Strategic capacity planning with workforce flexibility to deal with seasonal and variable demand

A case study in the agriculture sector at Company A

Harmen Denekamp, 01-07-2020, University of Twente

Author
Harmen Denekamp
Master Industrial Engineering and Management
Production and Logistics Management

Supervisors
Faculty of Behavioural Management and Social Sciences
Dep. Industrial Engineering and Business Information Systems (IEBIS)
Dr. E. Topan
Dr. Ir. J.M.J. Schutten

Company A Group, part of Croda International
M. Mulder, Supply Chain Manager
R. Boots, Planning Manager

Date: 10-05-2020

Preface

This master thesis was my first glance in the agriculture industry. Company A has been a great window in the world of seeds. Their pioneering and commitment to the craft of seed enhancements is inspiring. Despite being part of a large multinational since 2015, the atmosphere remains closer to a family-owned business. I want to thank Michiel and Rowan for their time and expertise, you were always ready to answer questions and provide new perspectives, which helped me tremendously. And of course the entire supply chain department for the fun we have had together. Not only the various dinners and trips, but especially the conversations and jokes.

I also want to thank Engin and Marco for their guidance. Engin, I appreciate your kindness and the discussions we have had, it helped me to think more critically. Marco, I appreciate your honest and specific feedback, it especially helped me to improve the structure of my thesis.

I hope this thesis will be enjoyable to read, if it is your cup of tea, and contribute to the theory on capacity planning and the daily reality at Company A.

Management summary

Company A is the market leader in vegetable seed treatments. Their position is built on innovation, industry-leading quality, and high on-time delivery. In the past five years, growth has come to an halt, as the market became saturated. The current goal of Company A is to retain their market share by providing high on-time delivery and quality, while reducing costs. Company A is struggling with the demand uncertainty and strong seasonality inherent in the vegetable market. They experience a reduced on-time delivery of 91% during peak season, down from 95% during low season. Within Company A, this problem is often attributed to a lack of capacity. However, they are currently unable to form a coherent strategic and tactical capacity plan to address this issue. The main research question of this research is as follows.

How can machine- and operator capacity planning deal with seasonal and uncertain demand to improve on-time delivery in a cost-efficient way?

We find that the underlying problem is three-fold. First, the current demand forecasts are unreliable; they assume demand is equal to last year without considering uncertainty. Second, the calculation of capacity demand from product demand is inaccurate, as the number of orders is not linearly related to processing time, thus product demand forecasts cannot be used for capacity planning. Third, capacity decisions are considered individually, resulting in a misaligned capacity plan. For example, when making machine procurement decisions, Company A currently does not consider that machine capacity can be increased through additional shifts. Strategic capacity decisions (i.e. machine investment) are related to tactical decisions (i.e. workforce planning).

We designed a capacity planning model that addresses these three problems. Our model first calculates historical capacity demand from historical sales orders to address the problem of inaccurate capacity demand calculations. By including all relevant details, some of which are unique to Company A, the capacity demand is calculated accurately.

Second, our model uses the historical capacity demand from the calculation part to generate future capacity demand forecasts. The forecasting part is based on the Error-Trend-Seasonality model by Hyndman et al. (2008). We use one-tailed upper prediction intervals to reflect demand uncertainty and seasonal components to reflect demand seasonality. The prediction interval covers the actual capacity demand with a certain coverage probability. We define several capacity demand scenarios that each correspond to a coverage probability. We add two methods to the forecasting part to include judgmental forecasts: adjustment factors and future sales orders. An adjustment factor is the expected percentage change from history, as caused by external factors such as legislation or technological innovation. The second method uses future sales orders, when there is no historical data, such as for new products. The prediction intervals, seasonal components, and judgmental methods address the problem of unreliable demand forecasts by modeling demand variability.

Third and finally, our model determines the optimal capacity plan that deals with uncertain and seasonal demand in a cost-efficient way. The optimization part is primarily based on models by Bihlmaier et al. (2009) and Fleischmann et al. (2006). We use three capacity demand scenarios as input (i.e. coverage probabilities of 50%, 70%, and 90%) to generate three alternative capacity strategies. The strategic and tactical capacity decisions are jointly optimized to find the optimal capacity strategy for each scenario. Each capacity strategy is evaluated by fixing the strategic decisions and optimizing the tactical decisions for a certain scenario.

