therefore, we create a new scale that allows for a better comparison, and also more clearly shows which materials are most critical.

By analysing the *Bill of Materials* (BOM) in the ERP system, we can observe which materials (buy or make) are required for certain make items and in which order these are made. By manually combining the BOM with all lead times from the ERP, we can perform a critical path analysis. A buy material is on the critical path if it is the material that has to be ordered earliest for a whole product, taking into account both procurement lead times and manufacturing lead times of all make materials. More specifically, we say that the material has no *slack to the critical path*.

If we know the critical path, we can also determine the slack to the critical path of other buy materials. To be clear, a buy material with e.g. 6 weeks slack to the critical path may be ordered 6 weeks later than the first material (disregarding delivery lateness). Figure 2.9 visualises this characteristic, and Appendix C gives the mathematical notation. Slack to the critical path is more useful than procurement lead time on its own, as it puts procurement lead times in perspective to other materials of the product. Holding inventory for materials with no slack directly reduces customer lead times. At the same time, these are also the materials at which increasing safety lead time directly increases customer lead time. Therefore, this measure helps a lot when identifying *critical materials*. By critical materials we mean the materials that we expect most problems with (e.g. much delivery lateness or defects) combined with how well we can anticipate on those problems. From now on, we often refer to slack to the critical path with just *slack*.

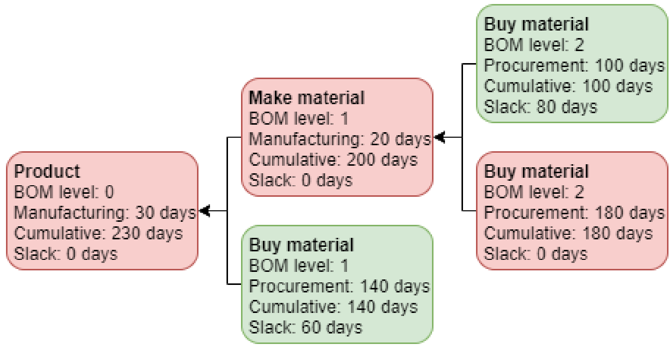

Figure 2.9. Slack to critical path explained.

2.4.2. Material costs

Material costs are definitely important to include in any decisions, be it considering either supply or demand uncertainty. Activities should not be delayed due to a lack of very cheap materials, while at the same time the expensive materials should not be stocked in abundance. There is a lot involved with the costs of materials. However, it would not be worthwhile to thoroughly describe all the different facets. We choose to further elaborate only on the fraction that is directly important to possible decisions in this research.

During the procurement phase of bid preparation (see Figure 2.1), we make agreements for the more expensive materials. We often also have *Blanket Purchase Agreements* (BPA) with suppliers, which basically indicates predetermined prices given a certain quantity ordered. That means that we have *all-unit quantity discounts*, where the price of all units decreases once a certain quantity threshold is ordered. This discount is based on the fact that the supplier also has less costs, e.g. fewer set-up costs or material discount in his turn. These discounts differ heavily, from no discounts to over 80% if a sufficiently large order is placed. Especially at mechanical parts where the set-up time is relatively long, we see high discounts. While these quantity discounts decrease the purchasing costs, the holding costs are not always taken into account with these decisions. This may lead to the point where quantity discounts actually do not outweigh these holding costs.

Next to that, we also see fixed costs of ordering. The majority of the fixed costs are at Thales generally referred to as *non-recurring expenditure* (NRE). These include for example a mould that the supplier has to manufacture the first time a material is produced. Furthermore, it also includes inspection costs. There is one important type of quality inspection that is mainly applicable for some more expensive and larger materials, namely the *First Article Inspection* (FAI). This is a standardised inspection in the aerospace and defence business, which should be conducted in any of the following cases:

- First production run of a new material
- Changes in process, design, tools, location, used materials, or sources
- More than 2 years since last production of the material

Conducting an FAI involves extra costs and possibly extra procurement lead time. For that reason, such an inspection may be waived by our customer, if our supplier has proven to be very reliable. The costs of FAI relative to total material costs differ, from being just a minor part of the total costs up to a significant amount. Therefore, it could be interesting to include it as a separate parameter in purchasing decisions on uncertain demand, where purchases between demand could result in lower total costs. However, we do note that we do not have clear information on possible changes in process etc., and we cannot easily determine the period since last production at the supplier. Therefore, we most likely exclude it as a separate issue in this research.

Fixed costs of ordering that are made indirectly by Thales and which are not included in the NRE are determined through the *Material Handling Charge* (MHC). It is a fixed percentage of buy material unit price that is determined yearly based on expect purchase volumes for the next year. It depends on four rates: purchase rate, inventory rate, accounts payable rate, and aging rate. The first three rates are determined on personnel or working hours indirectly needed for in this case Aerospace. Purchase rate is about personnel of the procurement department, inventory rate for personnel of the warehouse, and accounts payable for work hours of the finance department. Finally, aging rate discounts for dead stock. For Aerospace, this rate is 0%, as materials are only purchased for sales orders and the customer pays for all stock that we purchase for them.

If we would aim to reduce fixed costs of ordering, using the MHC as an indication would be wrong in our opinion. Consider that two shipments of 50 materials with €100 material costs incur the same MHC as one shipment of 100 materials with €200 material costs (as the MHC is fixed). In reality, shipment costs would probably be lower in the second case. Also the other way around, if we get high quantity discounts we pay less MHC because of the lower unit price, but handling and shipment costs do not decrease (much). Determining actual shipment and handling costs is very difficult however, as these costs are not directly relatable to certain material purchases. In addition, the materials are often small and demand is still relatively low, even when possibly increasing lot sizes based on future demand. Therefore, concerned experts believe very little is to be gained on this point. We agree with that, and therefore exclude it from further analysis.

Furthermore, we also consider price speculation. If we would order materials based on forecasts for up to four years, we may have to include possible price changes. An expected increase of prices of e.g. 5% per year does mitigate some of the holding costs that result from higher order quantities now. On the other hand, prices could also decrease, as we often see with electrical consumer products. We discussed these possibilities, and heard that trying to predict this or even find a trend is not feasible. However, a risk mark-up for expected price increase is now already added in case we do not know exact material costs (see Section 2.1). We suppose that these are more or less correct assumptions. We therefore believe that this parameter should be included in decisions of purchasing for a longer period of demand.

In our calculations, we use the material costs from the ERP system. This is based on actual prices from the past. Furthermore, based on the BOM of which a material is part of, we calculate what the total material costs per product are given the required total quantity of that material in a product. This tells us more than just the costs of one unit material. We consider these costs the material costs for further use. Just as with slack to the critical path, it means that we may have different material costs for a certain material, depending on the product we are evaluating. Next to that, another key take-away is the cost reduction offered by quantity discounts, but that holding costs should also be taken into account. Even with known demand, these discounts can improve our situation if used correctly.

2.4.3. Commonality

Another characteristic that we analyse is the shared usage of materials, or commonality. By commonality we mean: To what extent are certain materials that are used in the assembly of one product, also used in other assemblies? It is an important factor for both the analysis of safety lead times per material, as well as in procurement and inventory decisions. As described earlier, Aerospace products are made specifically to the wishes of a certain customer. If the customer no longer requires the product from Thales, it becomes obsolete and also the materials needed specifically for that product. If materials are also used for other Thales products, the value of these materials would remain significantly higher once such a product is no longer required. In addition, pre-assembling subassemblies used for multiple products in times of excess capacity may be more valuable.

To analyse the commonality of a certain material, we distinct three categories. A material may be used in another product in the same system (if applicable). These commonalities are least valuable, since orders are mostly placed for all products in a system. Second, a material may be used in a product from another system within Aerospace. Third, a material may be used in another product outside Aerospace. Naturally, combinations of these may also exist. We analyse commonality by joining all product BOMs from the ERP system and counting multiple occurrences in Excel.

Of the 488 unique materials used in the products we analyse, we find 32 materials that have commonality within the system. These all are within two systems, of which we know they are almost always sold as one system with all the underlying products. This was confirmed during analysis of sales orders for forecasting. Within other systems of Aerospace, we only find 8 common materials. These materials cost on average €0.70 per product, and slack to their critical path exceeds well beyond 40 weeks. Finally, only a type of washer and a label are used in products outside Aerospace as well. All materials with commonality are buy materials, so there is no extra value in pre-assembling. As there is hardly any commonality in Aerospace, we believe there is very little added value to include it in any further analysis. Hence, we consider it out of scope for Aerospace. We do recognize that it could be important to take into account outside Aerospace.

2.4.4. Shrinkage

Another important characteristic is shrinkage or attrition. By that, we mean that materials may break in the assembly process, get lost, or any other reason why the yield is not 100%. Shrinkage can cause serious delay at production. If we only see defects when the material is needed and we then have to purchase again, it can take very long until we have the new materials again. Therefore, inbound materials may receive quality inspections (possibly double with FAIs), but this is not often the case.

We can also deal with shrinkage by purchasing extra. Materials may include a certain shrinkage factor. For example, if we know that we need on average 10% more of a material, we preferably have at least 10% additional materials. This factor is generally accounted for in the cost price to the customer, although the calculated costs are at times too little according to experts. With this extra sum, either fixed extra demand is added to the first purchase order, or safety stocks are created. Especially in the case of the latter, however, it is not always clear that the safety stock is used to cope with shrinkage. And while the shrinkage factor may be too low at times, we at the same time observe that actually built safety stocks may be (way) too high given the shrinkage factor.

We believe that shrinkage and quality inspections may have a significant impact on costs and tardiness. Optimising the combination of the correct extra quantities and materials to inspect could be an interesting topic. The data on this is limited however, as we do not have defect data for many materials and we cannot e.g. generalise defects to a supplier. It would therefore require a different approach. We further consider this in Section 2.6.

2.4.5. Warranty periods

The supplier gives a certain warranty period for their materials. If materials are used after the warranty and turn out to be defect, the supplier is not liable anymore. The warranty period may vary heavily depending on the type of material. Materials are preferably used within the warranty period, as defects are then still refunded. However, these defects may often be detected only later in the manufacturing process, when valuable manpower and time that was already spent is a much larger cost factor than the material itself. We believe that it could be an important parameter to include when considering if a material should possibly be stocked or not, especially if we consider quality issues. That is, if a material is known to have much quality issues and also have a short warranty period, it may not be wise to keep it on stock (for a longer period of time).

2.4.6. Revision risk

The risk of a buy material being revised is an important characteristic. As there is practically no commonality in Aerospace, all buy materials that are replaced by another material become obsolete. In

Aerospace, the materials that are revised mainly belong to products that are still rather new (i.e. at most 10 years old), or if it contains lead (which is replaced for health risk reasons). Most products are already manufactured for a long time, and the risk of revision for the underlying buy materials then is rather small (but do occur).

If we wish to quantify the revision risk, we could transform the expected number of revisions during a time period (revision rate) to a probability. We would then first have to retrieve the revision rate. Within the ERP system, materials have a revision number attribute. It may include a number and a letter, in which the number should mean a functional change and the letter a change in documentation or the like. By considering the increase in number over a certain period, we are supposed to get an indication of the revision rate. Unfortunately, we observe that this does not hold in practice. Seemingly unimportant changes in the ERP system turn out to be actual material changes. Exact information on all changes is stored in the PLM system. However, retrieving this information on a large scale for many selected materials is not possible. Therefore, if this risk is to be considered, we suggest to use expert opinion rather than this approach but realise that it remains difficult to assess.

2.4.7. Salvage value

If a material still has some value once it becomes obsolete for production, we can deduct that value from the loss we would elsewise have. Possible salvage values may result from selling the stock, or from selling the raw materials after shredding the items. Be it unfortunate for Thales, we deduct from interviews with experts that salvage value is negligible. At times, stock may be sold if a buyer is found, but that is hardly ever the case for Aerospace materials. Stock that cannot be sold is generally shipped to a company who takes care of shredding or the like, without any payments. Therefore, we conclude that it is irrelevant to further consider in any analysis.

2.5. Managing items and inventory policies

The way that items are managed and inventory policies are closely related. We introduce a few new concepts that are essential to know, in order to understand the used inventory policies. We then further investigate what the current policies are, and how Thales currently decides which policy a material gets.

2.5.1. Inventory policies

Materials may be *hard pegged* or *soft pegged* in the ERP system. A material that is labelled as hard pegged is directly linked (pegged) to a specific project (or sales order). It cannot be used for another project and excess stock is not even seen as available, unless a Master Planner manually reallocates the material to another project or assigns it as reservable. These manual interferences are a downside of hard pegged materials, in terms of easily overviewing the stock levels. Soft pegged materials are not dedicated to any project in general. However, soft pegged materials may be reserved for a project on the short term. An issue that arises with soft pegged materials, is that materials can be assumed to be plenty on stock to fulfil demand of a certain project, but are unexpectedly used for another project. For hard pegged materials, this is not the case.

Except for some of the differences as described in Section 2.2.2, the Aerospace materials that are included in the MRP system have a PTO policy. No optimal order quantities are determined consistently, and we observed no more than a few safety stocks at Aerospace. Per the availability differences between soft and hard pegged materials as described, we consider hard pegged (MRP included) materials as a separate policy next to soft pegged MRP materials. The third policy is the min/max policy for materials that are not included in the MRP. Materials with a min/max policy are reordered when stock level reaches a certain minimum level, up to a maximum level. It is mainly the case for floor stock materials, such as wire, sleeves or washers. The materials are stored in boxes at the work cells, and are replenished from the *central warehouse* (CW) through a Kanban system. When a box is empty, a new one is provided from the CW. If the inventory level at the CW drops below the minimum level, an order is sent to the supplier such that the order quantity and the remaining stock level together add up as the maximum level. Note that the min/max policy applies to the stock at the warehouse, which is different from the Kanban system at the work cells.

In order to determine the applicable minimum level, the average demand per day of the past year for that material is determined. The minimum level then should be higher than the procurement lead time in days multiplied by the average demand per day; in other words, the expected total demand during procurement lead time. Therefore, it is basically a safety stock. However, this is the procurement lead time from the ERP system. No variability of demand or procurement lead time is taken into account, which should be in order to account for uncertainty. Still, the minimum may also be adjusted manually to cover more than the expected total demand during procurement lead time.

Each quarter, the demand of materials and their procurement lead times are evaluated. If the minimum levels are (much) higher or lower than the expected demand during procurement lead time, they are adjusted. The reasons behind these anomalies should be checked such that the minimum levels are not increased or decreased too much due to outliers. Then again, since Thales has over 100,000 different materials in its ERP system, things could possibly go wrong here. Also, min/max policies based solely on usage in the past pose a serious risk of delaying manufacturing activities. Irregularly high demand can result in stock-outs, and possibly waiting long periods for cheap materials. For that reason, there is an on-going project to also take into account future demand to determine the minimum level. Next to the minimum level, the maximum level of the policy has to be determined. It can be based on several rules, but is often equal to the MOQ.

To recap, the three material management/inventory policies we include in further consideration are: hard pegged MRP materials, soft pegged MRP materials, and (soft pegged) min/max materials. How does Thales determine what policy a material should have? Generally speaking, hard pegged MRP materials are more expensive materials (> €250), and materials which usage has to be tracked precisely (e.g. materials requiring export licenses). Soft pegged materials cover the remainder, with min/max policies for floor stock and MRP inclusion for the rest. However, we also find MRP-included floor stock materials in some cases. A reason could be that there has been no demand for min/max policy materials in the recent past, hence MRP inclusion is more reliable.

2.5.2. Evaluation of inventory policies

In order to more closely inspect the three inventory policies, we plot the procurement lead times of buy materials versus the material costs in Figure 2.10. We include all materials from the BOMs of the products named in Section 2.3.2, so some materials appear more than once. The log value of material costs is taken for better visualisation because material prices range heavily from tenths of cents to over €25,000 per unit. About one third of these materials are hard pegged and MRP-planned, 26% is soft pegged and MRP-planned, and the remaining 41% has a soft pegged min/max policy. As could be expected, the more expensive hard pegged materials generally have a longer procurement lead time as well. However, there are also quite some soft pegged cheap materials with a min/max policy having procurement lead times of more than 40 weeks.

As described in Section 2.4.1, we consider slack to the critical path more important than absolute procurement lead times. Figure 2.11 plots the same material costs of buy materials versus named slack. Again, this is done for all BOMs of the products we consider. A material appearing more than once may therefore have different slacks, dependant on the product it is part of. What we can see clearly here, is that materials that are cheap can still have large influence on delaying customer lead times hence on OTD. From this new point of view, there are many more soft pegged materials – both MRP-planned and min/max – with such high impact. It emphasises the need to use slack to the critical path when assessing criticality of materials, and therewith inventory policies.

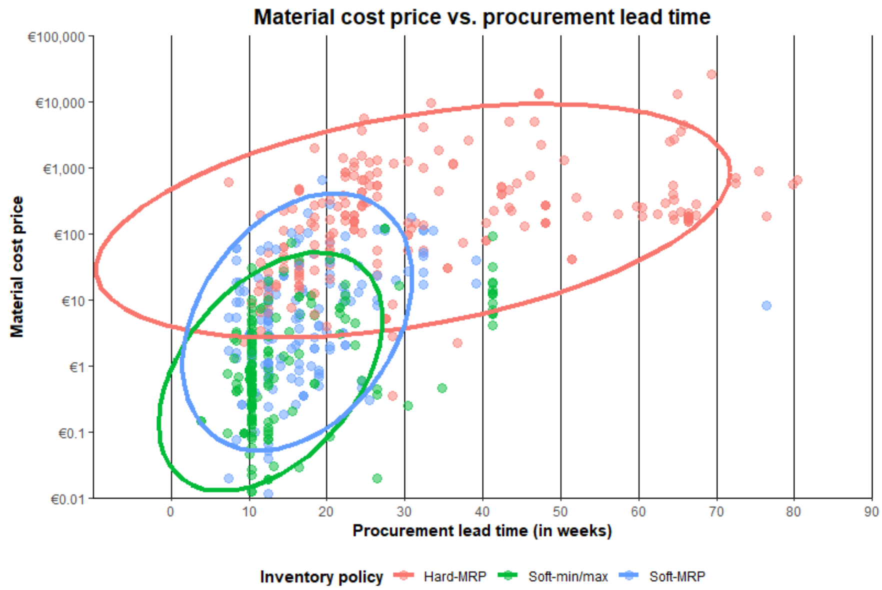

Figure 2.10. Materials costs versus procurement lead times, categorised by inventory policy.

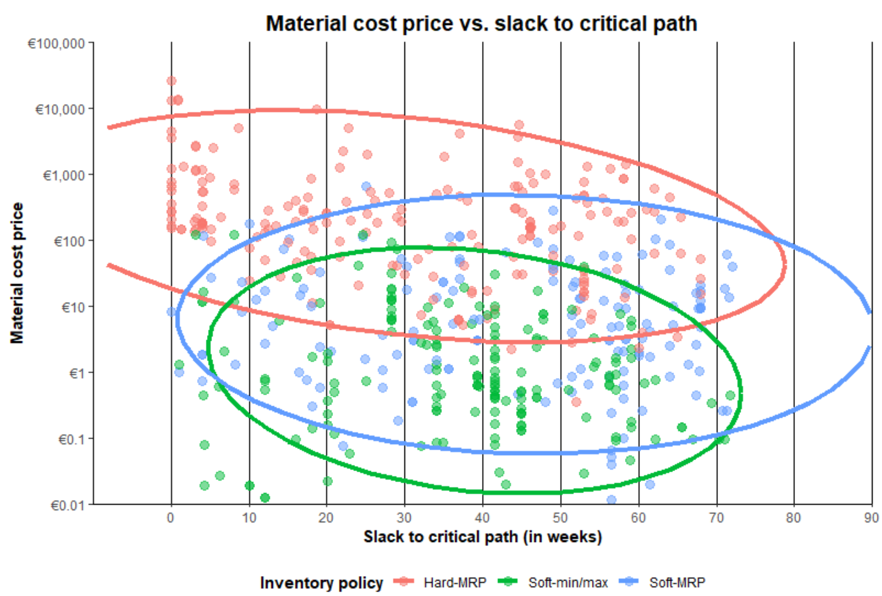

Figure 2.11. Material costs versus slack to critical path, categorised by inventory policy.

We also observe other problems with regards to safety stocks and min/max materials. Although we only see a few safety stocks, their levels were all high or even extremely high: A safety stock of 2,500 units while total demand for the coming years is below 1,500. The reason for the safety stock is not always clear either. It usually is for uncertainty in deliveries, but it can also be used for shrinkage. For both cases, we think that the safety stocks are quite high, and therefore most likely *too* safe. A problem that we see at min/max materials at Aerospace is that usage is not always registered well. As a result the min level may be too low, resulting in stockouts.

To summarise, we note that min/max policies on past demand can be unreliable, soft pegged MRP-planned materials may be unexpectedly used for other projects, and hard pegged materials have shortcomings in showing material availability and reallocation. Although there is a clear positive relation between material costs and procurement lead times, we do not see such a thing when considering slack to the critical path. We conclude that slack to the critical path is crucial to take into account in decisions, and that we must improve current inventory policies or find new ones to overcome the downfalls of the current policies. Other analyses may be interesting to do as well, such as assessing a material's delivery lateness (mean and variance). However, we do not have data for most of the materials specifically, so we cannot execute that analysis now.

2.5.3. Consumables

A special case of buy materials is the consumable. Consumables include material such as glue and chemicals that are used in the assembly process. At consumables, we find a policy similar to the min/max policy. At the work cells, the consumables are stored in cabinets in their original container (e.g. tubes or pots). A few people are responsible for periodically reviewing the number of chemicals present and their expiration date. When consumables must be replenished, new containers are picked at the chemical warehouse. The Manager Chemical Warehouse then decides on placing purchase orders based on the work cell review lists, the stock levels at the warehouse, and experience on demand in the past.

Consumables are not registered entirely correctly on the BOMs in the ERP system. Most of the times they are registered on the BOM, but in some cases they are missing. And even if they are not missing, the quantity required per (sub)assembly may be missing or incorrect. For that reason, the calculated demand from the ERP system cannot easily be translated to the need of consumables for a project. The most complete and up-to-date BOMs are registered in the PLM system. However, we cannot simply combine the BOMs on that system with demand registered in the ERP system. Therefore, it is crucial that all consumables are registered correctly in the ERP system, including the required quantities. Fortunately, in contrast to the past, there have been no significant production stops due to the lack of consumables during the past few years. Nevertheless, the process seems largely based on timely communication, experience and knowledge of those involved and usage in the past. Just like the min/max policies, it is prone to stock-outs with irregular demand.

2.6. Research methodology

In this chapter, we perform exploratory research to further examine research opportunities with a special focus on supply and demand uncertainty. We sketch the process flow from bid management to receiving materials, and analyse data of both procurement and sales. We observe that the forecasts are not (yet) good enough to use as input for decisions on uncertain demand. At the same time, we see that safety lead times are roughly the same for most materials – independent of actualised deliveries or material costs. These safety lead times are definitely not optimal. Within Thales' PTO strategy, good safety lead times can have a large impact on product tardiness and costs. Therefore, we decide to build a safety lead time optimisation model as the core of this research.

While we disregard dealing with uncertain demand, we still observe that high quantity discounts can offer large cost reduction. When deciding to purchase higher quantities for these reductions, holding costs or other parameters are not always taken into account. Therefore, we also include a *Total Cost of Ownership* (TCO) model that allows to quickly find the optimal order quantity in a simple setting. Unlike the safety lead time optimisation model, we do not investigate results from the TCO model, because we just see it as a simpler supportive tool to ease purchase order decision making. Furthermore, we find that safety stocks are hardly used but if so, they are likely high considering their goal. Additionally, we see that our cost-slack analysis shines a new light upon assessing materials. In order to ease decision making on materials, and their safety stocks in particular, we decide to design a material classification method. Finally, we will investigate what process improvements can be made to the sales and procurement activities.

One specifically interesting point that we also see, is shrinkage and quality inspections. Defects during production may cause significant delays, and quality inspections are not often performed. We are convinced that an optimisation of quality inspections, or at least a study on how to deal with shrinkage, could be a promising research subject as well. However, we think that optimising safety lead times is the highest priority in the PTO strategy, with better expected improvements. Therefore, we disregard shrinkage and quality inspections in the remainder of this research.