The results of our model answer the research question. We find that the current capacity levels are sufficient to fulfill capacity demand with a coverage probability of at least 90%. Moreover, we find that Company A can realize the same coverage probability at a lower cost, by reducing machine investment and increasing the use of workforce flexibility. Specifically, by using double shifts in all months except August and September, the number of C0414 dryers can be reduced to 12 and the number of P100 coating pans can be reduced to 6. Company A needs only one coating pan of all other types. Compared to the current situation, the total savings over ten years is €874,000 for the most conservative strategy (90%) and €1257,000 for the 70% capacity strategy. We find that the judgmental forecasts do not change the optimal capacity strategies, because the impact of judgmental forecasts on capacity demand is too small.

To deal with demand seasonality, Company A must have sufficient machines to deal with peak demand using double shifts, while reducing workforce flexibility in case of low demand to save costs. The coverage probability of prediction intervals for capacity demand forecasts is an intuitive and practical way to deal with demand uncertainty. Using this method, Company A can decide on the trade-off between coverage probability and costs. We recommend to use either a 70% or 90% coverage probability.

From a practical perspective, the results are especially useful for the replacement of dryers, which Company A aims to finish in 2023. Company A can purchase 12 C0414 dryers, instead of 14, while maintaining a high coverage probability. The largest savings can be realized when replacing coating pans. We recommend Company A to reevaluate the capacity plan every year, for which we designed a simple to use dashboard to update data, run the model, and view the results. Company A can use this tool for tactical workforce planning as well, by fixing strategic decisions to the current situation. We recommend two future research directions for Company A. First, a more advanced scheduling method and tool can improve on-time delivery and enable Company A to increase the utilization of machines and operators. Second, demand smoothing can further reduce the need for machine capacity. For example, the peak capacity for C0414 dryers is in March. If this can be smoothed towards April, where demand is much lower, Company A can satisfy demand with fewer machines.

Our research and model make three contributions to theory. First, the use of a detailed calculation model for capacity demand has shown to be useful when a piecewise linear transformation from product- to capacity demand is not accurate. Furthermore, this method can be used to calculate the capacity demand for new machines with different characteristics, based on historical sales orders. Second, adjustment factors and future sales orders are practical methods to include external factors not reflected in historical data. These methods enable Company A to determine the number of machines for upcoming new products and assess the impact of, for example, legislation. While the impact of these judgmental forecasts is currently low, there have been various cases in the past where these methods would have been very useful. Third and finally, the prediction interval for capacity demand is an intuitive and practical method to consider demand uncertainty in strategic capacity planning. In our literature review we found no research that uses prediction intervals for strategic capacity planning. Instead of the common stochastic models that use scenarios with probabilities to generate one optimal capacity strategy, prediction intervals allow us to generate alternative capacity strategies for various coverage probabilities, using a simpler linear model. This enables companies to make a trade-off between increasing the certainty of having sufficient capacity and the associated costs.

Table of Contents

1 Introduction ... 8
1.1 About Company A ... 8
1.2 Problem introduction ... 9
1.3 Problem identification ... 10
1.4 Research questions and approach ... 12
1.5 Scope ... 14
2 Current situation ... 15
2.1 Production processes ... 15
2.2 Capacity demand ... 20
2.3 Current capacity planning ... 22
2.4 Conclusion ... 25
3 Literature review ... 26
3.1 Forecasting ... 26
3.2 Capacity planning ... 33
3.3 Conclusion ... 38
4 Model design ... 39
4.1 Model overview and motivation ... 39
4.2 Model description ... 44
4.3 Capacity demand calculation model ... 47
4.4 Strategic capacity planning model ... 49
4.5 Disaggregating capacity demand ... 53
4.6 Conclusion ... 54
5 Model results ... 56
5.1 Model validation ... 56
5.2 Capacity strategies for various scenarios ... 63
5.3 Impact of judgmental adjustments ... 68
5.4 Compare capacity strategies ... 69
5.5 Sensitivity analysis ... 70
5.6 Disaggregating capacity demand ... 74
6 Conclusion ... 77
6.1 Conclusion ... 77
6.2 Contributions to theory ... 78
6.3 Recommendations for practice ... 79
7 References ... 81
8 Appendix: model implementation choices ... 85

1 INTRODUCTION

The purpose of this research is to advise Company A on the capacity planning of machines and operators. We perform a quantitative study to create a strategic capacity plan that deals with demand seasonality and uncertainty that Company A faces. Our approach involves the forecasting of capacity demand and optimization of the capacity plan. The client is the operations management team. They want to use the results of our research for upcoming machine investment- and workforce decisions. We execute the research in collaboration with the supply chain manager and planning manager.

Section 1.1 introduces the company Company A. Section 1.2 introduces the problem and Section 1.3 identifies the core problem. Section 1.4 describes the research approach and research questions. Section 1.5 discusses the research scope.

1.1 ABOUT COMPANY A

Company A is a seed treatment company, leading in high-end vegetable seed treatments and expanding in field crop seed treatments. Company A has about 450 employees working all over the world, of which about 200 are based in the headquarters in Enkhuizen. Revenue was 27.7 million euros in 2018. In Enkhuizen, the main activities are various treatments of vegetable seeds on a make-to-order basis. Customers deliver their proprietary seeds, which are enhanced by Company A and then sent back to the customers. The most important treatments are priming, coating and upgrading. Figure 1-1 shows one step of each treatment. These treatments use patented technologies and materials, developed by Company A's R&D for specific seed types.

Figure 1-1. The main treatment processes: priming, coating and upgrading. (Incotec, 2020)

Around 1970, Company A's inventions were revolutionary for the agricultural sector: crop yield and quality increased, while the amount of chemicals required decreased. Company A's yearly growth was about 20% for years on end. To meet the demand, high capital investments were made and the number of employees was increasing rapidly. However, customers and competition started to catch up, developing their own seed treatments. As the market became saturated, Company A's growth halted around 2014. These days Company A remains market leader in the vegetable crop market, retaining about 50% market share of the outsourced seed treatments. Company A is in the premium segment; their treatments are still considered as industry standard.

This is where Croda stepped in, a chemical company that acquired Company A in 2015. Croda helped Company A to reshape their business strategy, which can be summarized as follows.

1. Develop the most sustainable and environmentally-friendly treatments
2. Expand in the field crop market by developing treatments for field crops
3. Retain market share and improve margins in the vegetable crop market through operational excellence and cost reduction

Croda and Company A have already taken several steps to realize the strategy. Regarding operational excellence and cost reduction, the most important step was to integrate Company A in Croda SAP (i.e. Croda's ERP software) in 2018. This enabled Company A to automate many business processes, resulting in significant cost savings and improved operational performance. Our research is another step in this strategic direction; it contributes to operational excellence and cost reduction with the purpose of retaining market share and improving margins.

1.2 PROBLEM INTRODUCTION

The strategic goal behind this research is to retain market share and improve margins in the vegetable crop market. To retain market share, new customers need to be attracted and current customers need to be retained. The key question here is: what attracts and keeps customers? In the seed treatment market, the answer is, in order of importance: quality consistency, delivery performance, and price.

Company A defines quality as the percentage of seeds that grow according to plan. Quality consistency depends on the process design by R&D and process control by operators. Process control has been an issue a few times, with costly consequences. However, this issue is out of scope for this research, due to the biology expertise required to understand the issues.

Delivery performance is crucial, because most customers have a time window of a few weeks between harvesting and sowing season. Customers deliver their harvested seeds and need them back before the sowing season. When Company A is unable to enhance the seeds in this time window, customers lose an entire season, which is extremely costly. To reduce risks, Company A's largest customers use both in-house treatment and outsourcing. Smaller customers do not have the scale for in-house production. Therefore, delivery performance is even more critical for these customers. If delivery performance is too low, customers will move to competitors or increase in-house production.

Company A measures delivery performance using the 'on-time delivery' metric, which they define as the fraction of orders that are delivered no later than the requested delivery date. The requested delivery date is provided by customers upon ordering. Company A aims for an on-time delivery of 95% in each month. Additionally, we define the 'almost-on-time delivery' metric as the fraction of orders that are delivered at most a week later than the requested delivery date. When orders are a few days late, it is usually agreed upon with the customer. For example, delaying an order such that it can be shipped with another order for the same customer. The almost-on-time delivery shows more serious delivery issues, because these are more than a week late.