2.6.1. Research questions

Given the research subjects named in the introduction of this section and the main research question from Section 1.5, we define the research questions for this thesis. We distinct four stages in our research, with corresponding research questions.

The first stage has just been completed as part of exploratory research. As we sketched in Section 1.5, we first needed to analyse the current situation before we could actually find the best research opportunities to consider in the remainder of this research.

1. What are the most promising research opportunities to consider in the remainder of this research?
    a. What activities are involved between receiving a customer request and receiving materials from suppliers?
    b. How are current S&OP forecasts on sales orders determined, and how do they perform?
    c. What are relevant material characteristics that play an important role in inventory and purchasing decisions?
    d. Which current inventory policies used at Thales can be identified, and how do they perform?

Before heading into solutions, we take a look into literature to gain more knowledge. After literature background on inventory management, we analyse how to deal with procurement lead time in a manufacturing setting. Then, we consider how to classify materials using multiple criteria. Chapter 3 presents this literature research.

2. What can we learn from literature with regards to:
    a. Inventory management in general?
    b. Dealing with procurement lead time variance in a manufacturing setting?
    c. Classifying materials using multiple criteria?

Once we have enough background information, we can propose solutions. These solutions are designed in Chapter 4. The core of the research is the safety lead time optimisation model, which we first discuss. Then, we describe the required process changes. The next step is to build the TCO model. Lastly, we explain the design of the material classification approach.

3. How do we build the safety lead time optimisation model?
    a. How can the relation between delivery lateness and safety lead times be modelled?
    b. How can we solve the model?
    c. How do we cope with uncertainty of parameters in the model?
    d. How do we implement the model?
4. What process changes are required to improve the current situation?
5. How do we build the Total Cost of Ownership model?
6. How do we design the material classification method?
    a. Which combination of parameters are criteria in material classification, and how will we rank materials in different classes?
    b. Which policies are most suitable for the different classes?

Finally, insight in possible improvements from the safety lead time model is crucial, which we consider in Chapter 5.

7. How does the safety lead time optimisation model perform?
    a. How can we measure improvement from the model?
    b. What improvement can we expect from the model?
    c. How does the model solution method perform?
    d. What is the impact of uncertain parameters on the model's results?

Figure 2.12 summarises the research approach. The subjects in the figure are keywords that are directly related to the research questions. The chapters in which these questions are answered are given. The final chapter serves to conclude on and discuss the research.

Stage	Current situation	Literature	Model and process design	Performance evaluation
Chapter	*2*	*3*	*4*	*5*
Subjects	• Activity overview • S&OP forecasts • Material characteristics • Inventory policies	• Background on inventory • Procurement lead time uncertainty • Classification	• Safety lead time optimisation model • Process changes • TCO model • Material classification approach	• Performance of safety lead time optimisation model

Figure 2.12. Overview of the research approach.

2.6.2. Research scope

As we first performed exploratory research before defining research questions, we already narrowed down our scope. Still, in addition to the research questions, this section defines more specifically what is included in the research and what is not:

- ➢ Coping with procurement lead time uncertainty is the main centre of interest. Other areas like the TCO model are also discussed, but in lesser detail.
- ➢ The focus of this research lies on buy materials. As we do have an assembly setting, we of course have to evaluate manufacturing activities up to a certain point. However, we mainly assume this as deterministic.
- ➢ Although Aerospace is the product chain that we aim to improve, proposed solutions should preferably also be implementable at other product chains of Thales.
- ➢ Calculations and tools are limited to be given in Microsoft Office programs, i.e. mostly Excel. Due to security reasons and license costs, different software is undesirable.
- ➢ Potential adaptations of the ERP system for the new approach of material management are not considered in technical detail.

3. Literature review

This chapter covers a literature review performed to gain a deeper understanding on dealing with procurement lead time uncertainty and classification of materials. First, we provide a background on inventory management in Section 3.1. In Section 3.2, we elaborate on how to approach procurement lead time uncertainty in mathematical models, applicable to the Aerospace setting. Next, Section 3.3 describes classification of materials.

3.1. Background of inventory management

An important theme of our research is inventory management. Although we do not primarily focus on actually building substantial stocks, it is essential to provide some background on the matter – not in the last place to show some disadvantages of holding inventory. We provide insight in *what* types of inventory exist and *why* inventory is held. We focus on the Aerospace context, which means that not necessarily all inventory types, costs and considerations are included.

3.1.1. Types of inventory

There are multiple means of categorising inventories. Two ways of doing so are by looking at the type of item, and their function in inventory. By the first, we distinguish the following types:

- *Raw materials:* These are the materials necessary for production activities of a company. Earlier referred to as buy materials.
- *Work-in-process:* Made up of all components, parts, assemblies and so forth, that are currently (waiting to be) processed in the system.
- *Components:* Also known as subassemblies. It is the stage between raw materials and finished goods, where items have already been processed but have not yet reached completion. In our context called make materials.
- *Maintenance, repair and operating (MRO) goods:* Items that are used in the process or support the process but are not directly a part of it. Our definition of consumables is part of MRO goods.
- *Finished goods:* The final products of the process that are shipped to the customer.

Considering the function of inventory, based on Silver, Pyke & Thomas (2017) we find the following categories applicable to our context:

- *Cycle inventories:* Inventories that result from ordering or producing in larger quantities than demand (in batches). A certain number of items is stored, and that number decreases over time due to demand. The time between two batches is the cycle, and inventory between these periods the cycle inventory.
- *Safety stock:* With uncertainty in demand or supply, safety stock can be used to account for this uncertainty and to be able to fulfil demand when procurement lead time or demand is higher than expected. Hence, safety stock is often based on a desired service level in addition to uncertainty parameters of demand and procurement lead time.
- *Pipeline:* Includes goods in transit between two stations in a process. It is proportional to the transit time between locations and the usage rate of the item.
- *Decoupling stock:* Inventory that is held to make up for out-of-sync demand and supply between two dependent operations. It allows to maintain workflow and decentralised decision making.
- *Congestion stock:* Due to limited capacity in e.g. production equipment, stock can build up while waiting to be processed.

3.1.2. Considerations for holding inventory

Not all types of stocks as described above always exist in a supply chain. For instance, congestion stock is undesirable and only results from named limited capacity. Per type of inventory, we might find different reasons why inventory is held. Based on Nahmias & Olsen (2015), we find the following motivations:

- *Economies of scale:* Through purchasing in larger sizes, quantity discounts may apply that reduce unit price. In addition, larger order sizes result in fewer orders, meaning less fixed costs of ordering. Related to that, a firm can considerably reduce costs by producing in larger batches as setting up production requires time and money. Even when there is no demand yet for a product, these lower costs could be an economic motivation to produce larger batches.
- *Uncertainties:* Due to uncertainties and opportunity costs, a firm can benefit from having inventory to avoid risk of stock-outs. These uncertainties include uncertainty in demand, in supply (whether it will be available at all), in procurement lead time, in labour availability, and so on.
- *Competitiveness:* By having components or finished products on stock, customer lead time can drastically decrease, on-time deliveries may increase and even costs may be lower. This can increase a firm's competitiveness.
- *Speculation:* Prices for items or conversion rates might negatively change in the future, due to which it might be beneficial to build inventory.
- *Smoothing:* In anticipation to peak demand, producing and storing inventory can aid in decreasing production disruptions.

Inventory may also appear for other reasons. For example, purchasing materials might be restricted to an MOQ which will result in stock if less than the full quantity is needed. Also, with increasing transportation times, pipeline inventories will grow. Of course, holding inventory also comes with costs and risk. Products might not be sold at all, materials or finished products could become obsolete, funds are not available for other use. Keeping inventory always boils down to a trade-off between cost and responsiveness. The goal of good supply chain design is to find the right level of responsiveness, with the right form, location and quantity of inventory at the lowest possible cost (Chopra & Meindl, 2013).

3.1.3. Costs associated with placing orders and holding inventory

Holding inventory results in several cost types. The main type of inventory costs is holding cost or carrying cost. It is the total of costs related to owning inventory, either physically on hand or in transit. The components of these costs are as follows:

- *Capital cost:* These costs are related to not being able to spend money and receive interest on other investments. It is generally considered the most significant part of holding cost. It is also quite subjective, as it is hard to say how much return on investment a firm would be able to get on different projects. Often determined as the weighted cost of capital.
- *Storage cost:* The actual costs for storing inventory. These are rental or mortgage costs, personnel costs for handling, electricity, and so forth.
- *Obsolescence cost:* Holding inventory also brings risk of obsolescence. Items might break, deteriorate or become completely obsolete over time. This risk is high for Aerospace, because once a customer no longer demands a product, related items mostly become obsolete.
- *Service cost:* The fourth component includes all 'service' costs, such as insurance and taxes.

Together, these costs are defined as a certain percentage of item cost that returns each period. In practice, one often sees that these costs are between 15-35% of the value of inventory per year. Figure 3.1 gives a possible combination of named costs.

Other costs that we include and found earlier are different fixed costs of ordering and inspection costs. When an order is placed, it may come with set-up costs regardless of the order size. Inspection costs apply especially to more expensive materials, and are often only necessary if the material has not been produced for a while or the production process has changed. Penalty costs depend on how a firm handles demand that cannot be filled from on-hand inventory. Excess demand can either be back-ordered or is lost sales. Different costs apply in both cases and are heavily case-dependent. In both cases, there is also a loss-of-goodwill costs, which can be very difficult to estimate. At Aerospace, excess demand is always backordered.

Best Practice Inventory Carrying Costs (Typical Ranges +/-10%)		
1 - Capital Costs	Interest, Opportunity	12.0%
2 - Inventory Service	Taxes, Insurance	1.5%
3 - Storage Space	Warehousing	3.0%
	Clerical/Recording Keeping	1.5%
	Inventory Counting	0.7%
	Material Handling	1.3%
	Depreciation	1.0%
4 - Inventory Risk	Obsolescence	4.0%
Total		**25.0%**

Figure 3.1. Example of inventory carrying costs structure (Weber, 2017).

Increasing batch sizes might result in a lower unit price due to quantity discounts, and will result in less set-up costs as the order frequency decreases. At the same time, cycle inventory increases and therewith total holding cost increases. This is a core trade-off at inventory management. Figure 3.2 shows this trade-off for a simple model with no uncertainties, procurement lead time and quantity discounts, and constant demand. In these cases, we find the lowest total costs when ordering costs are equal to holding costs. This quantity of materials to purchase is the *Economic Order Quantity* (EOQ). Although exact cost functions may differ in practice, we often see a similar trade-off.

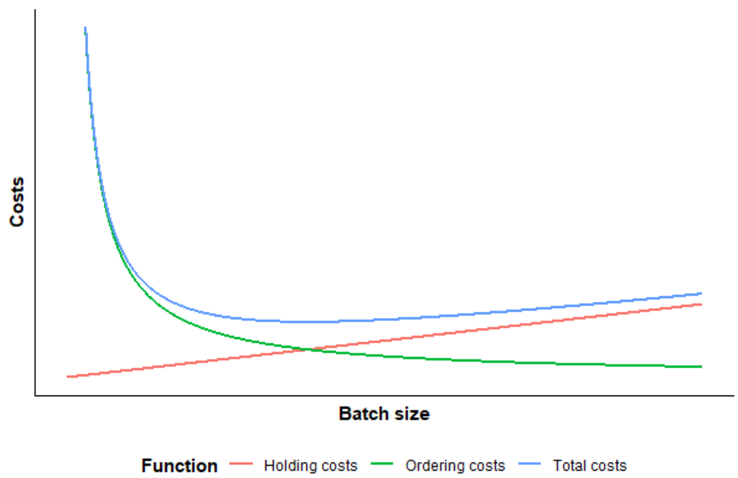

Figure 3.2. Order quantity trade-off.

3.2. Procurement lead time uncertainty

The main focus point in our research is coping with uncertainty in procurement lead time. We consider procurement lead time uncertainty from a manufacturing perspective, and in particular an assembly system perspective. In this section, we elaborate on literature that describes this perspective. We do not intend to be exhaustive, but rather show what model characteristics we can distinguish and how to best approach building our own model.

3.2.1. Safety stocks vs safety lead times

The most recent overview of supply planning with known demand but uncertain procurement lead times is given by Dolgui, Ben-Ammar, Hnaien, & Ould-Louly (2013). As it turns out, attention for this setting from literature is still quite limited, as already noted by Mula, Poler, Garcia-Sabatar, & Lario (2006). In fact, even until today, the majority of contributions to literature in this area were done by the authors of Dolgui et al. (2013). A very closely related research area is that of MRP and the *Master Production Schedule* (MPS). Simply put, the MRP determines a replenishment schedule based on gross demand given by the MPS and current inventories. The basic MRP assumes deterministic demand and supply, which is its downfall in practice. Therefore, much research has turned towards parametrisation of the MRP, to account for different kinds of uncertainty. Dolgui & Prodhon (2007) provide an overview of literature on MRP specifically, where we see many similarities with literature on procurement lead time uncertainty in general.

The two most common approaches to deal with uncertainty in procurement lead time are safety stocks and safety lead times (Dolgui & Prodhon, 2007). In an assembly setting, the first allows us to start assembling even though a material is delivered too late. Safety lead times for materials on the other hand aim to minimise delay of the production start time by building a buffer for expected material lateness. It is debatable which option is preferred. Whybark & Williams (1976) and Van Kampen, Van Donk, & Van der Zee (2010) believe safety stocks are to be used for uncertainty in demand, while safety lead times should deal with uncertainty in procurement lead time. In contrary, Grasso & Taylor (1984) concluded from their simulation results that safety stocks are preferred in both cases. Then again, according to Plenert (1999), one could reduce or even eliminate safety stocks with sufficient safety lead times. In our view, both could be necessary. Unfortunately, to the best of our knowledge, literature describes no model that includes both.

In terms of holding costs, safety stocks are generally much more expensive than safety lead times. Figure 3.3 shows this difference, where the green line is the inventory level with safety lead times and the red one with safety stocks. A safety stock is held at a more or less constant level throughout the year, incurring holding costs over that whole period. Safety lead times only result in more holding costs for the period that we actually hold stock longer. However, if we set safety lead times at the materials with the longest procurement lead time (or actually 0 slack, recall from Section 2.4.1), we also increase the customer lead time. This is of course undesirable. Nevertheless, these materials are only a very small portion of all materials. We choose for the approach to set safety lead times where possible, and check if it is viable to build a safety stock for the remaining few materials. For now, we are interested in optimising safety lead times in an assembly system. Translating the basic principles to safety stocks (if needed) should then not be too difficult.

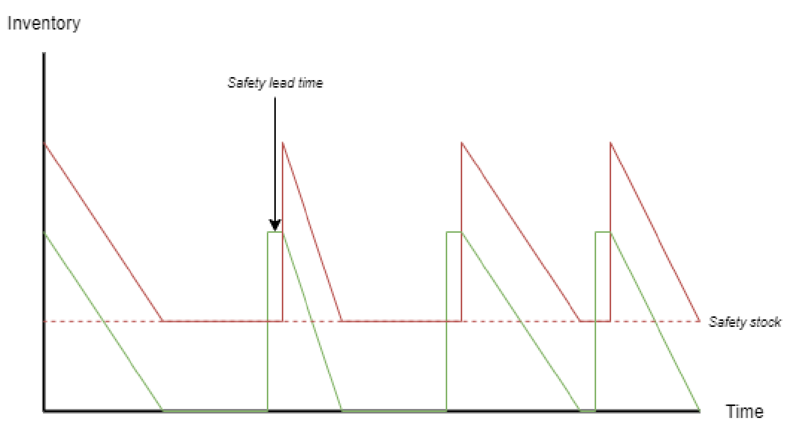

Figure 3.3. Safety stock vs safety lead time.

3.2.2. Model characteristics

There are a few model characteristics that we need to address before describing a general model design. Although we could distinguish on many, we for now consider the optimisation criteria, the number of assembly levels, and single versus multi-period as most important.

The first model of Yano (1987) and many others since then, e.g. (Hopp & Spearman, 1993; Song, Yano, & Lerssrisuriya, 2000; Ould-Louly & Dolgui, 2002; Fallah-Jamshidi, Karimi, & Zandieh, 2011), minimise the sum of holding costs of materials and *tardiness* (production delay) costs if the product is finished later than planned. While the earlier models of Yano (1987) and Hopp & Spearman (1993) do not include holding costs of materials while waiting for another late material, this is included in later models. However, as for instance Borodin, Dolgui, Hnaien, & Labadie (2016) note, it may be hard to determine tardiness costs in practice. Therefore, many remaining models minimise holding costs with a service level constraint. Then again, unless a customer specifically asks for a certain service level, determining a service level is often quite arbitrary as well. A third approach is to also consider setup costs (Ould-Louly & Dolgui, 2013), which allows to view the costs from the broader manufacturing perspective. Obviously the production batch size influences safety lead times through different holding costs, and by including the setup costs we may get closer to determining the ideal production batch size.

Concerning the number of assembly levels, we distinguish between two-level and multi-level assembly systems. Remember the difference from Section 2.2.1. Examples of the multi-level models are (Ben-Ammar, Marian, Wu, & Dolgui, 2013) and (Ben-Ammar, Dolgui, & Wu, 2018). The multi-level property changes the function of the product delivery delay. As we later more explicitly show, the product tardiness is generally assumed equal to the maximum lateness of all materials. In the multi-level assembly system, we can have multiple delays and therefore product tardiness is equal to the sum of maximum material lateness at each assembly level. One could also debate if the holding cost structure differs. Does a delay of one material result in an increase of holding costs for all other materials as we deliver the final product later, or do we limit the costs to the delay of the concerned subassembly?

The third interesting point is whether the model analyses a single period or multiple periods. Recall from Section 2.2 that we produce in multiple batches and also aggregate demand of 30 workdays with the FDS. Tardiness of one batch may delay subsequent batches too (actual capacity is not infinite), and material demand for a batch may have been fulfilled by an order for an earlier batch. Unfortunately, multi-period models are scarce and not feasible for our setting. These models are recently mostly described by the authors Ould-Louly & Dolgui (2008, 2009, 2011, 2013), but also earlier by others (Proth, Mauroy, Wardi, Chu, & Xie, 1997). All the models have a downside for which they do not apply to our case. Where one considers only one component (Ould-Louly & Dolgui, 2013), another assumes the same holding costs for all units (Ould-Louly & Dolgui, 2008), and a third analyses a combination with the EOQ (Ould-Louly & Dolgui, 2009). The most relatable model (Ould-Louly & Dolgui, 2011) approaches our FDS case with *order periodicity* (the fixed cycle in which material orders are placed), but it is not useful as it assumes that all different materials are purchased at the same moments. They also note that different periodicities for each material is future research work, which still has not been performed.

Using a single-period model to analyse our setting is an option. It will probably underestimate tardiness and both holding and tardiness costs, but its optimal safety lead times most likely still improve the current situation. In Section 4.1.1, we consider how to deal with this. The multi-level assembly system is an interesting field, but we limit ourselves to a two-level assembly system for now. It may not be fully correct, but it should approximate the situation well since we have at most four assembly levels and manufacturing lead times are small. We also consider only one product at a time. Commonality of materials with multiple products adds another dimension to the model, but we saw in Section 2.4.3 that Aerospace materials are generally only used for one product. In the two-level assembly system, we neither have to include manufacturing lead time when considering planned procurement lead times. We also believe a model with tardiness costs is more useful, and elaborate on the uncertainty of its value in Section 4.1.4. Set-up costs are considered out of scope for now. In the next section, we describe how such a model works.

3.2.3. Model design

We do not write mathematical notation for the model in this section, as Section 4.1.2 already does that. Instead, we illustrate the basic idea of the model and discuss how to solve it. Figure 3.4 visualises the holding and tardiness costs for a three-level assembly system. Lead times will in practice never be such that all materials are ready exactly for the next stage. Since all materials are needed to build an assembly, the product tardiness is equal to the maximum delay of all materials. Those materials that are earlier than that last material are held on stock, which incurs holding costs. With increasingly long delays of one material, we therefore also see higher holding costs for other materials. Lateness of materials also cause delay of the production start date, resulting in tardiness costs (e.g. inefficiency costs, penalties). Without safety lead times (or materials kept on stock) and just-in-time planned delivery dates, the risk of tardiness is very high as we have many materials.

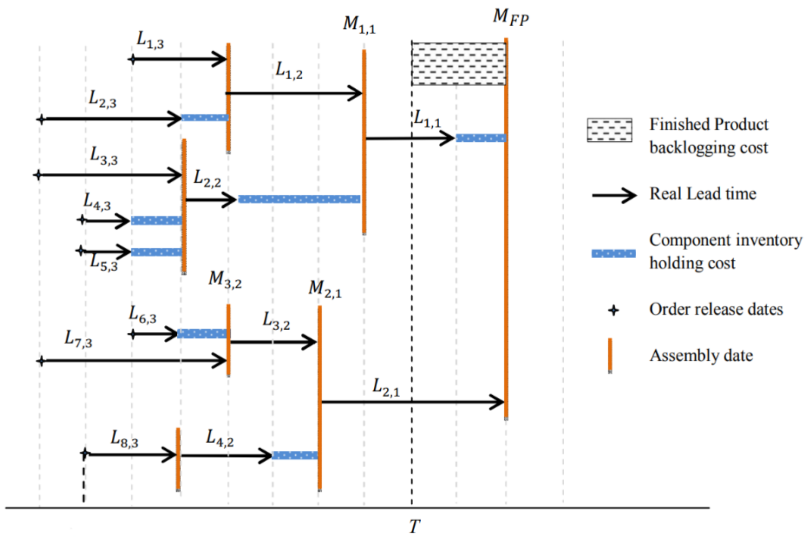

Figure 3.4. Holding and tardiness costs visualized for a three-level assembly system. The line above T is the moment production of the product should have started (Ben-Ammar et al., 2018).

So how do safety lead times help? As we increase safety lead time of one material, the risk of tardiness decreases. Therefore, for all other materials, the expected holding costs of waiting for named material also decrease. In addition, tardiness costs also drop. Of course, the trade-off is that the material with increased safety lead time has higher holding costs for being earlier than required. In the bigger picture, we want to find the right combination of safety lead times such that holding costs of materials being early (due to safety lead times) are balanced with holding costs of materials waiting for other materials and tardiness costs. The primary parameters that influence the optimal safety lead time for a material are its costs and its delivery lateness distribution function.

Although many models use this approach, in Section 4.1.2 we base our notation on Song et al. (2000) thanks to its clear explanation. They also describe a simple greedy heuristic which actually gives the optimal solution. The reason for that is the cost function is a non-negative linear combination of convex functions, which is convex itself as well. After we introduce notation and the heuristic in Section 4.1.2, we show this.

3.2.4. Conclusion

Section 3.2 serves to understand how we can deal with uncertainty in procurement lead times from an assembly perspective. Cost-wise, safety lead times are often preferred over safety stocks, but safety lead times may not be applicable in all cases. We also observe a limitation in current literature for our setting, especially when considering multiple periods. We decide to continue with a two-level assembly system as it is a good approximation compared to a multi-level system, and further investigate how to deal with multiple periods in Section 4.1.1. Finally, we find a simple heuristic to solve the model, which we describe in Section 4.1.2. In Section 4.1, we further consider what additions to the models from literature we can make to better fit it to our setting.

3.3. Material classification

For some materials we want to consider safety stocks. This will not be possible for all items. We want to consider how we can simultaneously regard material criticality and related risks. Also, intensively researching and updating model parameters for all materials is time-consuming and costly. Classifying materials offers a way to structure decisions for groups of materials. We consider both quantitative as qualitative classification models in this section.

3.3.1. Quantitative models

The most common way of quantitative classification is the ABC method, introduced by General Electric (Dickie, 1951). It was used for inventories of their finished goods. ABC classification is a Pareto based method, where Pareto's principle was that generally 80% of certain results come from 20% of the actions (Pareto, 1971). In the original ABC classification, it means that 80% of total sales volume comes from 20% of the items. This class is considered most critical, therefore A. 30% of the remaining items provide another 15% of total sales volume, and is labelled B. Finally, C consists of the last 50% of the items which only incorporate 5% of total sales volume. Based on this classification, most attention is of course spent on consecutively classes A, B and C.