Figure 1-2 shows the on-time delivery and almost-on-time delivery for 2018 and 2019. We observe a performance difference in low season (April through October) and peak season (November through March). On average, the on-time delivery is 96% in low season and 91% in peak season. June 2019 is an outlier for on-time delivery, but not for almost-on-time delivery. This outlier is most likely caused by deviating from the requested delivery date in agreement with customers, which is not a serious problem. Company A aims for an on-time delivery of 95%, thus is not meeting their target in peak season.

Action problem
On-time delivery was 91% on average during peak season in 2018 and 2019, which is well below the target on-time delivery of 95%.

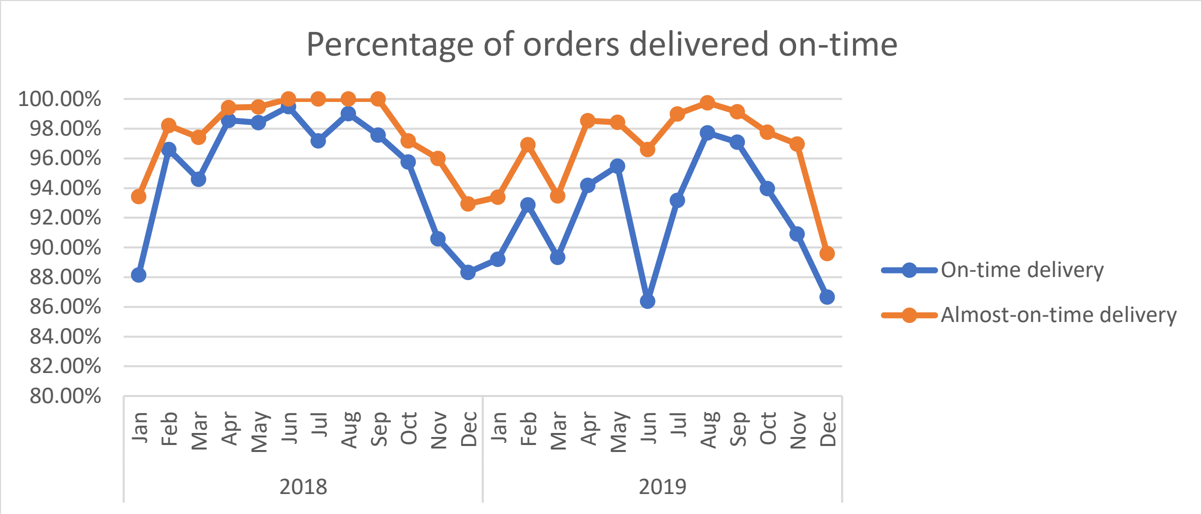

Figure 1-2. On-time delivery and almost-on-time delivery as % of all orders

Price is less important to customers than quality consistence and delivery performance. The added value of Company A's seed treatments far outweighs the costs. To illustrate, the value of one kilogram of tomato seeds is five times more than one kilogram of gold. However, prices are under pressure due to an increasing market maturity and competition. For that reason, cost reduction is the main way to improve margins for Company A. The action problem must be addressed in a cost-efficient way, otherwise Company A could simply double capacity levels to resolve most delivery issues.

1.3 Problem identification

To address the action problem effectively, the core problem must be identified. A problem cluster is useful to identify the core problem. Figure 1-3 visualizes the problem cluster. An important property is that the core problem must be influenceable (Heerkens & Van Winden, 2012). The causes from the problem cluster that are not the core problem are listed below.

1. Quality issues during production is one of the causes of late delivery. Recall from Section 1.2 that we consider quality issues as out of scope, due to the biology expertise required to address these issues.
2. The planning department schedules orders within two workdays of receiving them. They use a backward scheduling method, working back from the requested delivery date. Once an order is scheduled, which is the reservation of a timeslot for the required resources, this is not changed. The reason is that it is a manual and time consuming task to change the schedule in the current ERP system, which is SAP. Company A is currently not interested in changing the scheduling process and systems, because of the costs and risks associated with such a change.
3. Rejecting orders to improve delivery performance is not a feasible alternative. Company A forms partnerships with customers for many years. Rejecting an order hurts the partnership, as many customers rely on Company A's treatments.
4. Demand peaks can cause capacity shortages for resources that have low usage during other times of the year. This demand seasonality is part of agriculture; crops grow in specific time windows (i.e. seasons). Recall from Section 1.2 that orders must be processed within these time windows, thus demand smoothing through delaying is restricted. Furthermore, it is against Company A's business strategy to delay orders; they distinguish themselves by being a flexible partner.

5. Demand peaks cannot be addressed with inventory. Customers deliver their proprietary seeds after harvesting them and retrieve them right after treatment, it is not possible for Company A to have seeds in inventory.
6. From a capacity point of view, flexibility is limited because operators require at least two and up to 24 months of training. They learn sensitive information during this training. That is why each operator is a permanent employee.

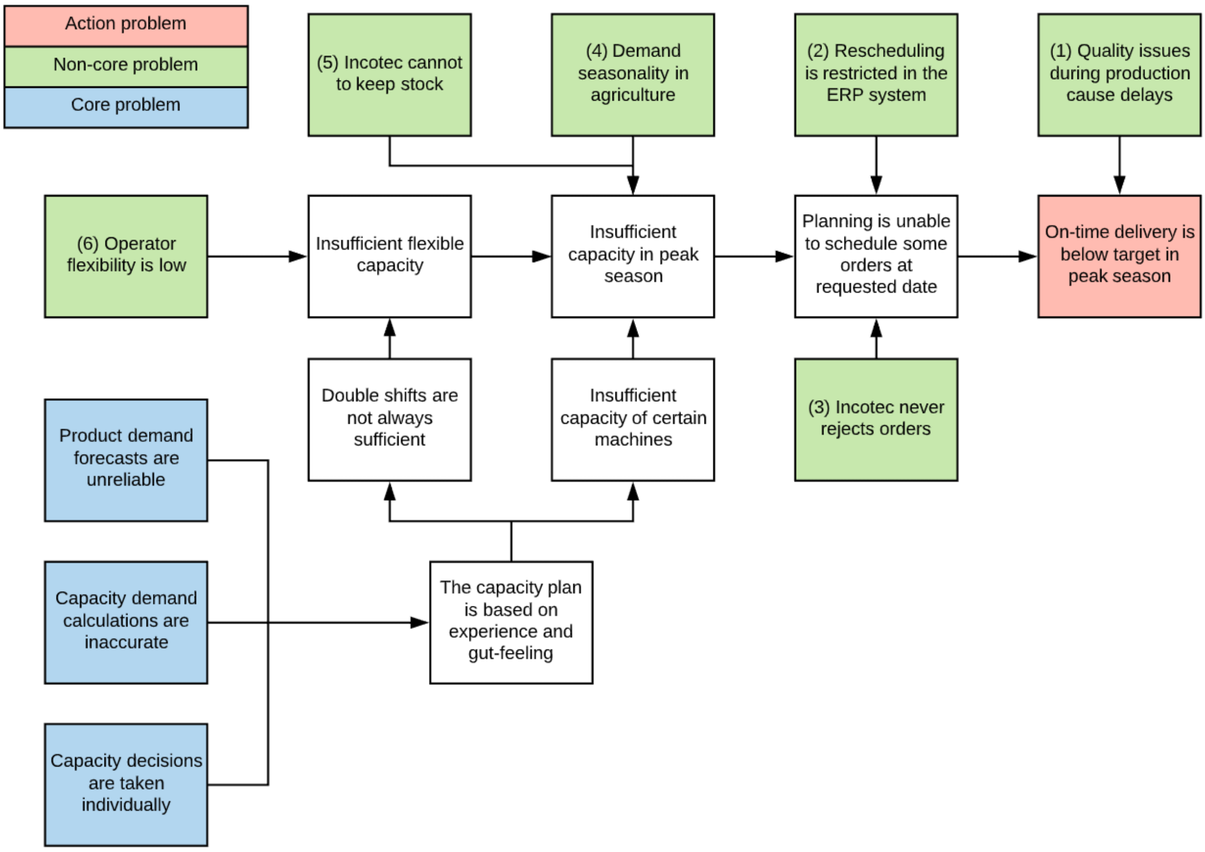

Figure 1-3. Problem cluster that identifies the core problem for the action problem.

Now the causes that are out of influence have been described, we identify the core problem. We find that the core problem is three-fold: unreliable product demand forecasts, inaccurate capacity demand calculations, and misaligned capacity decisions. We define capacity demand as the processing time required of each machine- and operator-type to satisfy product demand. In the remainder of this thesis, we refer to treatments as products, simply because Company A does so as well.