Classifying on sales volume makes sense when speaking of finished goods, but in our setting we have materials. As Silver et al. (2017) point out, firms may also include inexpensive and slow-moving items in class A if they are critical to business. But what do we define as critical to business? And can we fit that within one criterion? Flores & Whybark (1986) debated that more criteria should be included besides sales value, e.g. procurement lead time for purchasing items or commonality (shared usage) in a manufacturing setting. For two criteria, they proposed a joint criteria matrix as shown in Figure 3.5. An item is scored in a class for both criteria, as usual. If the same class applies for both criteria, the item is simply that class. The combinations outside the diagonal (green) line are reclassified to the three classes AA, BB and CC through managerial judgment, which could involve weights for a criterion. For example, cost is deemed more important in the figure, as higher cost classes generally get a higher combined class as well. It also presents a new problem: How to determine and include preference of one criterion over another? Besides, we can also do this for three criteria with a cube, but it gets increasingly more difficult as the number of criteria rises. How can we deal easily with more than two criteria?

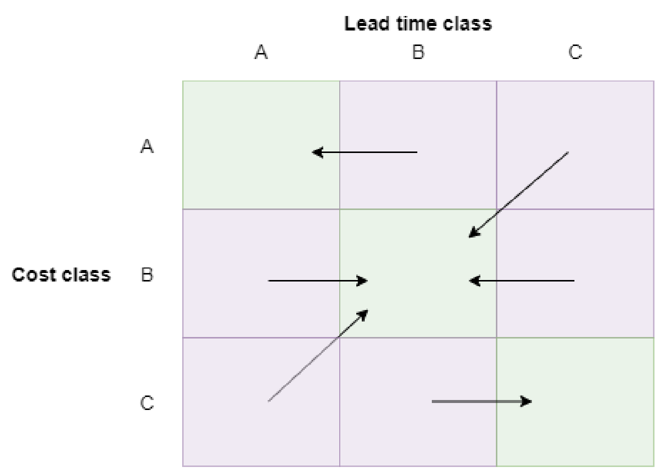

Figure 3.5. Joint criteria matrix, where mixed classes (e.g. A-B) are subjectively categorised as a pure class (i.e. A, B or C).

A partial answer to these questions was given by Ramanathan (2006). He designed a weighted linear optimisation model, similar to *data envelopment analysis* (DEA). It gives items a score between 0 and 1, based on the sum of actual performance at the criteria multiplied by calculated weights. To determine these weights, the model maximises per item the sum of the performances at criteria multiplied by the weights of those criteria. At the same time, the weights are constrained such that for all other items, the sum of their performances at criteria, multiplied by the determined weights of the evaluated item, never exceed 1. In other words, the model basically evaluates how well an item performs at each criteria compared to others, and selects weight values most favourable for itself. The score can be used to rank items, where once again we could apply the original 20-30-50 rule of ABC.

Based on Ramanathan's model, we find other developments. As the linear model has to be solved for each item and takes increasingly longer for more items, Ng (2007) proposed a simplification of this linear model. It has somewhat comparable results to the DEA-like model, but does not require linear optimisation while it needs a subjective ranking of criteria. Zhou & Fan (2007) note that the model of Ramanathan can lead to an item scoring as A with high performance on a less relevant criterion, due to its DEA properties. They suggest to also determine least favourable weights, and use a subjective combination of the most and least favourable weights for a new ranking. Following on Ng (2007), Hadi-Vencheh (2010) presents a nonlinear model which determines a common set of weights for the criteria. The model also needs subjective ranking of criteria, but is an improvement over Ng's model by keeping impact of the obtained weights in all cases. It has slightly different results at cost of less simple implementation.

Apart from the models described above, numerous contributions involved multiple-criteria ABC analysis from equally many perspectives. More examples can be given of weighted optimisation methods, of which also worth mentioning are an improvement to Zhou & Fan (2007) by removing subjectivity (Chen J.-X. , 2012) and the latest notable contribution using acceptability analysis (Li, Wu, Liu, Fu & Chen, 2019). Apart from weighted optimisation models, we for example also see Euclidean distance-based models (Bhattacharya, Sarkar & Mukherjee, 2007; Chen, Li, Kilgour & Hipel, 2008), and artificial-intelligence models (Yu, 2011). Next to ABC models, we find other quantitative models that often focus on demand uncertainty such as the XYZ method. However, that is not useful in our manufacturing setting as demand for all materials of a product is relatively the same. Quantitative models as described here have the shortcoming that it is hard to incorporate categorical measurements. For that, we turn to a combination of quantitative and qualitative classification: *Analytic Hierarchy Process* (AHP) based methods.

3.3.2. AHP methods

A powerful tool for making decisions with multiple weighted criteria is the AHP (Saaty, 1980). Simply speaking, it allows to determine weights of each criteria through pairwise comparisons, and has a consistency check included. The exact approach is elaborately explained and easily found on the internet. Flores, Olson & Dorai (1992) and Partovi & Burton (1993) were the first to use AHP as an inventory classification method; in fact, the scores obtained through AHP is input for the basic ABC classification. We observe many contributions since then, such as a combination of AHP with decision trees and categorical levels (Braglia, Grassi, & Montanari, 2004), fuzzy approaches (Cakir & Canbolat, 2008; Hadi-Vencheh & Mohamadghasemi, 2011), and a hybrid approach based on AHP and the K-means algorithm (Lolli, Ishizaka, & Gamberini, 2014).

It is seen as an advantage that through using AHP, management can actually decide what criteria are most important and also with some weight included. It was also considered a method that may well be used by any manager as no model optimisation is required, although Botter & Fortuin (2000) surprisingly concluded that AHP was considered too theoretical and not acceptable by management. Another point of criticism of using AHP at inventory classification is that criteria must be independent (Partovi & Burton, 1993). If not, they should be aggregated into a single criterion. It may also be considered a difficult task for the decision maker to accurately assign exact values to pairwise comparisons. As expert judgment plays an important role in weighing criteria thus ranking items, the ABC classification results may vary significantly.

3.3.3. Conclusion

In this section, we evaluate classification of inventory items through methods from literature. We find that quantitative models often require solving weighted optimisation models, and generally also solving a model for every single item. Some models do allow for full objectivity, where no human judgment is required. In other cases, a simple ranking of criteria is needed. Including categorical measurements is hard here, however, for which we rather use AHP based methods. In that case, through pairwise comparisons, human judgment is necessary, allowing to determine how much a criterion is preferred over others.

If we relate these findings to our context, it is difficult to say whether any of these methods will help us in our classification problem. Implementing the classification definitely should not be time-consuming or require additional software for solving the models. We also observe that there is little literature on classification in a manufacturing context, while there is much attention for sales items and spare parts. Although there are commonalities, all items have a certain high criticality in manufacturing already. As we are mainly looking for business improvement opportunities while requiring materials available at time of manufacturing, we perhaps need a different approach. In Chapter 4, we consider methods from literature from a practical perspective to our context, when designing the classification method.

4. Model and process design

This chapter involves suggested solutions to problems sketched earlier. In Chapter 2, we conclude that we will mainly focus on coping with supply uncertainty. We also saw that the primary way of doing so is to optimise safety lead times. Furthermore, in Section 2.2.2 we pointed out that there is delay caused by suppliers (delivery lateness) of which we have sufficient data, but also lateness of materials because e.g. we ordered too late. Our solutions towards coping with supply uncertainty are as follows. We build a safety lead time optimisation model to find the best safety lead times given delivery lateness data. To deal with other lateness of materials, we propose improvements to the sales and procurement process. Furthermore, we design a classification approach, to ease decision making on specific materials and assess whether to possibly include safety stocks in addition to safety lead times. Finally, we also build a *Total Cost of Ownership* (TCO) model. This does not directly help towards dealing with supply uncertainty, but helps reducing costs of purchasing and holding materials in various settings.

Section 4.1 describes the quantitative model to deal with delivery lateness. In Section 4.2, we propose the process improvements. Subsequently, we present the TCO model in Section 4.3. Finally, we consider material classification in Section 4.4.

4.1. Safety lead time optimisation model

In this section, we use the data of delivery lateness (recall from Section 2.2.4) to optimise safety lead times. We discuss the model in context of described literature in Section 4.1.1. In Section 4.1.2, we show the model. Subsequently, the delivery lateness distribution is considered in Section 4.1.3. Other parameters are described in Section 4.1.4, and we finally describe the implementation of the model in Section 4.1.5.

4.1.1. Literature context

In Section 3.2.3, we theoretically described a model from literature to optimise safety lead times in a two-level assembly system. We believe that this model very much suits our needs. Naturally, we tailor the decisions to our preferences, but these are details compared to the structure that the theoretical model offers. The changes include a different probability distribution function (delivery lateness versus procurement lead time) and a specific due date with related maximum safety lead time depending on slack. Figure 4.1 shows this last constraint. Given a fixed procurement lead time (as the ERP assumes), we see that the maximum safety lead time that we can assign is equal to the material's slack. If we want a larger safety lead time, we must either order before the purchase kick-off moment (not an option) or postpone the production start date on forehand and therefore increase customer lead time (not a preferred option). Therefore, we set the slack as upper bound for the safety lead time.

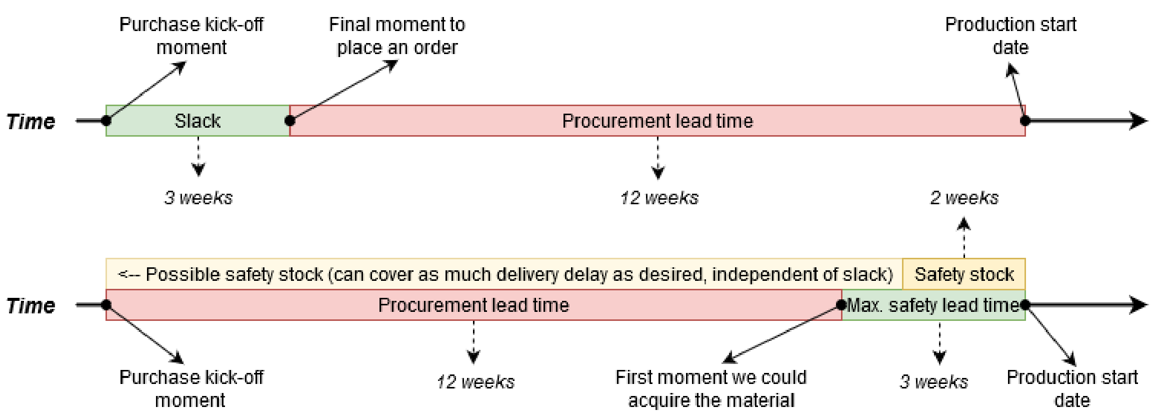

Figure 4.1. Slack constraint in safety lead times, and the role of safety stock at delivery delay.

The most important shortcoming of the model is that it considers a single period, while we have multiple periods (e.g. 12 batches a year) which effect each other. We also saw in Section 3.2.2 that literature offers no suitable models to deal with this issue in our context. Therefore, we attempted to model one ourselves, in which we include multiple batches and FDS (see Section 2.2.2). Unfortunately, our effort to create such a mathematical model was in vain. Nonetheless, we still incorporate the multi-period case in another way. We build a Monte Carlo simulation model that does include multiple batches and FDS. However, it is not nearly fast enough to evaluate the myriad possible combinations of safety lead times. Therefore, we use it to verify the results from the mathematical model, such that we know how accurate that model represents the more realistic setting. Section 5.3.2 discusses the Monte Carlo simulation model and its results.

Figure 4.1 also shows a safety stock. Note that although we cannot further increase safety lead time for some materials, we may still create an extra buffer to cope with delivery lateness by having materials on stock. This may be a safety stock held throughout the year, or e.g. extra materials purchased sufficiently early to be definitely on-time for the first batch (if possible for that material). Although both safety lead times and safety stocks may achieve the same goal, they have a different cost structure and work differently. For example, if delivery lateness exceeds safety lead times we still experience delay, while sufficient stock may allow to produce anyway. Again, literature does not offer a combination of both in the assembly setting, so we considered such a combined model ourselves. We did not achieve to analyse it for the multi-period setting, and a single-period model does not give useful results for the actual situation. Nonetheless, we gained some insights in using these safety stocks for supply uncertainty as briefly described in Section 4.4.3.

4.1.2. Safety lead time model description

Recall an explanation of the model in Section 3.2.3. This section covers the model description and the optimisation procedure. Prior to that, we give the assumptions that we make.

Assumptions

1. *We assume that the demand and supply quantity are deterministic.* In practice, there may be some material defects but often also a higher purchase quantity to correct for that. There can also be some rescheduling of production or demand from the customer, but we consider that very situational and out-of-scope for now. Thales considers to freeze the MPS in the future, due to which rescheduling on the shorter term should be avoided. Therefore it may well be justified.
2. *We assume a two-level assembly model*. This is technically correct for some products/systems in Aerospace, but not all. The two-level assembly model does allow to approximate a multi-level assembly system, in two ways. The first option is to model the full assembly as a two-level assembly. We then do not consider that there may be multiple materials that cause tardiness in different subassemblies, rather than only the latest material. This underestimates total product tardiness. The second option is to model all (sub)assemblies as individual two-level assemblies. In that case, a delay in any subassembly causes tardiness of the product. This overestimates product tardiness, since in practice a delay in one subassembly needs not cause product tardiness if another subassembly has a longer delay (depending on where they are used). We choose for the first option for simplicity reasons, but are aware of the issue when considering results.
3. *We consider batches independent from each other.* Recall our reasoning from the previous section, and the Monte Carlo simulation model to compensate for this issue.
4. *We do not consider initial stock*. This is in line with the previous bullet point. Also, the model shows tactical decisions rather than optimising at an operational level with small initial stocks.
5. *We assume that each material arrival is independent of other materials and of its quantity*. In practice, this is probably not completely true. A supplier who has to deliver e.g. 8 different materials will deliver worse compared to delivering 1 material, but also better than 8 different suppliers with 1 material. We do not consider this as a separate issue in the model, but we will come back to this when estimating model delivery lateness parameters in Section 4.1.3.

6. *We only consider MRP planned materials.* All min/max materials are assumed to be available at any time. Even though this is debatable, the settings of these parameters are a different issue.

Model description

We consider a two-level assembly problem with n distinct buy materials to produce one type of final product or system. Let for each material $1 \leq i \leq n$ the random variable L_i be the delivery lateness in weeks. As we will see in Section 4.1.3, working with weeks is more convenient than workdays. A negative number of L_i equals earliness and a positive number lateness. The cumulative distribution function of L_i is given by F_i. The holding cost rate (per week) is equal to h, the cost price for the required number of items per product/system of material i is c_i, the tardiness costs per week (or costs of delayed production activities) is given by D, and the batch size (in which product amount we place underlying material orders) is q. Our goal is to optimise for each material i the safety lead time x_i (in weeks), such that the sum of holding costs and tardiness costs is minimised.

Now, let us introduce some mathematical notation. We use a bold faced letter to express a vector of size n. We then have the decision variable set $\mathbf{x} = (x_1, x_2, \ldots, x_n)$ and material cost price set $\mathbf{c} = (c_1, c_2, \ldots, c_n)$. For any two vectors $\mathbf{a}$ and $\mathbf{b}$, let $\mathbf{a} * \mathbf{b} = \sum_{i=1}^{n} a_i b_i$. Furthermore, let z^+ be the maximum of $\{0, z\}$. We observe that $(x_i - L_i)^+$ is earliness of material i, and $(L_i - x_i)^+$ is the lateness of material i. As all materials are required before production may start in the two-level assembly system, the production delay due to waiting for materials (tardiness of the product) $T'(\mathbf{x})$ given vector $\mathbf{x}$ then is

$$T'(\mathbf{x}) = \max_i (L_i - x_i)^+.$$

In line with that, we also have the probability of waiting at most t number of weeks for materials as

$$\begin{aligned} P[T'(\mathbf{x}) \leq t] &= P[(L_i - x_i)^+ \leq t \ \forall i] \\ &= \prod_{i=1}^{n} P[(L_i - x_i)^+ \leq t] \\ &= \prod_{i=1}^{n} F_i(x_i + t). \end{aligned}$$

Knowing that, we can calculate the expected number of weeks waiting for materials, $T(\mathbf{x})$, as

$$T(\mathbf{x}) = E[T'(\mathbf{x})] = \int_0^\infty \left(1 - \prod_{i=1}^{n} F_i(x_i + t)\right) dt,$$

which naturally results in production delay cost of

$$D * T(\mathbf{x}). \tag{1}$$

For the holding cost of materials, we have two parts. The first part is obviously due to arriving early. Next to that, we also have extra holding costs for all materials that arrived earlier than the latest, as production is delayed and we receive payment for these costs relatively later. Together, the expected holding cost is

$$qh \sum_{i=1}^{n} c_i E[(x_i - L_i)^+ + (T'(\mathbf{x}) - (L_i - x_i)^+)].$$

Note that

$$E[(x_i - L_i)^+ - (L_i + x_i)^+] = x_i - E[L_i].$$

Also, let

$$l_i = E[L_i], \text{ and } c_0 = \sum_{i=1}^{n} c_i. \tag{2}$$

Then, we can write the expected holding cost as

$$qh(\mathbf{c}(\mathbf{x} - \mathbf{l}) + c_0 T(\mathbf{x})). \tag{3}$$

By combining equations (1) and (3), we have the following expected total cost function:

$$C(\mathbf{x}) = qh(\mathbf{c}(\mathbf{x} - \mathbf{l}) + c_0 T(\mathbf{x})) + DT(\mathbf{x}). \tag{4}$$

Our optimisation problem then is

$$\min_{\mathbf{x}} C(\mathbf{x}). \tag{5}$$

Finally, we add the restrictions. We discussed in Section 4.1.1 how a material's safety lead time may not surpass its slack. We notate this slack of material i by s_i, with the simple constraint

$$x_i \leq s_i, \forall i. \tag{6}$$

Optimisation procedure

We use the optimisation approach of Song et al. (2000), which we briefly mentioned in Section 3.2.3. It works as follows. Since creating a negative safety lead time is counterintuitive and hardly ever useful, we do not consider negative values for **x**. Therefore, we initially set $\mathbf{x} = 0$, and iteratively increase the x_i (by 1) that decreases the total cost function the most ('biggest bang for the buck'). If no further allowed increase of any x_i results in a reduction of the total cost, we arrive at the final solution, which is always the optimal solution.

To understand that this solution is optimal, we consider why the cost function is convex. Observe that the function $(L_i - x_i)^+$ is convex in x_i for all possible values of L_i. Then, the maximum of several of those convex functions is of course also convex, so

$$\max_{1 \leq i \leq n} \{(L_i - x_i)^+\}$$

is convex in **x**. As that is convex, then the expected value of tardiness $T(\mathbf{x})$ is also convex in **x**. Finally, as **c*x** is also a convex function in **x**, and the rest of the cost function of Equation 4 are fixed parameters, the total cost function is convex. As a result, the greedy algorithm will always find the optimal solution. In Section 5.1, we evaluate the heuristic in terms of time required for increasing assembly sizes.

4.1.3. Delivery lateness distribution

We have three model parameters that we need to estimate. The most important one is the probability distribution function for delivery lateness of materials. We cover this extensively and therefore devote a separate section to it, as it is crucial in getting reliable answers from our model. Several steps have to be taken before we can actually fit and work with these data:

1. *Data removal:* We already removed irrelevant data from our analysis as described in Section 2.2.3. Before continuing with the analysis, we consider additional removals.
2. *Fitting distribution over total set:* We assume that suppliers all follow the same probability distribution, be it with other parameter values. We first try to find this distribution on the total dataset.
3. *Fitting supplier distributions:* We check if this distribution indeed holds for individual suppliers.
4. *Missing data:* Our data are incomplete for many suppliers, and especially at a material level. We consider how to cope with that issue.

Data removal

The first thing that we have to consider is if we have to further clean up our data. One thing that catches sight are the different segments, see Table 4.1. By far, most of the data concern the segment *Electronics* and *Mechanical*. It turns out that purchase orders of *General expenses* and *General expenses & IT solutions* are for internal use, so we do not have to further consider those. The last two segments (*Systems & …*) include materials that are also seen in *Electronics* and *Mechanical* but for a different purpose, e.g. maintenance. Therefore, there is no reason to remove these. The other point of concern are outliers. In the dataset as a whole we observe less than ten possible outliers. Due to the size of the dataset, these outliers will not heavily influence the general probability distribution. It is especially at a supplier level at which we have less data that an outlier may significantly change the distribution's parameters. On the other hand, if an outlier is what we observe in practice it should not be removed just like that. We decide to not remove any outliers unless it is clearly a mistake.

Table 4.1. Statistical analysis of segment data.

Segment	Count	Mean	St. dev.	Max	Min	Median
Electronics	1802	-0.5	10.1	112	-31	-1
General expenses	7	6.7	14.9	33	-8	0
General expenses & IT solutions	186	0.5	14.0	89	-74	0
Mechanical	4564	-0.2	13.8	118	-364	-1
Systems & SW engineering & …	141	5.1	29.9	97	-144	0
Systems & Equipments (incl. their …	172	-7.6	36.7	79	-179	0

Fitting distribution over total set

What we can observe from Figure 2.5 is that the delivery lateness distribution function clearly behaves differently when it is negative compared to when it is positive. The right tail is much longer than the left tail which makes perfect sense, since suppliers will not start and deliver much earlier than planned but may actually have large delays. On first sight, it seems like a negative binomial or lognormal distribution. However, those do not fit because of the extra peak on the negative side. Fortunately, we only need a distribution function of positive lateness and the expected value of total lateness (as we disregard negative safety lead times). Bearing this in mind, we decide to split the distribution in two. That means, we consider a probability of being early (< 0 workdays lateness) with a distribution, and the probability of being late (>= 0 workdays lateness) with its distribution. This may not be completely statistically sound, but it should provide a good approximation.

Another decision that we make is to use weeks lateness instead of workdays. As we can see in Figure 2.5, we are talking about large numbers concerning positive lateness so separate days do not matter much. By using weeks instead of workdays, we vastly reduce the search space and computing power required. And as calculating these safety lead times already is a brand new project, it is fine to start off with weeks. Figure 4.2 shows the distribution of these observed positive lateness values in weeks. It clearly has a long right tail with high probabilities for little lateness. This is common for distributions in the gamma family. The fitted values in the figure is that of the general gamma distribution, which seems like a good approximation. With values of shape α = 0.243 and rate β = 5.146, we cannot determine a special case of the gamma distribution such as the exponential distribution. Although the gamma distribution is a continuous distribution and we have discrete data, we approximate it by assuming that $P(X = k) = F(k + 0.5) - F(k - 0.5)$. The Chi-Square value is with 34.8 below the test statistic of 39.4 and therefore, at 95% significance, there is no significant difference between the presumed gamma distribution and the data.

Lastly, we also define a maximum delivery lateness. Theoretically, given a probability distribution, there is always *some* chance for an extremely high delivery lateness. That is never true in practice. Experts suggested that the delay probably never exceeds more than once the procurement lead time p_i, as the supplier is then able to execute the whole process twice. However, we still see delays exceeding that limit in about 1% of the observed order lines (63 out of 6872 order lines) which we consider too much. According to the data, it makes more sense to choose the limit as the minimum of twice the procurement lead time and 20 weeks. We namely only observe a delivery lateness of twice the procurement lead time 7 times, and more than 20 weeks just 2 times. We decide to use this maximum lateness as a cut-off point rather than rescaling the whole distribution function, as we expect nearly no difference in outcome and it is much easier to implement.

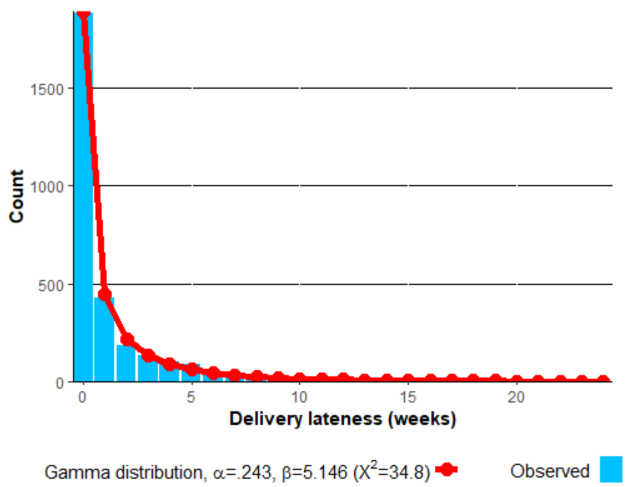

Figure 4.2. Fitting the delivery lateness distribution.