Company A's product demand forecasts are unreliable. Company A has sold 161 different products since 2015, most of which are sold infrequently (<5 times per year). For each product, Company A assumes demand for each month is equal to the same month previous year. This is the seasonal naïve method. This method does not provide accurate forecasts, because of demand uncertainty. Demand depends on harvest quantity and timing. These are different each year due to, for example, weather conditions. Another source of uncertainty is competition, which causes the customer portfolio to change each year. As an alternative to seasonal naïve forecasts, the top five customers provide judgmental forecasts that consider external factors (e.g. weather). However, the accuracy of these judgmental forecasts remains inconsistent.

Even if demand is known, the calculation of capacity demand is inaccurate. Currently, Company A measures capacity demand as the number of orders. Each order consists of a product type (i.e.

treatment) and seed quantity to be treated. The required resources and processing times vary depending on the product type and seed quantity. Therefore, the capacity demand for 100 orders for coating machine can be completely different than for dryer machines, depending on the product type and seed quantity of each order. The quantity of each product type depends on the season. For example, in winter the demand can be 80% chicory and 20% lettuce, while in summer 50% tomato and 50% lettuce. That is why some machines are only used during a few months of the year.

To clarify, the capacity demand cannot be calculated from the number of orders for each product type using a linear formula. The reason is that the seed quantity of each individual order determines the machine type and processing time. The total seed quantity cannot be used either, because orders for the same treatment must never be combined. The seeds within each order are unique, even for individual customers.

Due to unreliable product demand forecasts and inaccurate capacity demand calculations, Company A has been unable to create an aligned capacity plan. Instead, capacity decisions are currently taken individually. For one department, the demand planner decides on hiring decisions, while the production manager decides this for another department. For each machine investment, a new project team is set up to decide on capacity levels. The consequence is that Company A is unable to provide the capacity to fulfill demand in a cost-efficient way.

We summarize the three core problems using the following definition of the core problem.

Core problem
Unreliable product demand forecasts and inaccurate capacity demand calculations leave Company A unable to create a capacity plan that deals with seasonal and uncertain demand in a cost-efficient way.

1.4 RESEARCH QUESTIONS AND APPROACH

1.4.1 Main research question

The action problem is below-target delivery performance during peak season. The core problems underlying the action problem are unreliable product demand forecasts, inaccurate capacity demand calculations, and a misaligned capacity plan. Based on the action problem and core problems, we define the main research question as follows.

Main research question
How can machine- and operator capacity planning deal with seasonal and uncertain demand to improve on-time delivery in a cost-efficient way?

1.4.2 Research approach

The research approach describes how we answer the research question. The core problems are the starting point, from which the research is structured in three steps, as visualized in Figure 1-4. First, the capacity demand must be calculated more accurately. We calculate historical capacity demand in this first step to be able to forecast future capacity demand in the second step. Second, instead of product demand, the capacity demand must be forecasted. Forecasting product demand is not viable, because there is not enough demand data to forecast the demand and seed quantity distribution for each product. This seed quantity is needed to calculate capacity demand. Finally, the optimal capacity plan can be determined, based on the capacity demand forecasts and other relevant parameters.

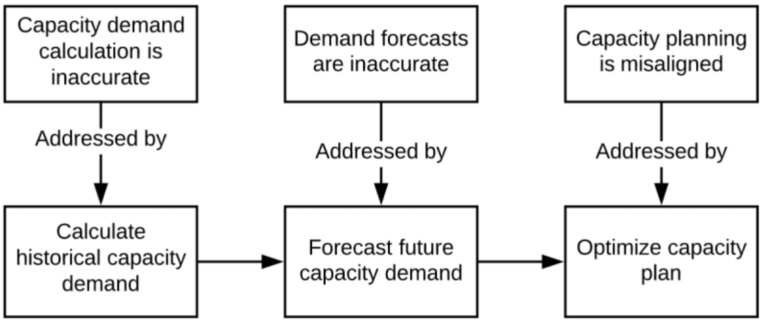

Figure 1-4. The research steps that follow from the core problems.

We structure our research in the following parts: problem identification; current situation; literature review; model design; model results; conclusion and recommendations. We combine this structure with the three research steps to form the research approach. Figure 1-5 visualizes this combination. Note that calculating historical capacity demand is not discussed in the literature review, because it requires a calculation model tailored to Company A's production processes.