Altogether, we then determine the delivery lateness parameters as follows (using the same gamma distribution parameters as Excel):

$$\text{Expected lateness } l_i = \text{Mean of total delivery lateness}$$

$$\text{Probability of negative lateness } \pi_i^{L-} = \frac{\#\text{rows lateness} < 0}{\#\text{total rows}}$$

$$\text{Probability of positive lateness } \pi_i^{L+} = 1 - \pi_i^{L-}$$

$$\text{Rate } (\beta) \text{ parameter of gamma distribution} = \frac{\text{Variance of positive lateness } (\sigma_i^2)}{\text{Mean of positive lateness } (\mu_i)}$$

$$\text{Shape } (\alpha) \text{ parameter of gamma distribution} = \frac{\mu_i}{\beta}$$

$$\text{Maximum positive lateness } M_i^{L+} = \min\{2p_i, 20\}$$

$$G(k, \alpha, \beta) = \text{Cumulative gamma distribution, with prob. equal to or lower than } k, \text{given } \alpha \text{ and } \beta$$

$$F(k) = \begin{cases} \pi_i^{L-} + \pi_i^{L+} * \left(G(k+0.5, \alpha, \beta) - G(k-0.5, \alpha, \beta)\right), & \forall k < M_i^{L+} \\ 1, & \forall k \geq M_i^{L+} \end{cases}$$

Fitting supplier distributions

The most important thing that we want to know is whether the gamma distribution is also applicable to individual supplier cases. To test that, we took the data of 7 important Aerospace suppliers and compared a gamma distribution with their own parameters to the observed values again (see Figure 4.3). Not all fits are representing the past data very well, but it again seems to give a fine approximation. For the bottom-left supplier we only had 20 positive lateness records (others all >100), therefore we only have approximations for 0 and 1 week lateness.

Missing data

A main issue is that we do not have data for all suppliers. In fact, we lack sufficient data for about 80% of our suppliers. There are multiple things we may do to deal with this problem. First of all, we could assume the supplier is an 'average' supplier, and take the total dataset as their distribution. It certainly will not give optimal solutions, but it could prove a fine approximation to work with. An improvement would be to include knowledge of the supplier in the artificial distribution. For example, we could also make an artificial distribution of a 'bad' and 'good' supplier, and use either one on expert opinion. A more sophisticated option would be to use data such as those presented in Table 4.1 to find patterns in distribution functions and other characteristics. A quick analysis shows that there is little correlation between delivery lateness and most characteristics, but perhaps further categorisation will lead to more information.

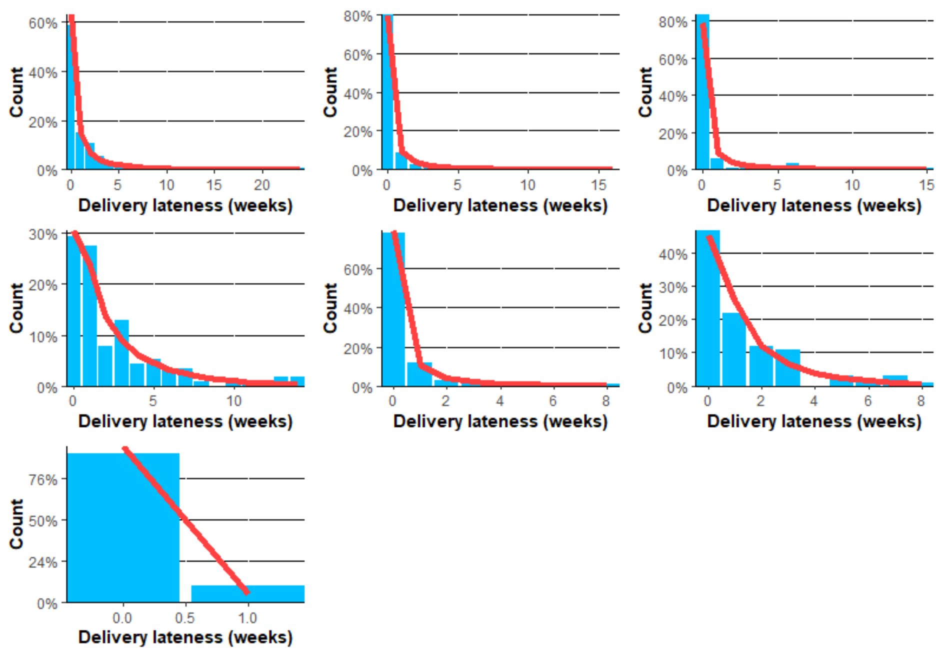

Figure 4.3. Observed (blue) vs. fitted (red) values for 7 important Aerospace suppliers.

Another problem that arises is that material delivery is not fully supplier dependent, according to expert opinion. A supplier may be just as capable of delivering 5 different materials together on-time compared to these materials individually, while struggling much more to deliver one specific material timely. At the same time, we simply do not have delivery lateness data for most materials, and if we do, there are not enough to determine reliable parameters on. Still, our model gives an advice for safety lead times, and the decision maker is assumed to be aware of the more difficult materials within an assembly. Adjustments will then always be made accordingly. Therefore, we consider each material individually with its supplier's delivery lateness data.

We could thoroughly explore the options to cope with missing data, but we choose not to at this moment. We would first like to know to what extent the model is influenced by a wrong distribution function. If it then turns out to be a major problem, we can provide further suggestions what the best actions would be. In Section 5.2, we will describe this analysis of missing supplier data.

4.1.4. Other model parameters

Other model parameters that we have to determine are tardiness costs and the holding cost rate. If production is delayed, there are numerous possible costs involved. Some examples are penalties, costs of unused capacity (or not enough capacity later), rush orders, costs of rescheduling, loss of goodwill, etc. It is very situation-specific, depending on many factors. Experts at Thales hint towards a few hundred to a few thousand euros per week, but also admit that we actually know very little of it. Therefore, in Section 5.4, we perform a sensitivity analysis to see how much these costs change the answers in general. This then provides insight in how to deal with this unknown parameter. It will also show if a study on these costs would be viable.

Concerning the holding cost rate, we choose not to consider in detail what Thales' holding cost rate might be. We discussed the MHC in Section 2.4.2, but this is just a small part of actual holding costs. Comparable companies mostly use a holding cost rate of 20-30%, and 25% is often assumed at Thales as well. Again, with the sensitivity analysis in Section 5.4, we measure how much the uncertainty in this value impacts the solutions of the model.

4.1.5. Model implementation

Building the model is one side of the coin, but we must also think of the implementation. For that, we need to answer the following questions:

- What is the role of the model in decision making?
- Who are involved with the use of the model and its outcomes?
- When and how often do we use the model?
- What changes are required to implement the model in the current processes and systems?
- How do we measure improvement resulting from the model?

Role of the model

The sole purpose of the model is to give an *advice* on safety lead times. The model makes several assumptions for which it does not fully reflect reality. Especially regarding specific materials that perform much worse or better than its supplier does in general, tailored decisions are required. Assessing all materials from their final assembly level seems to be a new approach, and calculating safety lead times (in this setting) definitely is. The optimal safety lead times resulting from the model could therefore be used as are as an experiment for the first time, as it most likely will result in an improvement (as we show in Chapter 5).

People involved

The primary user of the model is the Material Manager. At the same time, many others may be involved. Supply Chain Managers, for example, have much knowledge of specific assemblies and their materials, so their input will be vital in certain decisions. The same applies for e.g. Tactical Buyers regarding supplier performance, or the Inventory Controller if a combination of stock and safety lead time is analysed. Therefore, depending on the situation, the Material Manager will decide who else to involve.

Moment of use

It is undesirable to optimise safety lead times for each new purchase order, nor is that necessary. We suggest to evaluate the model each half year, as supplier performance can change quite a bit. For that reason, we also suggest to use supplier data of at most 2 years back as well. Furthermore, we propose to re-evaluate the model in one of the following cases:

- A change of materials in the BOM
- A switch in suppliers
- A significant change in production schedule (e.g. increasing batch sizes from 6 to 12 units per week)
- A significant change in a material's procurement lead time

Required changes to processes and systems

We are going to use safety lead times, but this option is not yet enabled in the ERP system. That is not major issue, since the option is available. Still, it requires about 3 months of testing according to the Subject Matter Expert. Then, all procurement post-processing lead times should be set accordingly, such that they represent the actual needed inbound time (i.e. expedition and possibly inspection time). Furthermore, it is crucial that delivery lateness data are registered better from now on. We observed that in some cases a buyer does not want a supplier to be 'punished' for a bad delivery that was not their fault, and then revised the promised date. We understand the good intentions, but it is essential that such changes are kept separate from the objective data on which we base quantitative decisions.

Measuring improvement

The model aims to reduce the sum of holding costs and tardiness costs. Both are difficult to measure as is. Currently, an important KPI for inventory management at Thales is the *Inventory Turnover Rate* (ITR), which indirectly relates to holding costs. We expect that the ITR increases due to the use of the model as more expensive materials should be on stock shorter. Hence, we suggest to measure the ITR of materials for a specific Aerospace assembly in a pilot case, and see if it does indeed increase. Additionally, we are interested to see if product tardiness does decrease. This may be measured by outgoing OTD, or Thales could initiate measuring average tardiness per batch. Again, we propose to also consider this in a pilot case.

4.2 Design of the sales and procurement process

We suggest additions to the process, both pre-order-acceptance and post-order-acceptance.

4.2.1. Pre-order-acceptance process

Recall from Section 2.2.2 that a typical issue may be that the given procurement processing lead time by the supplier may be outdated by the time that we actually request the material. In the worst case, we may even promise a delivery date to the customer that we cannot actualise. We believe that early communication between the CAT and Procurement as depicted in Figure 4.4 can help overcome that issue (recall legend and function areas from Section 2.1). However, it is not always quite straightforward. We concisely explain the desired process and mention the main difficulty.

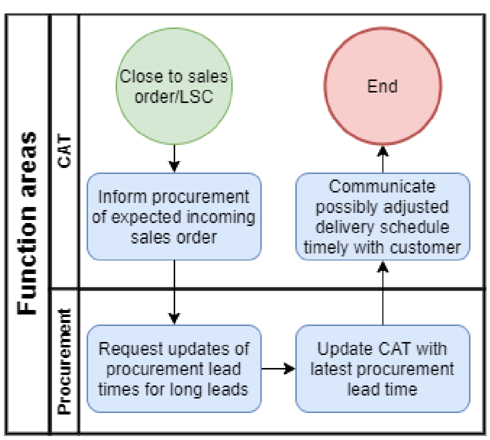

Figure 4.4. Suggested addition to pre-order-acceptance process.

The idea is that if the CAT knows or believes a sales order is close to being accepted, or an LSC/ITP will soon follow, this should be communicated with Procurement. They can then check again with the suppliers of the longest lead time materials (or more specifically the low slack materials) what procurement lead time they can now promise. Following on that, we can confirm a (possibly adjusted) delivery schedule to the customer. Now, recognising this moment is not always clear-cut. It could well be that Thales assumes an order is accepted soon, but it then turns out to take another half year. Still, we only consider the most crucial materials. Requesting an update for those few materials should not take much time, and it could definitely prove helpful. Therefore, we suggest to request such an update at any moment when there is enough confidence that acceptance will follow soon.

4.2.2. Post-order-acceptance process

Also after order acceptance or an LSC, we may have the same problem of outdated procurement lead times, but then in particular for those materials that we order later. In addition, we may encounter other problems, such as incorrect parameters or unknown information, due to which we order later than necessary to have the materials on-time. To improve this issue, we suggest a hybrid scheduling approach, rather than current backward scheduling (see Figure 4.5). This means that we would place orders once demand is known (similar to forward scheduling) but maintain a request date slightly before the need-by date (similar to backward scheduling). By sharing early demand information with the suppliers, they can anticipate on the orders in their own production activities. Also, we will notice issues from our side (e.g. missing information) earlier with less risk of late ordering.

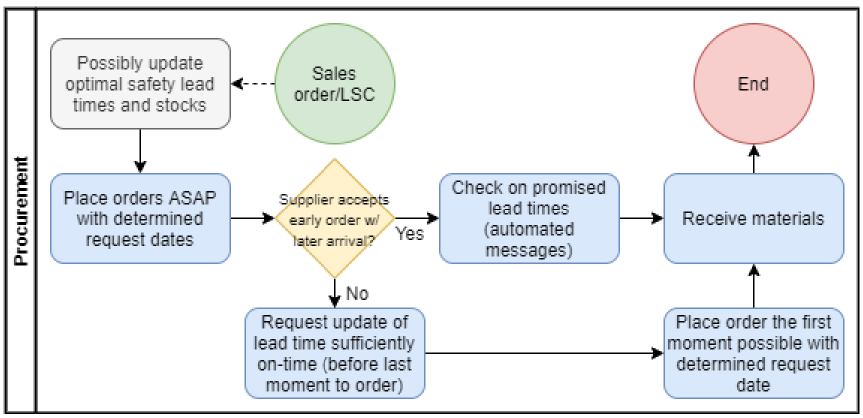

Figure 4.5. Suggested addition to post-order-acceptance process.

Since we are now placing orders earlier, it is crucial that we regularly check with the supplier if they can actualise their promises. Automating these checks (by automated messages to suppliers) should significantly reduce time spent by the buyers. The frequency of such checks is to be determined, but it should have the right balance of giving sufficient time to act without practically spamming suppliers. We do recognise that not every supplier will anticipate on early information. Some suppliers refuse to make promises on the longer term, as they may also receive possibly larger (more important) orders in the meantime, which could disturb the planning. An improvement in that case would be to at least check regularly with the supplier what the actual procurement processing lead time, and still try to place the order slightly earlier than required.

As a result, the promised arrival date will more often match the request date. Moreover, we might even assume that the orders have less delivery delay, as the suppliers received earlier notice. In other cases, Thales saw that suppliers actually delivered earlier than preferred as well, resulting in higher holding costs for Thales. It is very important that the request date is agreed on, else we may also have to resort to having an up-to-date lead time from the supplier and still place the order just-in-time. In any case, correct information is crucial.

Obviously, this process is not flawless and there will always be reasons why materials cannot be ordered earlier, suppliers causing disturbances, etc. However, it is an important step towards less nervousness in material planning. Combined with the following model, we can expect good improvement.

4.3. Total Cost of Ownership (TCO) model

This section covers the TCO model that we built. It is independent of the earlier proposed solutions, and rises from the need to quickly have insight whether we can profit from purchasing in higher numbers with related quantity discounts. It is a simple tool that determines, given a known demand now and a known demand in the future, the right balance between quantity discounts (now and then), fixed costs of ordering (now and then), an expected price increase and holding costs. No uncertainty, not many periods with many possible decisions, just allowing a quick advice. As it is more straightforward than the safety lead time model, we briefly give the model's notation now and explain after that.

We define the following decision variables that are optimised:

x^N : Number of materials to purchase now (units)
y^N : 1 if we place a purchase order now, 0 otherwise.

Next to that, we have the following parameters:

D^N : Demand now (units)
D^F : Demand in the future (units)
T : Time until future demand (years)
h : Holding cost rate (% of unit price per year)
O : Fixed costs of ordering (€)
q^0 : Minimum order quantity (units)
c^0 : Unit price at minimum order quantity (€)
q_j^D : Quantity discount bracket break for bracket j (units)
c_j^D : Quantity discount bracket unit price for bracket j (€)
α : Expected price increase (% of unit price per year)
M : Large positive number.

Then, we have the dependent variables:

c^N : Unit cost price now (€)
E^N : Excess stock after purchase and demand now (units)
x^F : Number of materials to purchase in future given future demand and purchase now
y^F : 1 if we place a purchase order in the future, 0 otherwise
c^F : Unit cost price of future purchase (€).

The objective function is:

$$\min\{x^N c^N + x^F c^F + (y^N + y^F)O + hE^N c^N T\}. \tag{15}$$

Finally, we have the following constraints:

$$x^N \geq D^N \tag{16}$$
$$E^N = x^N - D^N \tag{17}$$
$$x^F + E^N \geq D^F \tag{18}$$

$$x^N, x^F \geq q^0 \tag{19}$$

$$c^N = \begin{cases} c^0, & x^N < q_1^D \\ c_j^D, & q_j^D \leq x^N < q_{j+1}^D \text{ and } \forall j < J \\ c_j^D, & x^N \geq q_j^D \text{ and } j = J \end{cases} \tag{20}$$

$$c^F = \begin{cases} c^0(1+\alpha)^T, & x^F < q_1^D \\ c_j^D(1+\alpha)^T, & q_j^D \leq x^F < q_{j+1}^D \text{ and } \forall j < J \\ c_j^D(1+\alpha)^T, & x^F \geq q_j^D \text{ and } j = J \end{cases} \tag{21}$$

$$x^N \leq y^N M \tag{22}$$

$$x^F \leq y^F M \tag{23}$$

$$y^N, y^F \in \{0,1\} \tag{24}$$

$$x^N, x^F \in \mathbf{Z}^+ \tag{25}$$

Constraints (16), (17) and (18) ensure that current and future demand both are met. The minimum order quantities are given by Constraint (19). Constraints (20) and (21) show the cost price given quantities purchased, and potential price increases in case of future demand. With constraints (22) and (23) we ensure that fixed costs of ordering are included if we place a purchase order, and finally constraints (24) and (25) give some numerical limitations. As there are very few options to consider (POs between 0 and say max a few hundred), we solve the model by full enumeration.

There are other considerations that the user has to make. These are for example the revision risk of the material, or whether warranty plays a role in the longer period of holding materials on stock. Also, by purchasing for future demand, we might significantly reduce customer lead time (if it concerns a material on the critical path) and eliminate risk of delivery lateness for that next moment. If the material is one we experience many difficulties with, it would be even more profitable to purchase in quantities. In other words, this tool solely assesses the financial aspects of higher order quantities with two moments of demand. We do not further evaluate the model as with the safety lead time model, but refer to it with classification policies in Section 4.4.3.

4.4. Material classification

As Thales manages over a hundred thousand items, the need to quickly assess a material is high. Such an assessment should ease answering a question such as whether to stock a material or not, or if we should expect issues with an item – perhaps following with further analysis if unexpected problems emerge. To aid in answering these questions, we design a classification method for buy materials with a focus on material availability. First, we discuss which parameters to include in material classification. Then, we describe how we translate a combination of parameters to material categories. Subsequently, we briefly touch policies. We bear in mind that it is essential to keep classification simple, as it should improve decision making rather than complicate it.

4.4.1. Choosing parameters

When considering parameters to include in classification, we value both a certain measure of criticality (recall definition from Section 2.4.1) important as well as parameters that relate to risk. In Section 2.4, we bespoke several material characteristics that could prove useful here. In consultation with colleagues, we decided that the parameters to classify materials on are as follows:

- Slack to critical path
- Material costs
- Future demand
- Revision risk

Slack is an important discriminator of criticality in the PTO assembly setting. In line with the model and process which we designed, we know that most materials have sufficient slack such that safety lead times can cope with availability issues. Therefore it also clearly shows which materials possibly need stocks to fulfil the same purpose, and where we might expect problems. Whereas slack shows us criticality, the costs of a material are needed to assess the costs and therewith the risk of (stocking) materials. To be more exact, we use the material costs per product or system as described in Section 2.4.2. Furthermore, we incorporate whether any demand is expected for a material, as we of course do not want to take actions for materials we do not plan to use (soon). Finally, revision risk is also important, as we do not want to invest in materials that are expected to become obsolete soon.

Other options that could prove useful but which we did not include are delivery lateness and commonality. Delivery lateness embodies risk and criticality perhaps even more than slack does, and the combination of both is ideal. However, as we mentioned, we lack data for most suppliers and there may be considerable delivery lateness differences between multiple materials of one supplier. Therefore it is not possible to include. Commonality may be useful outside Aerospace, where many (smaller) materials are used in multiple products. Such materials have a lower risk of e.g. becoming obsolete, while a lack of those materials impacts more projects. Nonetheless, Aerospace hardly has commonality, so we skip that parameter now.

4.4.2. Creating categories

Creating categories is a mix of ABC classification as described in Section 3.3 and own insight. For material costs, we use the Pareto approach to roughly get a 20-30-50 division. Initially we considered doing the same for slack, but we concluded that having one cut-off point (two slack categories) better fits the purpose. The reason for that is that the slack parameter should show whether a material is possibly critical or not. So instead, we consider two categories: one with a realistic risk of delivery lateness being higher than slack, and one without. That means that we could anticipate on delivery lateness issues with safety lead times in the latter category, while we are limited in doing so in the first category. We saw in Section 4.2.3 that 20 weeks is one cut-off point for the delivery lateness distribution function. However, with only 2 out of 6872 order lines that exceed that number, we think it is too strict for classification. Therefore, we choose a boundary of 15 weeks, of which we still only saw less than 0.5% of the order lines exceed it.

For the Pareto analysis of material costs we look how these costs are distributed. Figure 4.6 shows this analysis, where the black line represents the cumulative distribution of all products combined. Based on that figure, we decided to have material cost categories of less than €10, between €10 and €250, and more than €250. These are tangible numbers, and €250 nicely coincides with the loose soft/hard pegged rule described in Section 2.5.2. As we can also see in Figure 4.6, the distribution for the different products follow the average line fairly well. Therefore, these categories would be fine for all products. Table 4.2 presents a distribution summary of the two categories combined.

It is important to note that we evaluate slack and material costs from a product/system perspective. Therefore, the same material may have different values for those parameters, depending on the product/system. For material costs, this effect is negligible as only a few very cheap materials have different quantities per system resulting in differences such as €0.20 per system instead of €0.33. At slack we do see differences of e.g. 50 weeks at one product and 20 at another, but never such that it the same material is in one classification category for one product and in the other for a second product. Then again, the current approach is to classify materials per product/system so it is not an immediate concern. If a global overview of all unique materials for multiple products/systems is required and we want to take this matter into account, we could for example take the minimum value of slack for the material as part of all products.

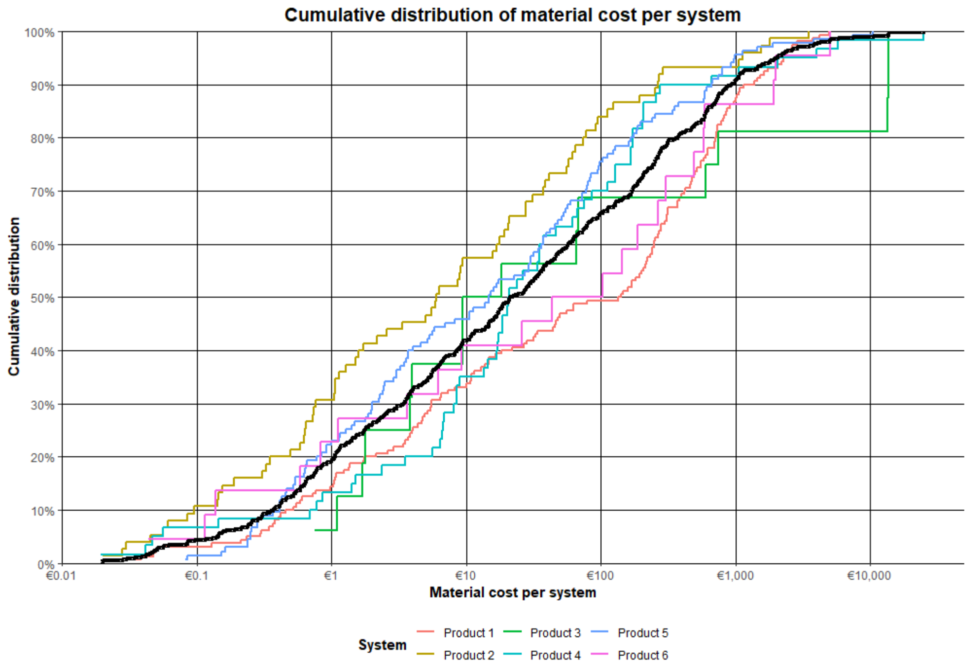

Figure 4.6. Cumulative distribution of material costs per system (costs are on log[10] scale).