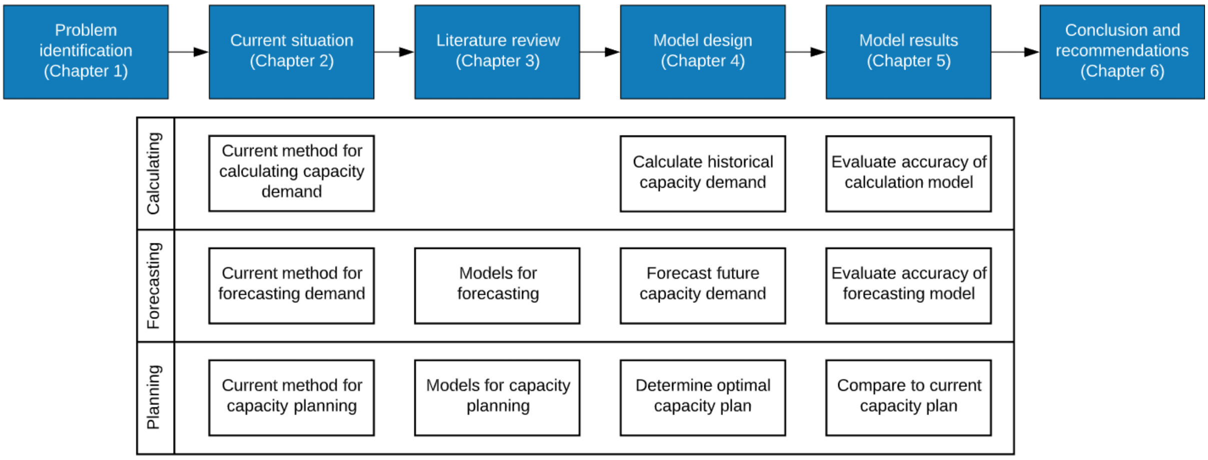

Figure 1-5. The research approach.

1.4.3 Research sub-questions

For each part of the research approach, we list the research sub-questions below. To answer these questions, interviews have been held with people at many different positions in Company A, including upper management and operating personnel. In addition, we analyze data from SAP and Excel to support these interviews and answers the questions.

Chapter 2 describes the current situation, where we answer to the following questions.

2.1 How are Company A's production processes currently organized?
2.2 How does Company A currently determine capacity demand?
2.3 How does Company A currently make capacity decisions?

Chapter 3 is a literature review, where we answer the following questions.

3.1 What are the top performing forecasting models from literature that use historical data to model uncertainty and seasonality?
3.2 What forecasting methods are available in literature that use human judgment?
3.3 How should forecasting performance be measured, according to literature?
3.4 What frameworks are available in literature to classify capacity planning models?

3.5 What capacity planning models are available in literature for strategic capacity planning with workforce flexibility?

Chapter 4 concerns the model design, which consists of three parts, as described in the research approach.

4.1 How can the capacity demand be calculated more accurately?
4.2 How can the forecasting models from literature be applied to Company A?
4.3 How can the capacity planning models from literature be applied to Company A?

Chapter 5 discusses the model results, where we answer the main research question through the following four questions.

5.1 How accurate can our model calculate capacity demand?
5.2 How accurate can our model forecast capacity demand?
5.3 How does the capacity strategy from our model compare to the current capacity plan?
5.4 What is the sensitivity of the model regarding parameters subject to uncertainty or change?

Finally, we conclude our research with recommendations on how Company A can integrate the model in their organization.

6.1 How can Company A integrate the designed model for future decision making?

1.5 Scope

Before diving into the analysis of the current situation, we first define the scope of this research. To fulfill demand, Company A depends on the capacity of machines and operators. The operator capacity impacts the machine capacity. Therefore, to create a useful capacity plan, both machines and operators must be included in this research. Chapter 2 discusses this in detail.

Machine investments have a high financial impact and machines operate up to 20 years, which is strategic capacity planning. Workforce planning at Company A concerns a one-year horizon and has a medium financial impact, which is tactical capacity planning. (Slack & Lewis, 2011) Therefore, this research concerns strategic- and tactical capacity planning. Operational planning activities, such as scheduling, are not part of this research.