Table 4.2. Slack and material cost classification categories.

		Slack (weeks)		
		0-15	*>15*	***Total***
Material costs (€)	*>250*	7.6%	16.8%	**24.4%**
	10-250	1.9%	31.4%	**33.3%**
	0-10	0.2%	42.0%	**42.3%**
	Total	**9.8%**	**90.2%**	**100%**

4.4.3. Policy suggestions

Based on the categories which we designed, we give several policy suggestions. These are not quantitative advices but rather suggestions what should be done with a material given its criticality and risk. The policies that we include are to stock items at Thales (and roughly to what extent), to stock the (raw) materials that our suppliers require at their location, purchase higher quantities, or do nothing extra. Figure 4.7 gives a visual representation of how such a classification matrix would be, where we now filled in the materials of Product 3 to further clarify (with dummy data for revision risk and future demand). In this case, we would propose the following policies:

- *Policy 1:* Critical materials with low costs. No-brainer: Stock abundantly, e.g. for months of work. Risking stock-outs for those materials never make up for inventory costs.
- *Policy 2:* Critical materials with medium costs. Possibly stock some materials, e.g. for a few weeks of work. Final decision depends on experience whether the material is likely to be delayed or not. Use of TCO model to possibly order in higher quantities has double value here, with reduced risk of lateness and purchase cost reduction (naturally TCO could be valuable for any €10+ material).
- *Policy 3:* Critical materials with high costs. Consider options of stocking (raw) materials at the supplier, to shorten the procurement lead time. Should still be affordable for a reasonable stock, as having just a few (raw) materials at stock at the supplier does not help much. Decisions on own insight.
- *Policy 4:* Non-critical materials with low costs. Not too much attention should have to be spent on these materials. If we see reasonable quantity discounts, stocking by purchasing less often (less fixed costs of ordering) in higher quantities will generally outweigh the holding costs.

➢ *Policy 5:* Non-critical materials with medium to high costs. Remain the current PTO strategy, including other advices given in this report. Give least attention to this category.

We repeatedly say to build stock, but we want to make a critical note here. It is very important to understand for which uncertainty we are building stocks. If we are building a safety stock for uncertainty in procurement lead time, we often only need enough stock to produce for a few weeks to cope with expected delivery lateness. When we consider shrinkage, we should order slightly more than demand to make up for defects. Only if demand is (very) uncertain and we want to react rapidly, larger stocks may be justified. In practice, we often see safety stocks that are much larger than required to deal with any of the first two issues, and uncertain demand is a minor problem for most materials in the long-term planning. We should still be placing purchase orders timely, such that by far most material demand can be fulfilled by those orders. In terms of not ordering more than required, one could also approach a safety stock as follows. For the first batch, we could already order materials for the last batch or last few batches. This way, we have a stock for contingencies throughout later batches, and in the end we have nothing extra when materials are used.

Another remark is that this classification tool is mainly a rough idea at this point. It would be interesting to also use it when materials are actually stocked out. If these materials are in *Policy 1*, *4* or *5*, we should thoroughly check why. Perhaps there are more (important) measures of risk that should be added to the classification. Also, different policies should be explored. For example, increasing the FDS for *Policy 4* materials automatically results in larger batches and fewer orders. But it also has it downsides with regards to new demand in the fixed supply period. Nonetheless, this proposal clearly shows a framework for classification, and highlighted important measures of criticality and risk.

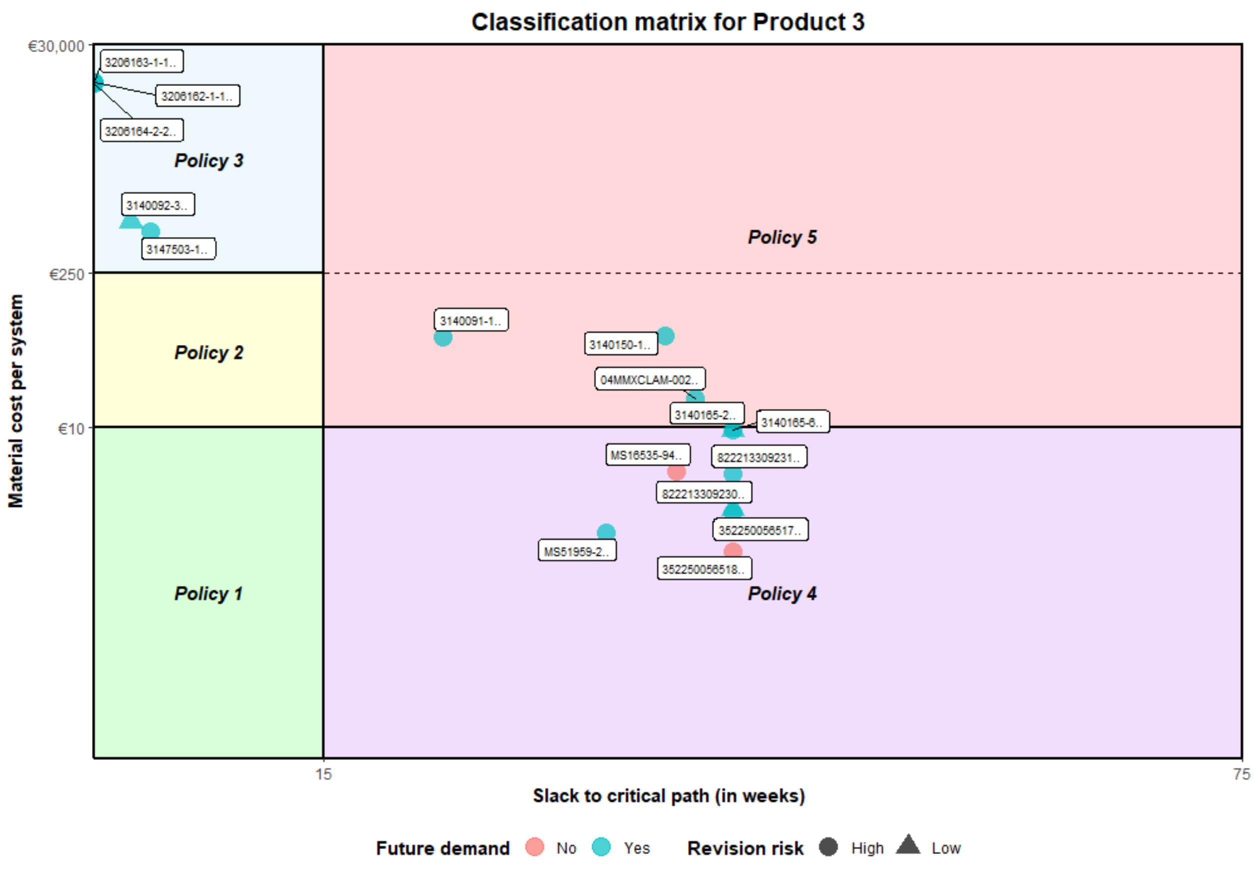

Figure 4.7. Example classification matrix, filled in with materials of Product 3.

4.5. Conclusion

This chapter describes the proposed solutions: dealing with supply uncertainty through a safety lead time optimisation model, process changes and a material classification method, and allowing overall cost reductions with a Total Cost of Ownership model. First, we design the safety lead time optimisation model. We consider the context in literature, mathematical notation and the solution method, fitting the delivery lateness distribution, uncertainty of parameters and implementation of the model. We are able to use a model from literature for the largest part. Unfortunately, we are not able to successfully extend the model with safety stocks or multi-periodicity. Still, we see a promising model that is easy to solve to optimality. Furthermore, we are able to fit a gamma distribution on the delivery lateness data, which fits well on both the total dataset as individual supplier data. We recognise that we do not have data for many suppliers for which we proposed several techniques, but we first analyse the impact of missing data in Chapter 5. We then also consider how uncertainty in holding cost rate and tardiness costs impact the model's solutions. Finally, we elaborate on the implementation of the model.

We also suggest additions to the communication process of sales and procurement. The safety lead times optimisation model considers delivery lateness, but supplier underperformance is not the only reason why materials are later than required. These process changes aim to improve late materials caused by Thales, by earlier and more communication between both sales, procurement and suppliers. Furthermore, we design the TCO model. It is kept simple, and allows insight in the total costs given current and future demand and an order quantity to purchase now.

In addition, we design a material classification method. Together with colleagues, we decided that important parameters to classify on are slack to the critical path, material costs per system, revision risk and forecasted demand (whether we expect *any* demand in the next year). We create the classes by a combination of the ABC classification method and binary categories (low/high, yes/no). We suggest five policies for now, which mainly focus on using safety stocks and the TCO model. The classification method is only a start, however, and we conclude that it would be interesting to also use it when materials are actually stocked out. By checking if these materials fall in the categories where we would expect contingencies (or not), we may adjust categories or add parameters. Additionally, we propose to further expand or even quantify policies.

5. Performance evaluation

This chapter describes a further analysis of the results from our safety lead time optimisation model. First, we evaluate our heuristic in Section 5.1. In Section 5.2, we consider how to deal with missing supplier data. Section 5.3 covers the improvement that we may expect from the model. In addition to the model's results, we also use a Monte Carlo simulation to get more reliable answers. Subsequently, in Section 5.4 we perform a sensitivity analysis to observe how our model reacts to uncertainty in certain parameters. Furthermore, Section 5.5 covers a trade-off between extending customer lead time and allowing more safety lead times for low slack materials. Finally, Section 5.6 concludes on the results.

Throughout this chapter, we base our evaluations on four systems. These include three systems from Aerospace and one radar system. The Aerospace systems are both systems that are already manufactured for a long period (>10 years) and a relatively new one. Our focus lies on Aerospace, but we at one point also consider the radar system in order to show what we may expect there. Table 5.1 gives an overview of these systems. The batch sizes and number of batches that we analyse are more or less what we currently have planned now, or what relevant numbers would be. At the radar system, we consider the production of singles. Note that these products do not align with names given earlier in this report.

Table 5.1. Evaluated system characteristics.

System	***Product 1***	***Product 2***	***Product 3***	***Product 4***
Product chain	Aerospace	Aerospace	Aerospace	Radar systems
Number of MRP buy materials	12	66	22	1,116
Number of unique suppliers	5	19	8	>70
Procurement lead times (min/max)	6 / 48 weeks	4 / 45.6 weeks	3 / 63 weeks	2 / 68.2 weeks
Batch size	8	6	10	1
Batch cycle time	1 per 4 weeks	1 per 1 week	1 per 4 weeks	-
Number of batches	12	49	14	1

5.1. Heuristic evaluation

The heuristics that we describe earlier are a simple necessity. In every assembly, we can for each material have any integer value of safety lead time between 0 and at most 20 weeks (the maximum of either slack or maximum lateness, see Section 4.1.3). In the smallest real scenario we have only 12 materials to evaluate, which in that case already results in over 6 billion different combinations to consider (while accounting for constraints). Since the model includes an integral over lateness of all materials, the calculations also become increasingly complex. It is therefore not possible to fully enumerate any reasonable assembly setting.

We already know that the cost function is convex, which fully justifies a greedy algorithm. However, we would still like to know how fast our heuristic calculates the optimal safety lead times, in order to get an idea what assembly sizes can be considered in reasonable time. Since we do not have many different assemblies to gather data on, we create virtual assemblies with material parameters randomly sampled from realistic ranges that we see in practice (see Table 5.2 for these parameter ranges). We do not include a joint distribution of parameters in these samples (e.g. a higher variance of positive lateness if we draw a higher mean of positive lateness). The correlation between the parameters from Table 5.2 are sufficiently low (absolute correlation values of 0.05-0.3) to disregard that given the information that we seek. Furthermore, we experiment for virtual assembly sizes including 10 to 250 materials with increasing steps of 10 materials at a time. We repeat the experiment for 10 times, with new randomly sampled materials each experiment.

Figure 5.1 shows the results of the run-time evaluation. We compose the total run-time of the number of iterations required for the final solution, and the run-time per iteration. One iteration is seen as a full cost evaluation of increasing any material's safety lead time by a week, and then choosing the best material to increase. We see that the number of iterations is linear with the assembly size, namely about 8-9 iterations required per material. The time per iteration is slightly exponential but mainly linear with the assembly size.

It ranges from circa 0.0007 seconds per iteration per material at a size of 10, to about 0.00085 seconds per iteration per material at a size of 250.

Table 5.2. Material parameter ranges (only these parameters are needed for random sampling materials as input for the experiments, rest automatically follows)

Parameter	Notation	Range
Procurement processing lead time	p_i	~U[5, 50]
Material costs per product	c_i	e^x with x~U[-1,10] (for relatively more cheaper items)
Expected lateness	l_i	~U[-3, 1]
Probability of negative lateness	π_i^{L-}	~U[0.15, 0.70]
Mean of positive lateness	μ_i	~U[0, 3]
Variance of positive lateness	σ_i^2	~U[0, 15]

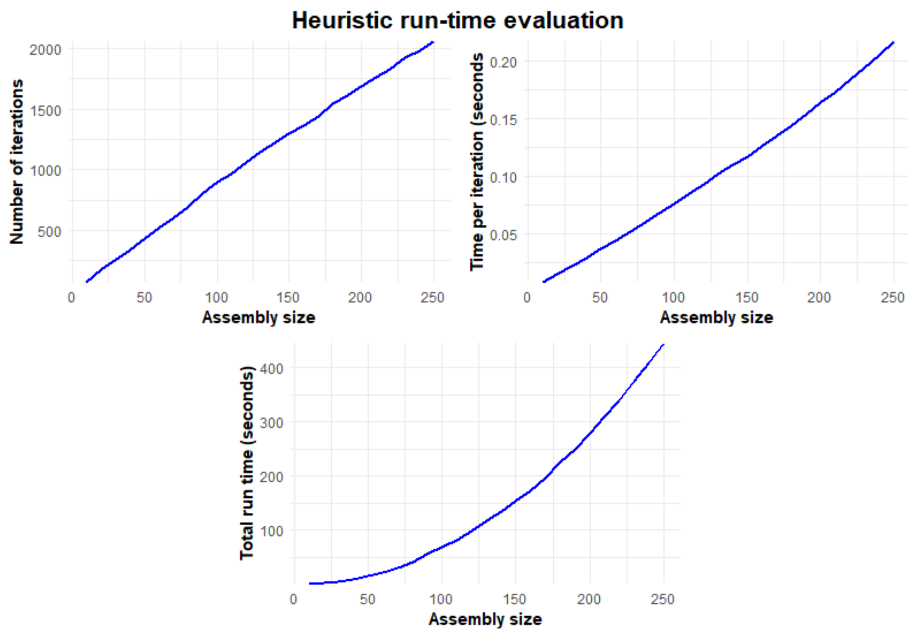

Figure 5.1. Heuristic run-time evaluation. Number of iterations (top-left corner), time per iteration (top-right corner), total run-time (bottom).

By using this decomposition method, we can also determine the expected run-time for larger assemblies by extrapolation. For example, with 1,000 materials we expect about 8,000 to 9,000 required iterations. If we extrapolate the increase of time per iteration per material, we expect roughly 1.5 second per iteration, with a total run-time of more than 3 hours. With a size of 2,000 materials the number of iterations again double, while the time per iteration exponentially increase to 5.5 seconds resulting in almost 27 hours total run-time. These exact numbers of course depend on the used hardware, which in this case was a fast laptop with an 8th gen i7 processor and 16GB internal memory. Then again, as we now know the ratios, we can find it for any configuration once we know run-times for lower assembly sizes. We expect that for any configuration, optimising the model for large radar systems (often far more than 1,000 materials) cannot be run quickly. Reducing the final assembly to smaller subassemblies is therefore crucial in that case. For Aerospace products, we do not foresee any problems.

5.2. Dealing with missing supplier data

We note in Section 2.2.4 that we do not have (sufficient) data for most of the suppliers. In this section, we analyse to what extent that is an issue. Although it is actually part of the sensitivity analysis, we need to determine how to deal with lacking supplier data first in order to run results. Here we test what the impact is of using the whole dataset of purchase order lines (all 6872 observations, recall from e.g. Section 2.2.4) in case we do not have (enough) data for a supplier. By doing that, we assume that the supplier performs exactly average given the data that we have.

Fortunately, for all suppliers that provide materials of *Product 1* and *Product 3*, we have enough data. Therefore, we can do the following. We first solve the model using the *actual* supplier data (so truly observed values for that supplier). Then, we again solve the model but now using the *total* dataset of all suppliers as input for the distribution functions. Subsequently, we observe if and how much the resulting safety lead times in both cases vary, by filling the optimal safety lead times of both solutions in the model with the actual supplier data. We can then compare the expected total costs. We use 4 different base settings in this comparison, so we can observe the differences in multiple scenarios. Table 5.3 shows these base settings, and Table 5.4 gives the results of the analysis.

Table 5.3. Test base settings for sensitivity analysis. First two cases are most realistic, while the last two cases serve to evaluate more extreme cases.

Base setting	Holding cost rate (yearly)	Tardiness costs (weekly)
Case 1	20%	€0
Case 2	25%	€2,000
Case 3	25%	€5,000
Case 4	30%	€8,000

Table 5.4. Expected total costs given optimal solution resulting from using either actual supplier data or total supplier dataset. Differences are negligible.

Base setting	Product 1 – actual data	Product 1 - total data	Product 3 - actual data	Product 3 - total data
Case 1	€50,208	€50,388	€18,192	€18,624
Case 2	€103,740	€103,776	€55,704	€57,228
Case 3	€164,712	€164,736	€85,944	€88,620
Case 4	€238,260	€238,296	€120,636	€124,404

The results that we see are very clear. Using the total dataset hardly affects the total costs that we may expect. There are some small differences in the safety lead times, but these seemingly have a minor impact. By using the total dataset, there is always some probability of a material being more than 10 weeks late. Therefore, cheap materials will always get a high safety lead time because it pays off. The largest difference that we see accordingly is that cheap materials have a higher safety lead time than 'required', but this effect is negligible on the total cost function. On the other hand, expensive materials will in both cases have little safety lead times, as it is more costly. Altogether, material costs play a much bigger role as the differences can be very large, compared to the differences in distribution functions with the actual and total dataset. And in the end, the model gives an advice but changes for specific materials will be made on experience. Therefore, we conclude that using the total dataset of all suppliers as input – if insufficient data are available for a certain supplier – is a good option.

We immediately apply that knowledge in the next section, where we use the total dataset for missing supplier data of *Product 2*, and use the total dataset for all suppliers of *Product 4*. The reason for the latter is that for most suppliers the data are missing, so retrieving the data for the few suppliers of which we do have data is not worth the effort nor will it improve insight or the solutions.

5.3. Expected improvement

The most important question remains: What improvement may we expect from the implementation of this model? We measure these improvements two-fold. First, we consider the costs and tardiness that the model gives to get an initial idea. However, the model assumes independency of production and delivery batches and ignores batching through the FDS. In practice, this plays a significant role. Therefore, as mentioned in

Section 4.1.1, we also created a Monte Carlo simulation. Its goal is to assess how accurate the model represents the more realistic setting.

We do not know the values of holding cost rate and tardiness costs, but our model of course needs that input. For this analysis, we use a holding cost rate of 25% as that is mostly used at Thales. For tardiness costs, we decide to analyse both for no tardiness costs (*Setting 1*) as €2K tardiness costs per week (*Setting 2*). The first is especially interesting, as we will prove that we can already significantly reduce the holding costs alone. The latter is picked somewhat average given the range that we bespoke in Section 4.1.4. Section 5.4 investigates the sensitivity of the model with regards to these two uncertain parameters. In case we lack sufficient supplier data, we use the total data as input as described in the previous section. Therefore, the results are not entirely true, but we also saw that it has minor impact on what we could expect in reality. The relevant results to measure are the expected total costs (for all batches together) and expected average tardiness per batch. We compare the model's optimal solutions with those resulting from the current post-processing lead times (minus actual time required for inbound).

5.3.1. Model results

Figure 5.2 and Figure 5.3 respectively show the expected total costs and expected average tardiness per batch that the model has determined. For all products in both setting, we see a significant reduction of expected total costs ranging from 18.4% to 56.7% (percentual improvements are further discussed in Section 5.3.2). Even if we do not include tardiness costs, the model indicates that we would still benefit heavily. Tardiness also drops heavily except for *Product 1*. The reason for that is that the current safety lead time assigned some safety lead time to materials with zero slack, which our model restrains from doing so.

We observe from the optimal solutions that cheap materials are mostly assigned high safety lead times (often maximum), as their holding costs for many extra weeks outweigh even the slightest chance of tardiness. This is a very practical insight that is perfectly logic, but not always applied. It emphasizes the need to look at the bigger picture (i.e. the concerned (final) assembly) when considering material availability or the desired service level of a material. It is also in line with observations from the previous section, where we see that the material price has a much larger influence on safety lead times than the parameters of delivery lateness.

Obviously, we are content with these outcomes. However, we place a critical note that these results are not so realistic for the Aerospace products for reasons we cannot influence or investigate. First, these costs are given that materials are ordered exactly to plan. In practice, this may differ for all sorts of reasons (e.g. MOQs or buyer's own insight). That means the current costs are actually *different*, although that does not necessarily mean the improvement is *less* than what we see now. Second, we have multiple assembly levels – especially at *Product 4* – that are now modelled as one due to our model structure. In general, this underestimates tardiness and therefore even higher safety lead times are required for optimality. Then again, this also implies that the relative improvement from our model could be even larger. Third, the distribution of delivery lateness of a specific material may vary compared to other materials of the same supplier. Still, we saw that the material price has a much larger impact on safety lead times so we do not expect too much difference.

In addition, remember the gap between the request date and promised arrival date from Section 2.2.2 (in particular Figure 2.4). The model does not include that now, as we do not have reliable data to incorporate that issue. It can definitely change optimal safety lead times and costs. However, if the processes are more in control (possibly helped by implementing the suggested process changes from Section 4.2), we can still come to the point where mainly suppliers cause delivery lateness. Then the results become much more valid. Furthermore, we consider batches independent from each other which is not the fact. Especially in the case of *Product 2* where a new batch should start each week, these can and will overlap. Fortunately, the Monte Carlo simulation helps in analysing this situation, which is the subject of the following section.

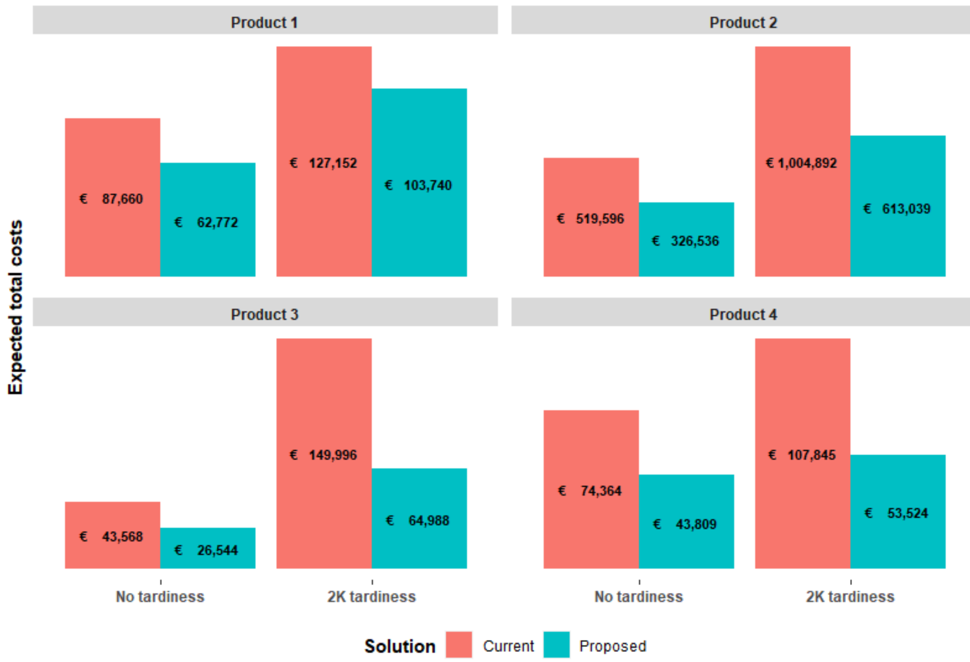

Figure 5.2. Expected total costs given by mathematical model for all tested products, settings and solutions. For all cases, we see a large decrease of expected total costs.

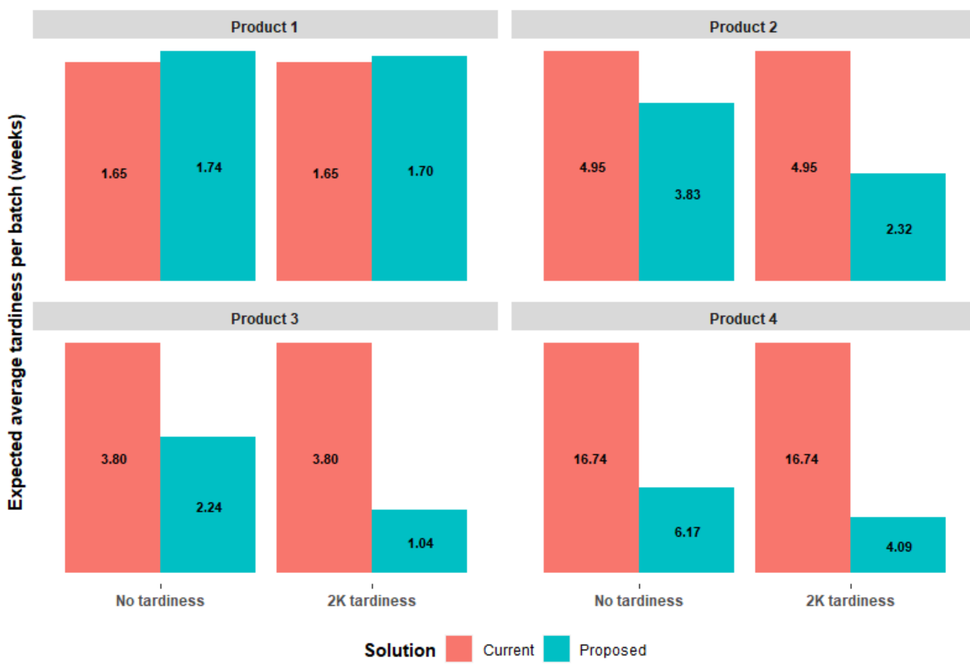

Figure 5.3. Expected average tardiness per batch given by mathematical model for all tested products, settings and solutions. Tardiness decreases heavily for all cases but Product 1.

5.3.2. Monte Carlo simulation

We described in the previous section that the model includes only one period, while we have a multi-period setting in reality. We can say for certain that total costs are higher, as batch delays stack over subsequent batches, and FDS settings result in higher holding costs. Although we did not succeed in including multiple periods in one mathematical model, we can still observe what results we can expect in the more realistic setting through a Monte Carlo simulation. In this simulation, we try to model the process of material arrivals, production activities and safety lead times and then simulate procurement lead times for numerous times to analyse how the system behaves. While mathematics can become very difficult to find exact answers with certain assumptions or restrictions, it is much more easier to describe these relations (see Appendix D for details) and see what happens when we draw random lead times. The main points that we could not capture in the *mathematical model* (referred to as MM) (recall from Section 4.1.1) but can include in the *simulation model* (referred to as SM) are the following:

1. *All materials have an FDS of 30 workdays*. Recall the FDS from Section 2.2.2. Due to that FDS, materials are purchased in larger quantities (for multiple batches) and less frequently. The simulation model does include this. In practice, even larger purchase orders may be placed due to MOQs or on buyer's own insight, but we exclude that.
2. *Workforce capacity is not infinite*. In practice, if we have a delay at one batch, most of the following batches are also delayed. At the same time, purchase orders are not often postponed to arrive with the new production schedule, as they are often order a long period earlier and it requires much effort. In our model, we assume that the delay at each batch also delays the possible start date (and delivery) of the following batches as well. The materials are still planned to arrive at the moment the batches were planned to start.
3. *Purchase orders may not cross in time*. This is always the case in practice. We disregard capacity issues at suppliers, which means that when a purchase order is much later its size may already include the next purchase order quantity as well.

Now, there are some differences in how we measure the results compared to the MM. First, as we now simulate, we include confidence intervals to the results. We run 10,000 replications (10,000 times randomly drawn lead times for all material arrival batches) for each configuration. The mean values of costs and tardiness resulting from these replications is what we may expect, but confidence intervals also show to what extent these numbers may vary. Next to that, we do not examine *Product 4* in this analysis. We consider these as separate batches of 1 which do not interfere which each other, so the SM does not add value compared to the MM here. Figure 5.4 shows the expected total costs and Figure 5.5 gives the expected average tardiness per batch.

Most importantly, we again see a significant reduction of total costs for all products and settings. Two other things also catch sight: total costs and tardiness are indeed (much) higher in the SM than observed in the MM, and the relative improvements seem quite similar in both models. We further analyse the last point in Table 5.5. This shows that the relative improvement is always (a bit) lower in the SM results than in the results of the MM, since costs in both the current and the proposed solution are much higher and therefore the difference relatively less. The differences are larger when tardiness costs are not included. Fortunately, both results never contradict each other: If the MM determines improvement, this is also the case for the SM.

We would also like to know how much higher the total costs are in the SM compared to the MM. Table 5.6 shows the ratio of SM to MM costs. The MM heavily underestimates the total costs, especially at the proposed solutions where the SM costs are always at least twice as high as the MM costs. We believe that the mean reason for this underestimation is that the model does not include FDS (which increases holding costs) and the stacking product tardiness with multiple batches. Slightly the same ratios also apply for tardiness, which we see by comparing Figure 5.3 with Figure 5.5. So, what can we conclude from all of this? Although the relative improvements are actually smaller in the more realistic case (SM), the actual total savings for the first three products are at least as much (~€235K vs. ~€276K with no tardiness costs) or much larger (~€500K vs. ~€730K with €2K tardiness costs) since the MM underestimates total costs. Even in the worst case our total costs will drop by more than 20%, while tardiness also decreases severely.

Figure 5.4. Expected total costs resulting from Monte Carlo simulation. Mean costs are given in bold in the middle, with a 95% confidence interval lower and upper bound at the top of each bar. We again see that total costs decrease in all cases with the proposed solution.

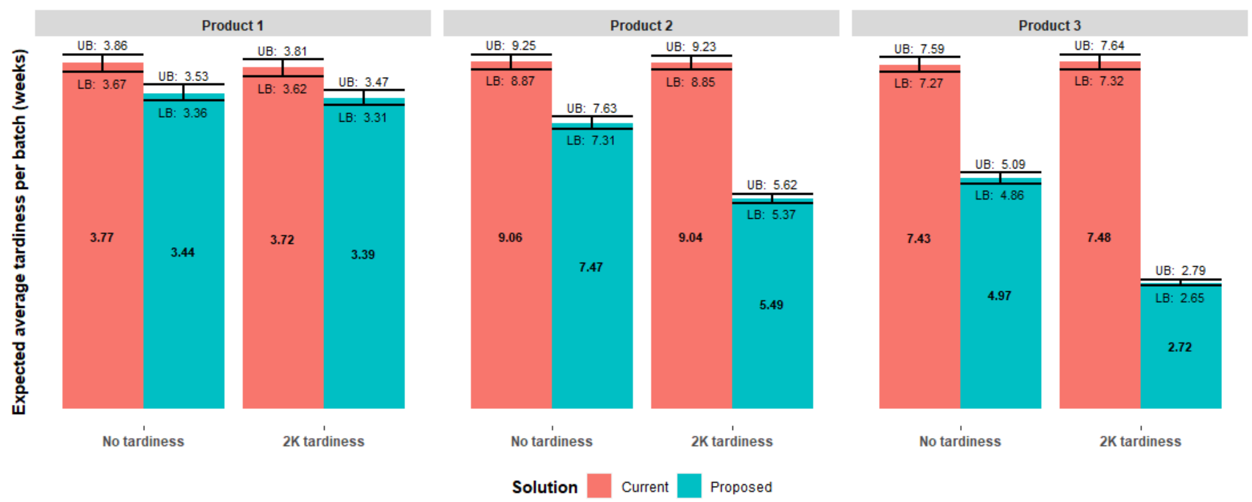

Figure 5.5. Expected average tardiness per batch (weeks) resulting from Monte Carlo simulation. Mean tardiness is given in bold in the middle, with a 95% confidence interval lower and upper bound at the top of each bar. Contrary to the MM, we also see tardiness reduction at Product 1.

Table 5.5. Expected total cost reduction per product, model and setting.

	Solution	Product 1	Product 2	Product 3	Product 4	Total
Math. model: no tardiness	***Current***	€87,659	€519,576	€43,564	€74,364	€725,163
	Proposed	€62,772	€326,549	€26,542	€43,809	€459,671
	Savings	**€24,887**	**€193,027**	**€17,022**	**€30,555**	**€265,492**
	%	**28.4%**	**37.2%**	**39.1%**	**41.1%**	**36.6%**
Math. model: €2K tardiness	***Current***	€127,150	€1,004,893	€149,997	€107,845	€1,389,886
	Proposed	€103,742	€613,024	€64,993	€53,524	€835,283
	Savings	**€23,409**	**€391,870**	**€85,004**	**€54,321**	**€554,603**
	%	**18.4%**	**39.0%**	**56.7%**	**50.4%**	**39.9%**
Sim. model: no tardiness	***Current***	€210,560	€905,448	€89,725	-	€1,205,732
	Proposed	€176,916	€686,870	€65,597	-	€929,382
	Savings	**€33,644**	**€218,578**	**€24,128**	**-**	**€276,350**
	%	**16.0%**	**24.1%**	**26.9%**	**-**	**22.9%**

Sim. model: €2K tardiness	*Current*	€298,466	€1,790,511	€299,615	-	€2,388,592
	Proposed	€258,157	€1,258,232	€142,699	-	€1,659,088
	Savings	**€40,309**	**€532,279**	**€156,916**	**-**	**€729,504**
	%	**13.5%**	**29.7%**	**52.4%**	**-**	**30.5%**

Table 5.6. Ratio of SM to MM total costs results. In all cases, we see that the mathematical model absolutely underestimates total costs.

	Model	*Product 1*	*Product 2*	*Product 3*
Current solution: no tardiness	*MM*	€87,659	€519,576	€43,564
	SM	€210,560	€905,448	€89,725
	Ratio	**2.40**	**1.74**	**2.06**
Proposed solution: no tardiness	*MM*	€62,772	€326,549	€26,542
	SM	€176,916	€686,870	€65,597
	Ratio	**2.82**	**2.10**	**2.47**
Current solution: €2K tardiness	*MM*	€127,150	€1,004,893	€149,997
	SM	€298,466	€1,790,511	€299,615
	Ratio	**2.35**	**1.78**	**2.00**
Proposed solution: €2K tardiness	*MM*	€103,742	€613,024	€64,993
	SM	€258,157	€1,258,232	€142,699
	Ratio	**2.49**	**2.05**	**2.20**

5.4. Sensitivity analysis

As we are not certain of the values of some parameters, we investigate the impact of these parameters. The parameters that we speak of in this case are the holding cost rate and tardiness costs. Their values may vary situationally, or we simply do not know the exact values. We assume that the holding cost rate probably lies somewhere between 20-30% of an item's value per year, as we note in Section 4.1.4. The tardiness costs are much more difficult to get an estimation of and are very situation-dependent. For our analysis we used both no tardiness costs and €2K tardiness costs per week, for which the area in between most likely reflects reality (based on expert opinion). We believe that the costs never exceed €10K per week and are never negative (profit for tardiness), so that is the range in which we test.

The sensitivity analysis in this section basically answers the question: If the uncertain parameters are in reality different than what we believe it to be, what does that cost us? So, if we solve our model with assumed parameters of holding cost rate and tardiness costs and they turn out to be different, how much higher expected total costs will we have compared to if we had solved the model with the correct parameters? To measure that, we use the following procedure:

1. For each tested combination of holding cost rate and tardiness costs (consecutively 20/25/30% and €0-10K with steps of €500), we solve the model. We therewith obtain optimal safety lead times for all tested combinations (63 combinations total).
2. We use 4 different base settings to compare the obtained solutions with. In other words, we say that we have 4 cases at which the *correct* or *true* parameters are known, and check the increase in expected total costs if we would use the *wrong* values of safety lead times with their different parameter settings. We can then observe how much the expected total costs deviate from the expected total cost of the optimal solution for each of the 4 base settings. As base settings, we use the combinations as shown in Table 5.3, which has both extremes as more likely tardiness costs.
3. We follow this procedure for the three Aerospace products that we now consider. Performing the same analysis for *Product 4* is not an option, as solving that model this often takes multiple weeks of running. The first three products should offer sufficient insight.

Figure 5.6 shows the results of the sensitivity analysis for the holding cost rate. It is obvious that assuming an incorrect holding cost rate on its own has a negligible effect on total costs, since we at most deviate about 1% from the optimal solution. Figure 5.7 presents the same results for tardiness costs. We immediately see that this is a much more important uncertain parameter. Especially if there are no tardiness costs while we do assume there are (high) costs, our solution will be far from optimal. Finally, Figure 5.8 depicts these results for the interaction effect with both holding cost rate and tardiness costs. We then observe that a wrong holding cost rate begins to have impact on the deviation from the optimal solution as well.

There are three points that we want to discuss given the figures: the insensitivity of *Product 1*, the seemingly illogical interaction effects, and what we may conclude from this analysis. The reason for the insensitivity of *Product 1* is simple. 4 out of 12 of its materials have a slack close to 0, which also happen to be the most expensive materials. In general, even more expensive materials will have an increasing safety lead time as tardiness costs rise. However, the low slack restrains these materials from having a higher safety lead time, resulting in practically the same solution and expected total costs. This is interesting, as it suggests that the safety lead times are not optimal because of the slack constraint. We further discuss this in Section 5.5.

If we look at *Case 1*, *Product 3* in Figure 5.8, we see the highest deviation from the optimal solution at the true value of the holding cost rate (20% in this case). Also in other cases and products, we observe that the true holding cost rate does not give the relatively best solution when tardiness costs are wrong. To understand why, we need to know the impact of both parameters on safety lead times. Lower holding costs generally result in higher safety lead times, except for the case of no tardiness costs where they are exactly the same for all holding cost rates (higher costs but same proportions). Higher tardiness costs naturally also result in higher safety lead times. When we assume too high tardiness costs in *Case 1*, the low holding cost rate results in extra high safety lead times, even further off from the optimal solution. The opposite also holds in other cases. So even though the model seems insensitive to the holding cost rate uncertainty from Figure 5.6, it is important to keep the interaction effect as seen in Figure 5.8 in mind.

The main conclusion that we can draw from this analysis is that it is very important to at least get some estimate of the tardiness costs, as it can have a large impact on the resulting solution and expected total costs. An inquiry to these costs may definitely be worthwhile. While these costs remain uncertain, it may be wise to include some tardiness costs such as €1K per week to avoid risk. When we optimise the model with such values, we will never deviate too far from the optimal solution. Furthermore, concerning the holding cost rate, using 25% per year is a fine assumption in all cases. Finally, the results from Section 5.3 will differ if we there wrongly assume tardiness costs, but they are then still an improvement as the deviation from the optimal solution is smaller than the measured improvement.

5.5. Impact of manually added slack

In Figure 4.1 we showed that the safety lead time is constrained by the slack of a material. Not being able to order earlier is a fact, since we then do not yet have the budget, know exact quantities, etc. Expediting the promised date to the customer is undesirable, but it is an option. Assuming that the customer date is determined such that the longest procurement lead materials will arrive just in time, then manually expediting this date by a week will give us a week extra to purchase earlier – allowing to increase safety lead times by a week. We call this manually adding slack or increasing customer lead time, as slack currently restrains materials from a certain safety lead time. Adding this slack could be viable, for example if we face high penalties for not meeting our promises or if inefficiency costs due to tardiness are high. This section analyses the impact of manually adding slack on the total costs.

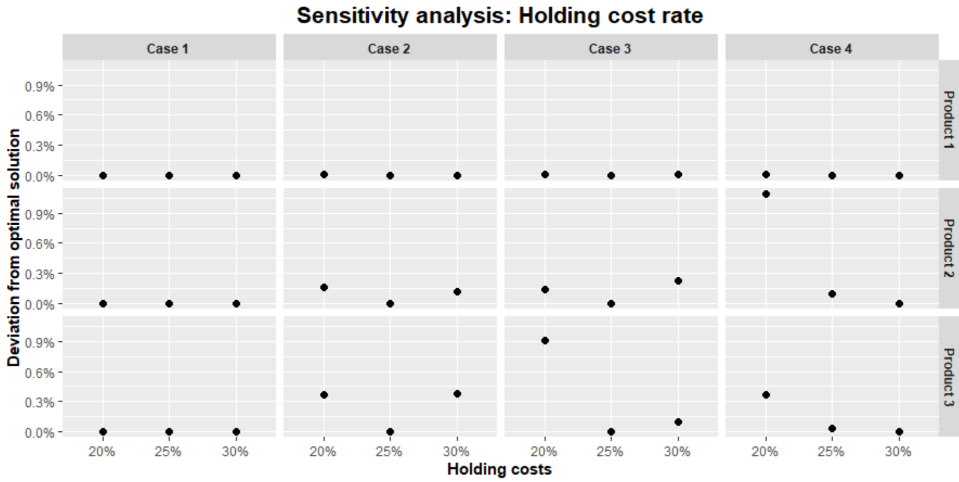

Figure 5.6. Sensitivity analysis of holding cost rate. Model is insensitive to this parameter alone.

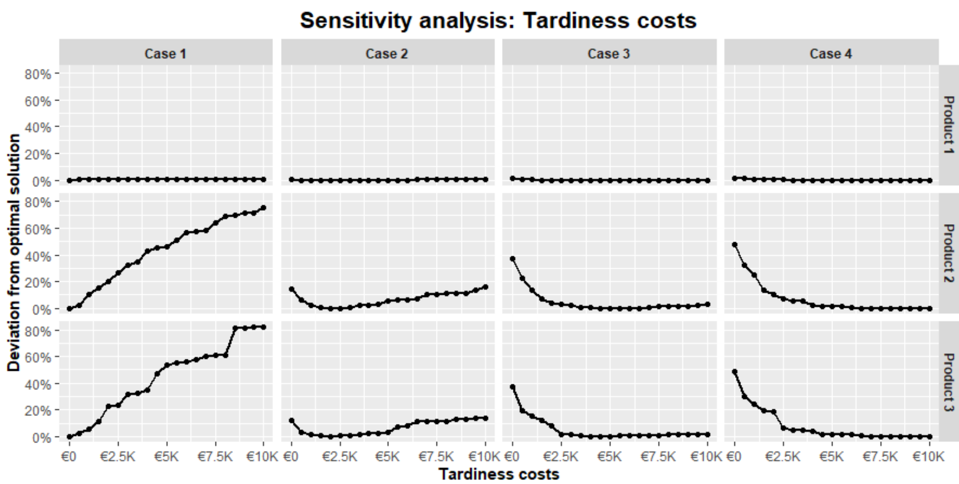

Figure 5.7. Sensitivity analysis of tardiness costs. Model is very sensitive to this parameter alone, especially when there are no true costs but we estimate (high) costs.

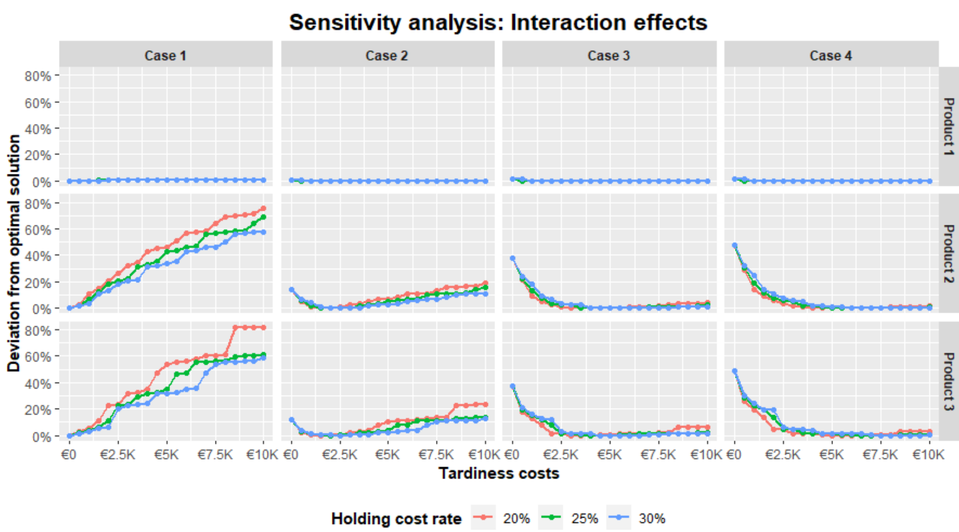

Figure 5.8. Sensitivity analysis of both parameters combined. We see that the holding cost rate now does have impact on the deviation from the optimal solution, when tardiness costs are wrong.

For this analysis, we again use the same four base settings of Table 5.3, such that we measure for different holding cost rates and tardiness costs. We currently use 0 weeks manually added slack, and measure to what extent we see a reduction of total costs by using 1 to 10 weeks manually added slack. Figure 5.9 shows these reductions for the first three products (again, running this often for *Product 4* simply takes too long) and the four cases. We observe that adding slack may definitely outweigh the longer customer lead time when there are considerable tardiness costs. The cost reduction per week added is by far largest at just one added week, and this also has minor impact on a customer lead time that is often already 1 to 2 years. Naturally, this constraint is mainly for the first batch, as later batches' safety lead times are not constrained by slack. We may increase safety lead times of later batches and further reduce total costs without adding this slack. Still, the first batch is most important, because lateness of that delivery delays all subsequent batches while other materials will still arrive at priorly planned moments. We also want only one safety lead time to fill in the system. Additionally, our advice should be adoptable without requiring many manual steps. To conclude, we believe that manually increasing the customer lead time to expand the possibilities of safety lead times can definitely decrease total costs and should be considered if tardiness costs are high.

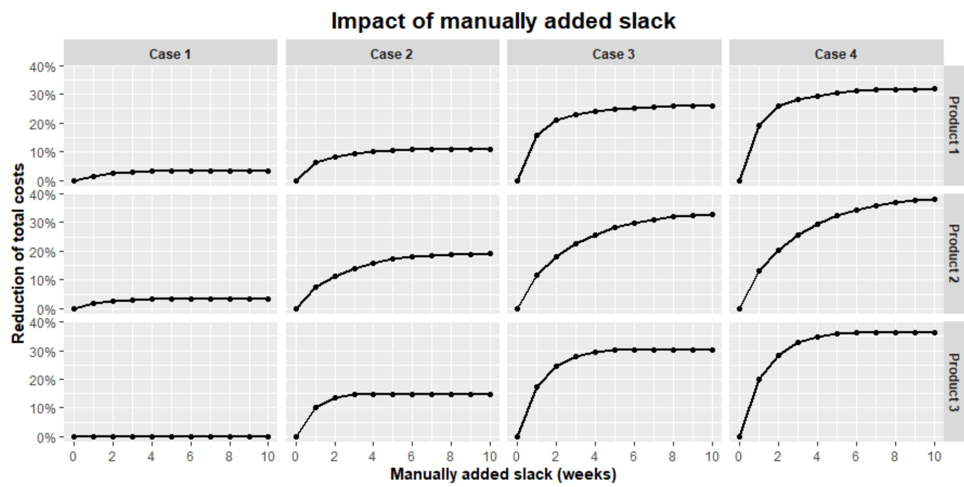

Figure 5.9. Impact of manually adding slack. When tardiness costs are high, adding slack can significantly reduce total costs.

5.6. Conclusion

In this section, we evaluate the performance of the core of our research – the safety lead time optimisation model. First, we evaluate the heuristic. We see that it is fast enough to easily evaluate multiple settings for Aerospace products, but it is too slow to evaluate a complete radar system. For those systems, a further decomposition of the most important subassemblies to consider is required. We also see that the impact of not having sufficient supplier data is minor, as material costs seem to play a much larger role in optimal safety lead times than the exact delivery lateness probability function. We conclude that it is a good option to use the dataset of all suppliers as input for that function if sufficient supplier data are lacking.

Both with the mathematical model and the Monte Carlo simulation model, we observe large possible reductions of total costs and product tardiness. We also consider to what extent the mathematical model represents the more realistic setting determined by the simulation model. The mathematical model overestimates relative improvements compared to the simulation model, but it also heavily underestimates total costs. These two facts combined make us believe that the expected improvement from the mathematical model are a good approximation or even lower bound. Given the products and batches presented in Table 5.1, we see improvements as shown in Table 5.5, promising hundreds of thousands euros total costs reduction. Extending this to even more products of Aerospace (and much more of Thales general), we could even expect savings of more than a million euros per year.

A sensitivity analysis shows that it is crucial to have a (somewhat close) estimate of the tardiness costs, while a wrong assumption of the holding cost rate has a much smaller impact on the optimal solution. While assuming e.g. €1K tardiness costs when using the model in the meantime most likely gives acceptable results, it is probably worthwhile to further analyse these costs. Finally, we investigate to what extent we could profit from manually extending customer lead time, such that we have more slack to increase safety lead times on materials. We observe that adding just a week can already make a large difference, especially when we see high tardiness costs.

6. Conclusion and recommendations

In Chapter 1, we introduced research opportunities and the research objective. It was clear that more explorative research was required to find the most suitable solution areas. This explorative research started in Chapter 2, where we analysed the current situation. Our analysis included mapping relevant sales and procurement processes, measuring forecast performance, analysing specific material characteristics, and evaluating current inventory policies. We then chose to focus on four solutions: safety lead time optimisation, more general process changes, a *Total Cost of Ownership* (TCO) model, and material classification. In Chapter 3, we gave background on inventory management, and reviewed literature on supply uncertainty in assembly settings and classification. Subsequently, we described the design of named solutions in Chapter 4. Finally, in Chapter 5 we analysed how the core of this research – the safety lead time optimisation model – performs for several products of Thales.

This chapter serves to share insights and discuss the research. Section 6.1 summarises and concludes on the most important findings. Subsequently, we give recommendations in Section 6.2, both to consider now and directions for future research. Finally, the discussion in Section 6.3 elaborates on the limitations of our research.

6.1. Conclusion

The main request from Thales was to aid in improving the on-time delivery by increasing material availability while reducing costs. We first entered explorative research before defining research questions, as there were numerous options to consider which required thorough analysis to assess their feasibility. Uncertainty is generally the hardest thing to deal with in a supply chain. Therefore, we did an analysis on possible improvements regarding both demand and supply uncertainty. On forehand, we already saw that there are many risks when making decisions with uncertain demand in the contract assembly setting. Working with the current purchase-to-order policy eliminates those risks. Still, large benefits are possible with regards to costs and customer lead time if we could anticipate on uncertain demand. For that reason we evaluated forecast performance and examined various risks that play a role. We concluded that Aerospace is not (yet) ready to base material decisions with high impact (e.g. purchases of expensive materials) on these forecasts.

At the same time, we observed that safety lead times are roughly the same for most materials. As each material's delivery lateness behaves differently and its holding costs also differ, this can never be optimal. We decided that optimising these safety lead time would be the core problem to tackle in this research. We recognised here as well that the data is not as we could hope for. Only for the delay caused by suppliers we see reliable data, and we lack sufficient data for most suppliers. Fortunately we were able to deal with both, as we explain later. Furthermore, we saw the need to quickly identify materials on their importance and risks, for which we chose to design a classification approach. Lastly, we decided to build a TCO model to ease decisions on purchase quantities given quantity discounts and holding costs (with known demand), and to propose process changes.

We learned multiple things from literature which we also applied in this research. Concerning supply uncertainty, we saw that safety lead times are mostly the preferred option over safety stocks. However, literature closely related to our setting is also scarce, especially when considering multi-period models. Nonetheless, it still offered a good model that approximates our problem, which formed the basis of our own model. With regards to material classification, we discussed several options to approach our classification design. Although many options were considered, amongst which some more complex ones, we decided to only include ABC classification in our own design.

Once we had all the information, we successfully built both models that resemble our situation. The models can be solved in Excel, which allows easy implementation at Thales. Unfortunately, the safety lead time optimisation model cannot capture all delay between the material need and when it actually arrives. Only the delivery lateness (time between supplier's promised date and actual arrival) is incorporated, as that is where we have more or less reliable data for. We were able to approximate these data with a gamma distribution, which is essential input for the model. In order to deal with the delay that the model does not capture (e.g. late purchase orders from Thales), we suggested changes to the communication process of sales and procurement. For the classification approach, we decided that important parameters to classify on are slack to the critical path (see Section 2.4.1 for a short definition), material costs, revision risk and forecasted demand (whether we expect *any* demand in the next year). We concluded that it offers a nice categorisation of materials with related policies, but that it also should be developed further.

Subsequently, we evaluated the performance of the safety lead time optimisation model. One specific useful insight is that it is not a big deal that we lack sufficient delivery lateness data for most suppliers. The material's cost price is a much more important factor on the optimal safety lead time than its exact delivery lateness distribution function. If we lack sufficient data for a supplier, we can assume that the supplier performs average (compared to other suppliers) and still get good results. The most important conclusion for this model is that it definitely promises a great reduction of holding costs and tardiness costs (costs of production delay). Both when evaluating the model's results as a more realistic simulation model, we see cost reductions of at least a few hundred thousand euros, with possibly more than a million per year when extending it to more products. We also saw that the model is very sensitive to a wrong estimation of tardiness costs, so it is worthwhile to investigate these costs.

Aerospace has been the primary focus of this research, but generalising these solutions to the rest of Thales is important as well. We conclude that all of our proposed solutions can very well be used for radar systems. For the safety lead time optimisation model, we especially require a different approach towards materials that are used in multiple final assemblies (commonality). These materials hardly exist in Aerospace, so we did not analyse that issue. Also, we saw that the model takes long to find the optimal solution for large radar systems. It may be necessary to decouple such a large final assembly into smaller subassemblies to analyse individually. Finally, more parameters – such as commonality – may be interesting for material classification outside Aerospace.

6.2 Recommendations

First things first, we obviously recommend to implement the proposed solutions from Chapter 4 (models, process changes and material classification). As a pilot, the exact safety lead times from the safety lead time optimisation model could be used, while further improvement may follow from experience in practice. Next to the solutions, we also gained more general insights which we shall share. Section 6.2.1 gives these managerial implications. In Section 6.2.2, we describe directions for future research, which we also place in an impact effort matrix.

6.2.1. Managerial implications

Regarding this research, we have various managerial implications for Thales:

1. *Consider the bigger picture when making material choices*. For the first managerial implication we refer to the title of this thesis. It is essential to consider the final assembly when making material choices. Do not attempt to optimise for separate materials, but consider the impact on the production of their final assembly. For example, inspecting a material that has no defects 90% of the time may not seem worthwhile for the material itself, but if it results in fewer delays for the final assembly, it could be. Also, a material with a procurement lead time of 20 weeks can be critical if its final assembly only has other materials with smaller lead times. At the same time, a material with 40 weeks procurement lead time may be much less relevant if other materials of that assembly exceed

60 weeks lead time. Only by bearing the final assembly in mind, we can properly assess each material and make the right decisions.

2. *Record delivery lateness data very carefully and use it in more decisions*. It seems that these data are now mainly used to assess whether a supplier is performing well and if actions are required. However, as supply uncertainty is the major type of uncertainty in a contract assembly setting, it is an absolute must to include it in various types of decisions such as safety stocks. These data should also not be polluted, as now can be the case due to e.g. not wanting to make a supplier look bad. Data to be used in quantitative decisions must absolutely be kept separate from such actions.

3. *Consider the type of uncertainty for which a safety stock is held*. If a safety stock is built to deal with procurement lead time variance, we often only need enough stock to work for a few weeks. In case it is built for expected defects, we should have just enough (or a bit more) for those expected defects. Uncertain demand is rarely the case, especially for Aerospace: Most materials can supposedly be ordered in time. Together with the safety lead times, when keeping this in mind we should be able to both increase material availability and reduce stocks.

6.2.2. Directions for future research

We identify the following directions for future research:

1. *Extend the safety lead time optimisation model*. The model now offers a good basis, but there are many extensions possible to increase performance. First of all, to use the model outside Aerospace, we should consider how to deal with commonality of materials: What is the optimal safety lead time for materials used in multiple assemblies? Other interesting factors to include are to make it a multi-level assembly model and to include multiple periods (optimise over multiple batches instead of one). It all helps in creating a more realistic model, which should allow to further reduce costs. Time spent on extending the model may vary from a little to a lot, and impact coincides.

2. *Optimise Fixed Days Supply (FDS) settings for materials*. Thales currently uses a FDS of 30 workdays for all materials. Recall from Section 2.2.2 that, because of that, demand of 30 workdays is aggregated into one purchase order. For very expensive materials, that means we have hundreds of thousands on stock for several weeks before it is used – *as a standard setting*. For these materials, it is most likely better to have a lower FDS resulting in more purchase orders with a lower average stock level in terms of holding costs and fixed costs of ordering. Also vice versa, increasing FDS and therewith decreasing purchase orders may decrease work required for buyers on cheap materials. Although finding the best FDS settings may be quite difficult, some changes on basic rules (such as material costs) may already have much impact.

3. *Quantify tardiness costs*. We saw in our performance evaluation that to find the optimal safety lead times, it is crucial to have at least some estimate of the tardiness costs. We recognise that finding these costs can be very difficult, as it includes both tangible costs (e.g. penalties) and non-tangible costs (e.g. inefficiency costs). A study to get a grasp on these costs, perhaps per product or customer, may definitely be worthwhile to get better safety lead times and therefore lower costs.

4. *Consider quality inspections from a quantitative approach*. Most materials are currently not inspected upon arrival at Thales. If we only find material defects at production, we lose valuable time. Early detection means we can already reorder and start production earlier. By considering probability functions of material defects, inspection costs, delay due to inspection, tardiness costs and time won, we could again make the trade-off between inspection or not. In the bigger picture, we could attempt to find if it is feasible to hire employees for the job. Considering these quality

inspections, and perhaps designing a framework to decide which materials to inspect, could be the subject of a future master's thesis.

5. *Further analyse delivery lateness distributions*. We now assume that any material of the same supplier has the same delivery lateness distribution function. In practice, some materials may have much more delivery issues than others. Having sufficient data for all separate materials would be ideal, but due to our high-mix low-volume material demand that is not realistic. A possible option for example would be to have a scaling factor on the distribution function. Having a more realistic distribution function results in better safety lead times and less required manual adjustments from the model's user. We expect that this should not take too much time effort, but also has moderate impact as concluded in Section 5.2.

6. *Further consider factors and policies for material classification.* The classification method we now designed is only the start. It would now be interesting to see if materials that we experience issues with are actually in a 'risk policy'. If they fall outside of that, we should check what other parameters we might be missing – what else defines criticality and risk? In addition, we now suggest a few basic policies. These may be extended or even quantified. We do not expect much impact on costs, OTD or such KPIs, but it may decrease required effort on the long run as it eases decision making.

7. *Include other influences of material delay in the model*. We now only include the delay caused by suppliers, but we also saw that it is only a fraction of total delay. If we register all types of delay correctly, we have more accurate input for our model and can further increase the reliability of its resulting safety lead times. Setting up registration correctly may take some effort, and it is difficult to assess whether we can actually model all other delays as they come from many different sources. However, if we succeed, we may expect even better safety lead times.

Figure 6.1 shows these directions for future research in an impact effort matrix, allowing a quick overview.

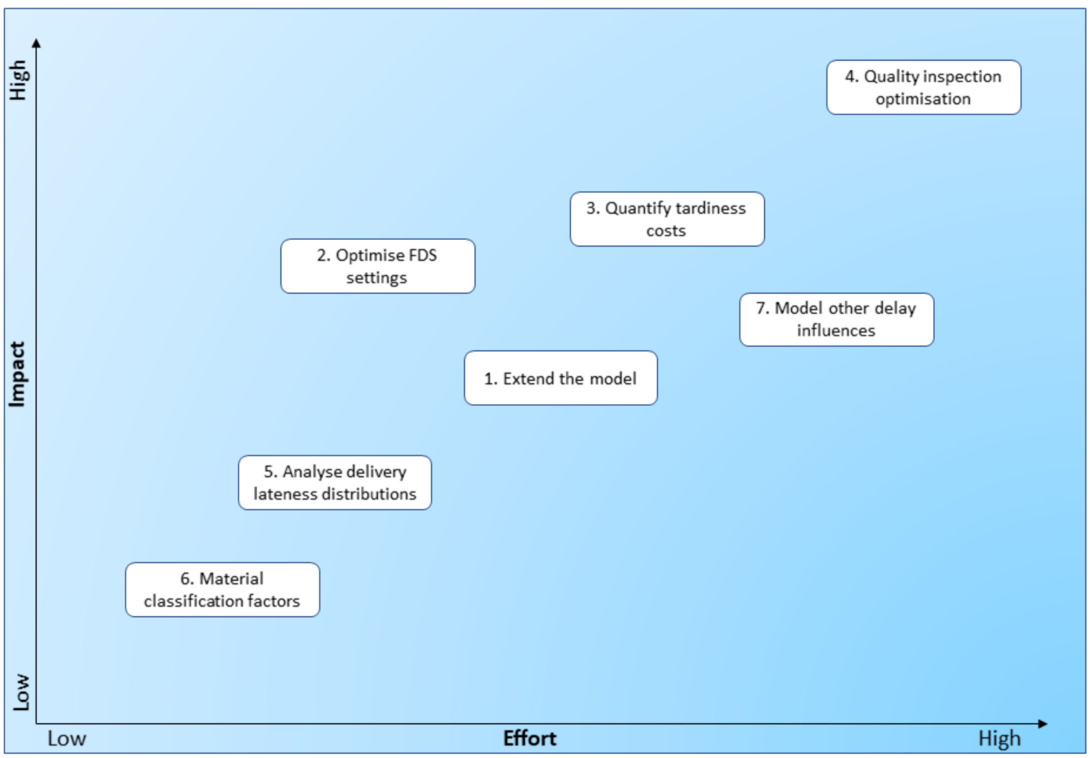

Figure 6.1. Impact effort matrix of directions for future research

6.3. Discussion

Despite the good results that we observed, there are limitations to our research. Some were already (indirectly) mentioned in the conclusion and recommendations, but we name them all to be complete.

The foremost limitations are those of the safety lead time optimisation model. We first consider the data of material delays. Our model considers delivery lateness, which we defined as the delay caused by suppliers. This was the best data that we could use, but we observed that there are other causes of delay (e.g. that we order too late such that the materials cannot arrive on time). Additionally, we assume that any material of a supplier is as likely to be late as the supplier's other materials. In practice, some materials may be (much) more difficult to deliver compared to others. Despite that we saw that material costs play a bigger role on optimal safety lead times than delivery lateness, this may still significantly change the model's solutions. On the other hand, the model gives an advice, and we expect that the model's user is aware of these materials.

Additionally, the model simplifies reality in different ways. The simplification with highest impact is that we consider a single-period model. This does not reflect reality, as we observe multiple batches where delay in one batch influences following batches. Although the Monte Carlo simulation verifies that also in the multi-period setting our safety lead times are still a large improvement compared to the current situation, the model most likely does not give the optimal safety lead times. Another simplification is that we treat all assemblies as a two-level assembly. This also underestimates the required safety lead times for (some) materials, though we believe that the impact is much smaller.

A further limitation of named model is that we assume that min/max materials are always on stock. This is often true in practice, but not always – making the results less valid. In addition, we disregard material quality issues in the model. Again, materials may have defects and this may cause delays, but at the same time we see extra material purchases or safety stocks to cope with that currently. Optimising quality inspections is definitely interesting, which is why we recommend it as a direction of future research.

Besides the optimal safety lead time model, the TCO model also has its limitations. We only consider two known moments of demand. In practice, we will have many moments of demand, and we could optimise with numerous moments to order and as many different order quantities. As a matter of fact, we built a model earlier that did exactly that analysis for uncertain demand. However, the goal of the TCO model is to offer an easy tool that determines the best order quantity in a simple setting, with known demand. Also, consistent demand over a longer period (often the case with the batches of Aerospace) can be approximated with the TCO model as well. And as long as it always reminds people to also incorporate holding costs when considering quantity discounts, it is already an improvement.

We also discussed that the material classification approach should still be developed further. Especially regarding the policies, it is quite limited. We currently suggest policies such as 'stock abundantly' for critical materials with low stock. If we want to significantly ease decision making, these policies should also provide a quantitative advice, such as 'keep safety stock of 4 weeks demand'. However, at this point we cannot give these quantitative advices, as we cannot assess the impact – especially combined with the new safety lead times. Therefore, it could be wise to experiment with such safety stocks for critical materials, and therewith find more exact policies for each class.

References

Ben-Ammar, O., Dolgui, A., & Wu, D. (2018). Planned lead times optimization for multi-level assembly systems under uncertainties. *Omega (U.K.) 78*, 39-56.

Ben-Ammar, O., Marian, H., Wu, D., & Dolgui, A. (2013). Mathematical model for supply planning of multi-level assembly systems with stochastic lead times. *IFAC Proceedings Volumes 46 (9)*, 389-394.

Bhattacharya, A., Sarkar, B., & Mukherjee, S. (2007). Distance-based consensus method for ABC analysis. *International Journal of Production Research 45 (15)*, 3405-3420.

Borodin, V., Dolgui, A., Hnaien, F., & Labadie, N. (2016). Component replenishment planning for a single-level assembly system under random lead times: A chance constrained programming approach. *International Journal of Production Economics 181 (A)*, 79-86.

Botter, R., & Fortuin, L. (2000). Stocking strategy for service parts - a case study. *International Journal of Operations & Production Management 20 (5-6)*, 656-674.

Braglia, M., Grassi, A., & Montanari, R. (2004). Multi-attribute classification method for spare parts inventory management. *Journal of Quality in Maintenance Engineering 10*, 55-65.

Cakir, O., & Canbolat, M. (2008). A web-based decision support system for multi-criteria inventory classification using fuzzy AHP methodology. *Expert Systems with Applications 35 (3)*, 1367-1378.

Chen, J.-X. (2012). Multiple criteria ABC inventory classification using two virtual items. *International Journal of Production Research 50 (6)*, 1702-1713.

Chen, Y., Li, K., Kilgour, D., & Hipel, K. (2008). A case-based distance model for multiple criteria ABC analysis. *Computers and Operations Research 35 (3)*, 776-796.

Chopra, S., & Meindl, P. (2013). *Supply Chain Management (5th ed.).* Essex: Pearson Education Limited.

Dickie, H. (1951). ABC Inventory Analysis Shoots for Dollars. *Factory Management and Maintenance 109 (7)*, 92-94.

Dolgui, A., & Prodhon, C. (2007). Supply planning under uncertainties in MRP environments: A state of the art. *Annual Reviews in Control 31*, 269-279.

Dolgui, A., Ben-Ammar, O., Hnaien, F., & Ould-Louly, M.-A. (2013). A State of the Art on Supply Planning and Inventory Control under Lead Time Uncertainty. *Studies in Informatics and Control 22 (2)*, 255-268.

Fallah-Jamshidi, S., Karimi, N., & Zandieh, M. (2011). A Hybrid Multi-objective GA for Planning Order Release Date in Two-level Assembly System with Random Lead Times. *Expert Systems with Applications 38 (11)*, 13549-13554.

Flores, B., & Whybark, D. (1986). Multiple Criteria ABC Analysis. *International Journal of Operations & Production Management 6 (3)*, 38-46.

Flores, B., Olson, D., & Dorai, V. (1992). Management of multicriteria inventory classification. *Mathematical and Computer Modelling 16 (12)*, 71-82.

Grasso, E., & Taylor, B. (1984). Simulation-based experimental investigation of supply/timing uncertainty in MRP systems. *International Journal of Production Research 22 (3)*, 485-497.

Hadi-Vencheh, A. (2010). An improvement to multiple criteria ABC inventory classification. *European Journal of Operational Research 201 (3)*, 962-965.

Hadi-Vencheh, A., & Mohamadghasemi, A. (2011). A fuzzy AHP-DEA approach for multiple criteria ABC inventory classification. *Expert Systems with Application 38 (4)*, 3346-3352.

Harris, F. (1913). How Many Parts to Make at Once. *Magazine of Management 10*, 135-136.

Hopp, W., & Spearman, M. (1993). Setting Safety Leadtimes for Purchased Components in Assembly Systems. *IIE Transactions 25 (2)*, 2-11.

Li, Z., Wu, X., Liu, F., Fu, Y., & Chen, K. (2019). Multicriteria ABC inventory classification using acceptability analysis. *International Transactions in Operational Research 26 (6)*, 2494-2507.

Lolli, F., Ishizaka, A., & Gamberini, R. (2014). New AHP-based approaches for multi-criteria inventory classification. *International Journal of Production Economics 156*, 62-74.

Mula, J., Poler, R., Garcia-Sabatar, J., & Lario, F. (2006). Models for Production Planning under Uncertainty: A Review. *International Journal of Production Economics 103 (1)*, 271-285.

Nahmias, S. (1979). Simple Approximations for a Variety of Dynamic Leadtime Lost-Sales Inventory Models. *Operations Research 27 (5)*, 904-924.

Nahmias, S., & Olsen, T. (2015). *Production and Operations Analysis (7th ed.).* Long Grove, Illinois: Waveland Press, Inc.

Ng, W. (2007). A simple classifier for multiple criteria ABC analysis. *European Journal of Operational Research 177*, 344-353.

Ould-Louly, M.-A. D. (2008). Supply Planning for Single-level Assembly System with Stochastic Component Delivery Times and Service Level Constraint. *International Journal of Production Economics 115 (1)*, 236-247.

Ould-Louly, M.-A. D. (2009). Calculating Safety Stocks for Assembly Systems with Random Component Procurement Lead Times: A Branch and Bound algorithm. *European Journal of Operational Research 199 (3)*, 723-731.

Ould-Louly, M.-A. D. (2011). Optimal Time Phasing and Periodicity for MRP with POQ Policy. *International Journal of Production Economics 131 (1)*, 76-86.

Ould-Louly, M.-A., & Dolgui, A. (2002). Generalized newboy model to compute the optimal planned lead times in assembly systems. *International Journal of Production Research 40 (17)*, 4401-4414.

Ould-Louly, M.-A., & Dolgui, A. (2013). Optimal MRP Parameters for a Single Item Inventory with Random Replenishment Lead Times, POQ Policy and Service Level Constraint. *International Journal of Production Economics 143*, 35-40.

Pareto, V. (1971). *Manual of Political Economy.* New York: A.M. Kelly Publishers.

Partovi, F., & Burton, J. (1993). Using the Analytic Hierachy Process for ABC Analysis. *International Journal of Operations & Production Management 13 (9)*, 29-44.

Plenert, G. (1999). Focusing material requirements planning (MRP) towards performance. *European Journal of Operational Research 119*, 91-99.

Proth, J., Mauroy, G., Wardi, Y., Chu, C., & Xie, X. (1997). Supply management for cost minimization in assembly systems with random component yield times. *Journal of Intelligent Manufacturing 8*, 385-403.

Ramanathan, R. (2006). ABC inventory classification with multiple-criteria using weighted linear optimization. *Computers & Operations Research 33*, 695-700.

Saaty, T. (1980). *The Analytic Hierarchy Process.* New York: McGraw-Hill.

Silver, E., Pyke, D., & Thomas, D. (2017). *Inventory and Production Management in Supply Chains (4th ed.).* Boca Raton, Florida: Taylor & Francis.

Song, J., Yano, C., & Lerssrisuriya, P. (2000). Contract Assembly: Dealing with Combined Supply Lead Time and Demand Quantity Uncertainty. *Manufacturing & Service Operations Management 2 (3)*, 287-296.

Van Kampen, T., Van Donk, D., & Van der Zee, D.-J. (2010). Safety Stock or Safety Lead Time: Coping with Unreliability in Demand and Supply. *International Journal of Production Research 48 (23-24)*, 7463-7481.

Weber, J. (2017, September 2). *People & Profits: Do You Know Your Carrying Costs?* Retrieved from Farm Equipment: https://www.farm-equipment.com/blogs/6-opinions-columns/post/14452-people-profits-do-you-know-your-carrying-costs

Whybark, D., & Williams, J. (595-606). Material requirements planning under uncertainty. *Decision Science 7*, 1976.

Yano, C. (1987). Stochastic Leadtimes in Two-level Assembly Systems. *IIE Transactions 19 (4)*, 95-106.

Yu, M. (2011). Multi-criteria ABC analysis using artificial-intelligence-based classification techniques. *Expert Systems with Applications 38 (4)*, 3416-3421.

Zhou, P., & Fan, L. (2007). A note on multi-criteria ABC inventory classification using weighted linear optimization. *European Journal of Operational Research 182 (3)*, 1488-1491.

Appendices

Appendix A. Procurement lead times agreements

Department	Transportation term	Pre-processing	Processing		Post-processing			
			Lead time supplier	Transport time	Inbound time	Inspection time	Safety time	Sub-total
-	Delivered at Thales	5	Product and supplier dependent	0	2	0	15	17
	Own transport (NL)		Product and supplier dependent	2				
	Own transport (EU)		Product and supplier dependent	3				
	Own transport (US)		Product and supplier dependent	10				
Other inspection codes	Delivered at Thales		Product and supplier dependent	0		10		27
	Own transport (NL)		Product and supplier dependent	2				
	Own transport (EU)		Product and supplier dependent	3				
	Own transport (US)		Product and supplier dependent	10				
Comment		Standard. Time for processing a planned order to purchase order.	Lead time provided by supplier	Transportation time from supplier to Thales	Time between truck to rack	Time needed for inbound inspection	Safety time for variance in processing time	Total post-processing time

Appendix B. Evaluation S&OP forecasts

B.1. Forecasted vs. actual amounts for 2016, 2017 and 2018

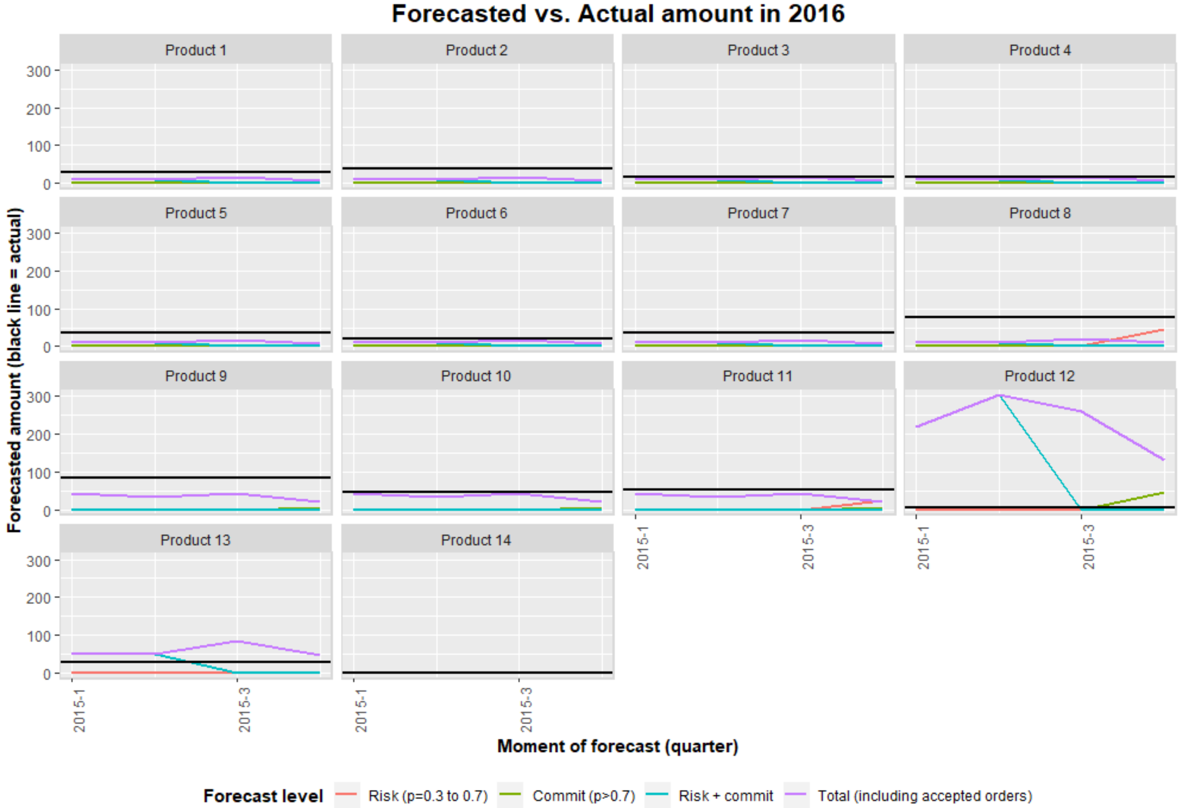

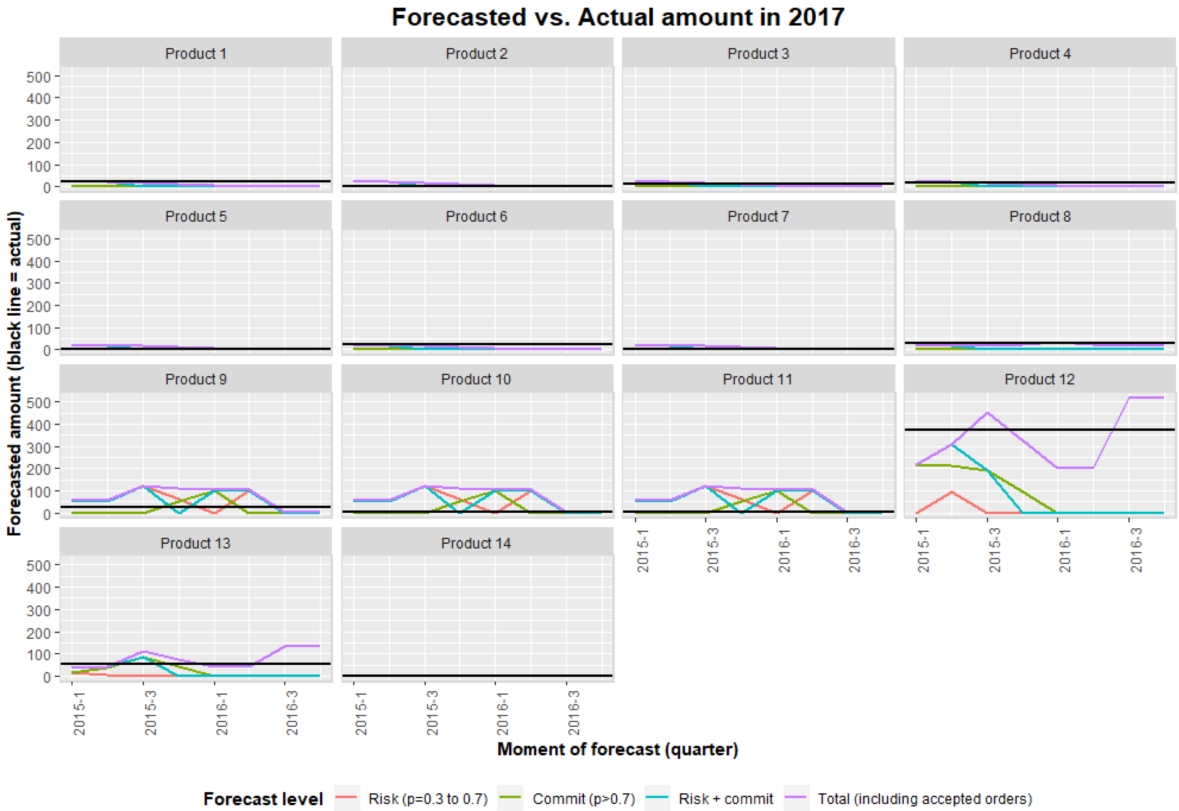
Forecasted vs. Actual amount in 2017
Product 1
Product 2
Product 3
Product 4
Product 5
Product 6
Product 7
Product 8
Product 9
Product 10
Product 11
Product 12
Product 13
Product 14
Forecasted amount (black line = actual)
Moment of forecast (quarter)
Forecast level
Risk (p=0.3 to 0.7)
Commit (p>0.7)
Risk + commit
Total (including accepted orders)

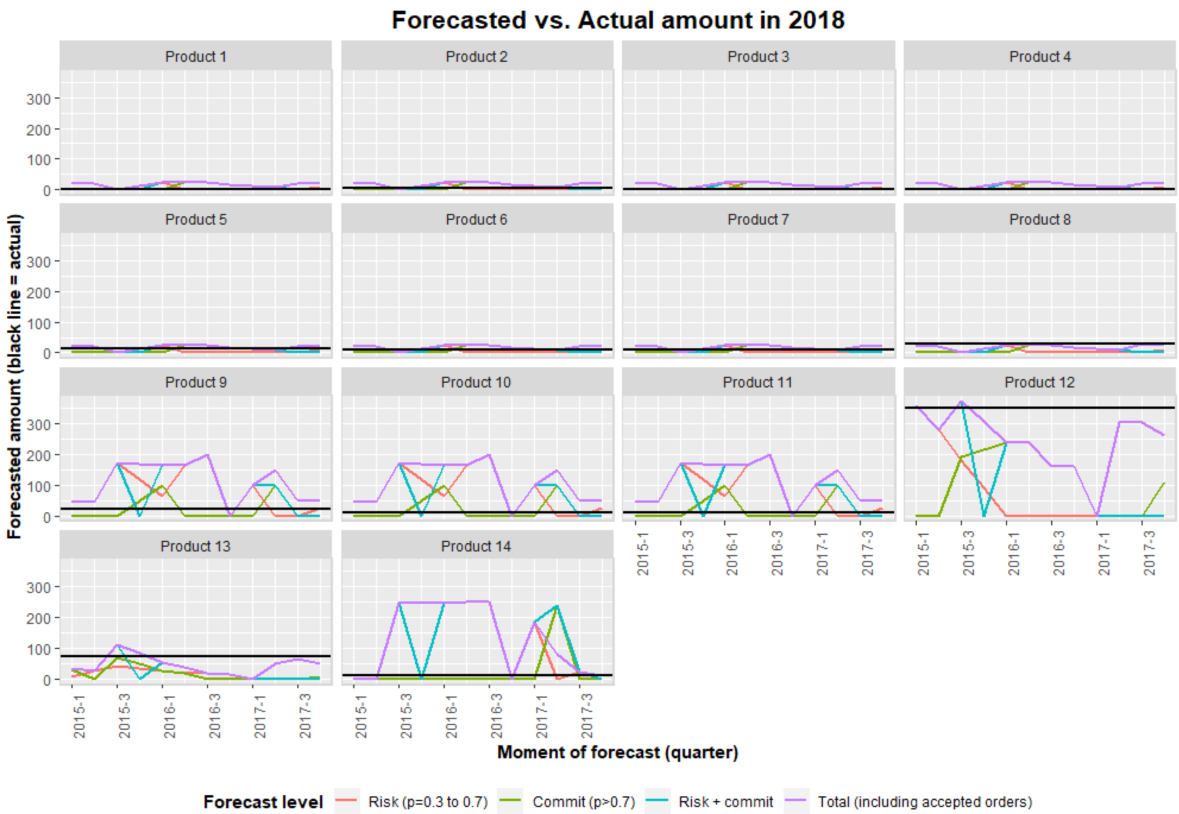
Forecasted vs. Actual amount in 2018
Product 1
Product 2
Product 3
Product 4
Product 5
Product 6
Product 7
Product 8
Product 9
Product 10
Product 11
Product 12
Product 13
Product 14
Forecasted amount (black line = actual)
Moment of forecast (quarter)
Forecast level
Risk (p=0.3 to 0.7)
Commit (p>0.7)
Risk + commit
Total (including accepted orders)

B.2. Relative forecasting bias S&OP forecast

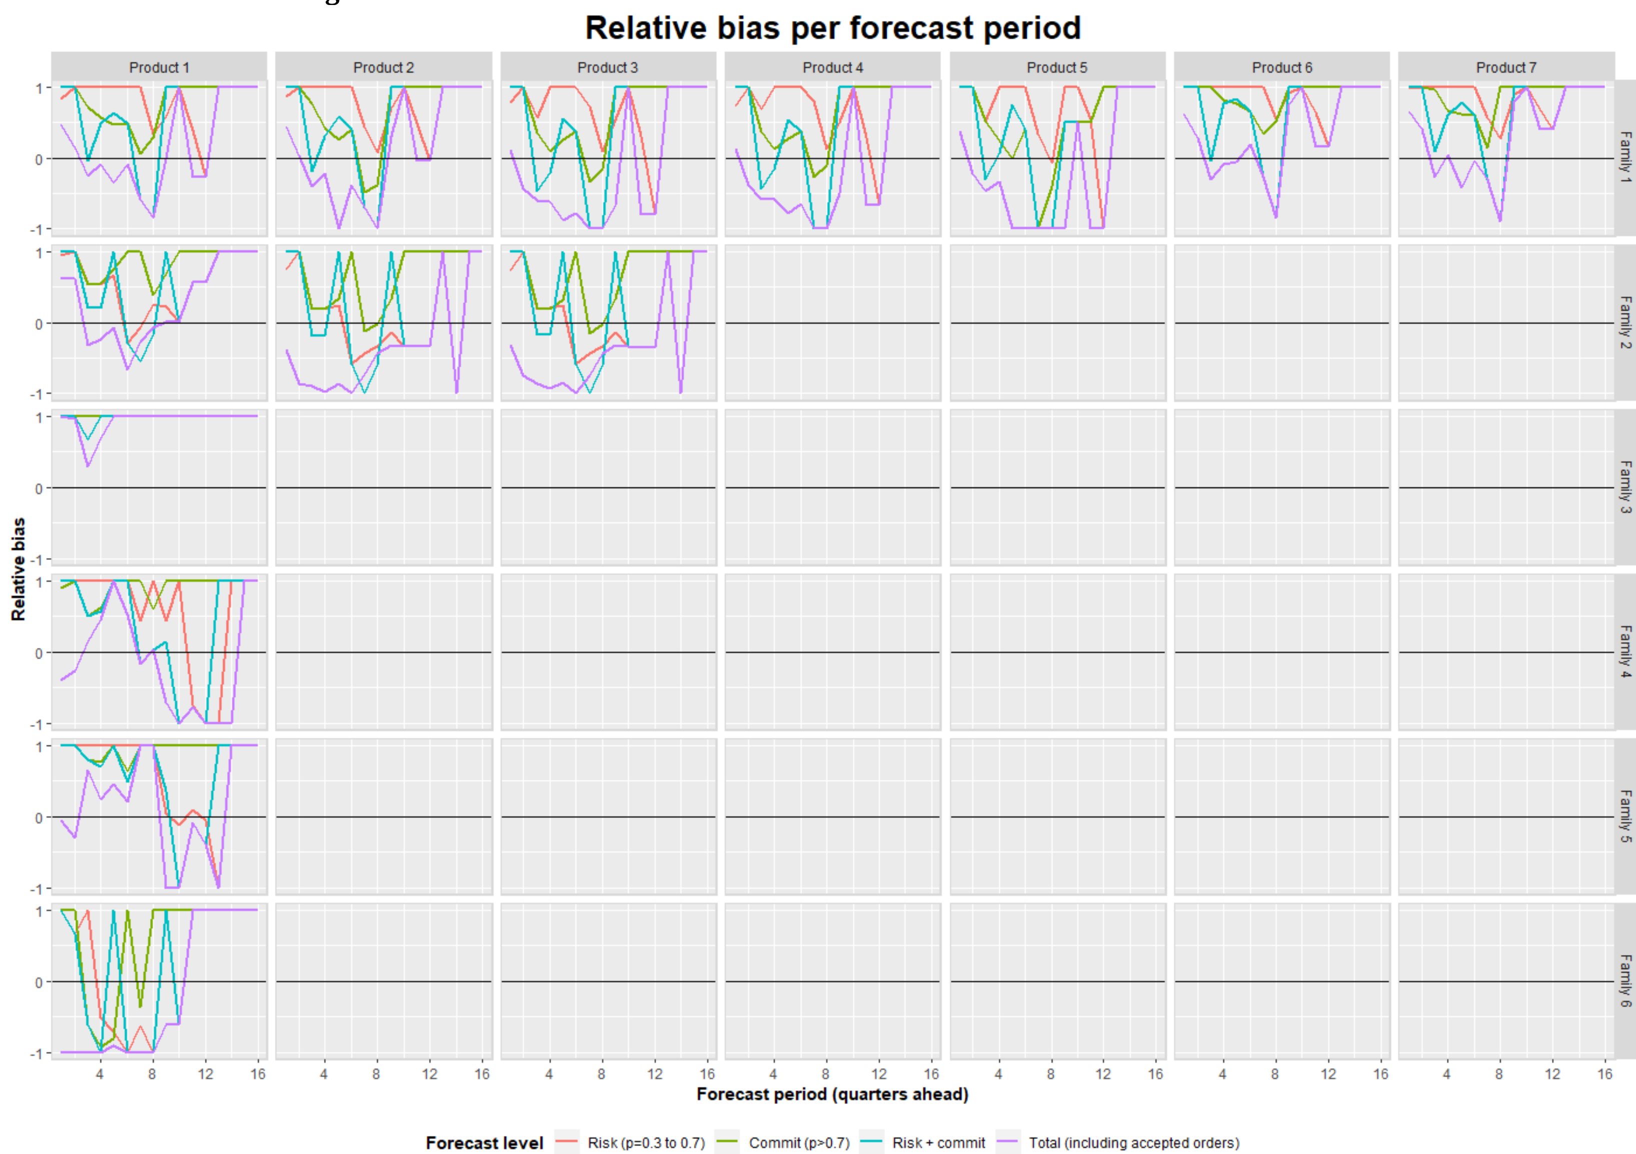

Appendix C. Mathematical notation of slack to critical path

Consider that a final assembly is made of *n* materials, both buy and make. Each material is used in an assembly, be it the final assembly or a subassembly (make material). The final assembly has index 0, and materials have index between 1 and *n*. Materials may be used in multiple other assemblies, but make materials are unique (not used in multiple other items). The index is given by either *i*, *j*, *k* or *l*. Recall from Section 1.3 that each material has a cumulative total lead time. We notate that material *i* used in item *j* has cumulative total lead time c_{ij}. The slack of material *i* used in item *j* is given by s_{ij}. Only for the final assembly, which is not used in any other item, we have cumulative total lead time c_0 and slack s_0. We can then determine the slack as

$$s_0 = 0$$
$$s_{ij} = s_{jk} + \max_{1 \le l \le n}\{c_{lj}\} - c_{ij}$$

Now, to translate it to materials in general (instead of materials used in another item), we of course have

$$s_i = \min_{1 \le j \le n}\{s_{ij}\}.$$

Appendix D. Technical description Monte Carlo simulation

In the Monte Carlo simulation, we basically describe relations with formulas and run the simulation many times with random numbers. In our case, the procurement lead times are random. In a logical order, the Monte Carlo simulation includes the following relations. We use Excel formula notation and notation as in Table 5.1 and sections 4.1.2, 4.1.3. We first give the notation for all parameters, and explain the simulation after.

Parameters

FDS : Fixed Days Supply (weeks)
b : Total number of batches
T : Batch cycle time (number of weeks between start of 2 batches)
D : Demand per batch (number of systems)
n : Number of batches aggregated into one purchase order
q_i : Purchase order quantities (number of systems) of purchase order *i*
π_k^{L-} : Probability of negative lateness of material *k*
e_k : Average earliness of material *k* (weeks)
m_k : Maximum delivery lateness of material *k* (weeks)
α_k : Alpha parameter of gamma distribution of material *k*
β_k : Beta parameter of gamma distribution of material *k*
c_k : Material costs per system for material *k* (€)
x_k : Safety lead time for material *k* (weeks)
d_{ik} : Delay of purchase order *i* of material *k* (weeks)
r_{ik} : Randomly drawn number for purchase order *i* of material *k*
A_{ik} : Arrival moment of purchase order *i* of material *k*
S_j^P : Planned start moment of batch *j*
S_j^A : Actual start moment of batch *j*
I_{kt} : Inventory level of material *k* at period *t*
C_{kt} : Holding costs made for material *k* at period *t*

Explanation

First, we have a number of material arrivals. We know that the FDS is 30 workdays (or 6 weeks). We also know the number of batches *b* to be produced, the batch cycle time *T* and the material demand per batch *D*. Therefore, we can determine for how many batches the material purchases are aggregated and how many purchase orders we have. The number of batches in one purchase order *(n)* is equal to 1 + ROUNDDOWN(*T/FDS*), so with 1 batch per 4 weeks and the FDS of 6 weeks, we get 1 + 1 = 2 batches of demand into one purchase order. Given *b* batches, we have *b*/*n* purchase orders. We use *i* for purchase orders (*PO*) 1 to *b/n*, *j* for batches 1 to *b*, and *k* for materials.

To determine the lead time of each purchase order *i* of material *k*, we draw a random number r_{ik} between 0 and 1: RAND(). The delay of the first PO (d_{1k}) is as follows (e_k is the average earliness of material *k* and maximum delivery lateness is m_k):

$$d_{1k} = \mathrm{IF}\left(r_{ik} \leq \pi_k^{L-}, e_k, \mathrm{MIN}\left(\mathrm{ROUND}\left(\mathrm{GAMMA.INV}\left(\frac{r_{ik} - \pi_k^{L-}}{1 - \pi_k^{L-}}, \alpha_k, \beta_k\right), 0\right), m_k\right)\right)$$

For subsequent POs this is basically the same, but we say that material orders may not cross in time. We need an additional variable to make that clear. We know that one purchase order aggregates *n* batches because of FDS. With that, we know that the cycle time of purchase orders is *nT* (e.g. with previous numbers, 2 batches per PO and 1 batch per 4 weeks, so one PO every 2 * 4 = 8 weeks). If a PO is delayed more than this nT, that means that the next PO is also delayed. Therefore, we get the following for i = 2, 3, …, b/n:

$$d_{ik} = \text{MAX}\left(d_{i-1,k} - nT, \text{IF}\left(r_{ik} \leq \pi_k^{L-}, e_k, \text{MIN}\left(\text{ROUND}\left(\text{GAMMA.INV}\left(\frac{r_{ik} - \pi_k^{L-}}{1 - \pi_k^{L-}}, \alpha_k, \beta_k\right), 0\right), m_k\right)\right)\right)$$

Knowing the delay, we can easily calculate the PO arrival moments (A_{ik}):

$$A_{1k} = d_{1k} - x_k$$
$$A_{ik} = i * nT + d_{ik} - x_k \; \forall i \geq 2$$

The PO quantities (q_i) naturally follow from the number of batches per PO and demand per batch:

$$q_i = nD \; \forall i$$

The first batch is planned to start at 0: $S_1^P = 0$. The subsequent batches are planned to start at:

$$S_j^P = (j - 1)T \; \forall j \geq 2$$

When the batches actually start, is at either the planned start date if there are sufficient materials, or later than that when all materials arrived. However, capacity also plays a role. There must always be the batch cycle time between two batches. So if one batch is delayed due to a lack of materials and starts two weeks later, then the subsequent batches also start two weeks later – regardless of the material availability. So, to find the actual start moment of batch j (S_j^A), we need to define the inventory level of material k at period t (I_{kt}):

$$I_{kt} = I_{k,t-1} + \text{SUMPRODUCT}\left(\text{IF}\left(A_{1k}\text{:}A_{\frac{b}{n},k} = t, 1\right), q_1\text{:}q_{\frac{b}{n}}\right) - \text{IF}\left(S_j^A = t, 1, 0\right) * D$$

And finally, the holding costs are calculated for every period that we have stock. So for material k at period t, the holding costs C_{kt} are:

$$C_{kt} = I_{kt} * 0.25 * \frac{1}{52} c_k$$

And the total costs are the sum of all holding costs for both materials and periods. The average delay is determined as the average difference between all S_j^P and S_j^A.