Within Company A Enkhuizen, there are nine production processes, each with their own machines and operators. In terms of strategic importance and cost, the most important processes are coating, priming, and drying. To be able to complete this research in six months, the scope only includes coating and drying. Priming is not in scope, because there is currently insufficient data to accurately calculate capacity requirements for priming machines. Company A ensures that priming capacity is never a bottleneck, because margins for priming treatments are very high. The other production processes are not a capacity bottleneck, thus can be left out of scope without any issues.

2 CURRENT SITUATION

In Chapter 1, we identified the three core problems. In this chapter, we aim to understand how the current organization and processes contribute to the core problems. Section 2.1 describes the production processes. Section 2.2 describes how Company A currently determines capacity demand, which is the basis for capacity decisions. Section 2.3 describes the processes for machine and operator capacity decisions.

2.1 PRODUCTION PROCESSES

In this section we answer the following research question.

(Q2.1) How are Company A’s production processes currently organized?

Capacity planning is the alignment of capacity and demand (Slack et al., 2013). Therefore, the available and required capacity must be known for each machine and operator. The following questions summarize the information required of each production process to plan capacity.

1. What resources are required in the process?
2. How to measure the capacity of the resources?
3. How can the capacity be increased or decreased?
4. What parameters must be considered for the capacity decisions?

Before answering these questions in detail, it is useful to have an understanding of the production processes at Company A.

2.1.1 Overview of production processes

Company A offers a variety of products, each product is a combination of seed treatments. Figure 2-1 gives an overview of the processes and possible production paths at Company A Enkhuizen.

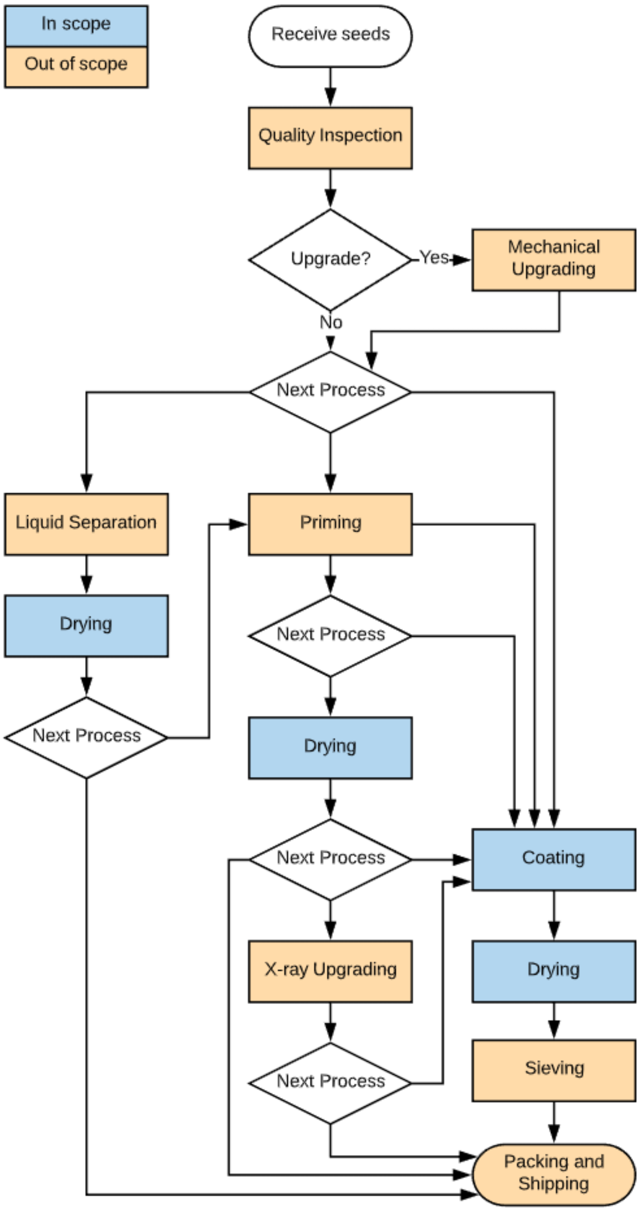

Figure 2-1. Production process for seed treatment in Enkhuizen.

To help understand the process, each processing step is described briefly.

1. When the seeds are received, they first undergo quality inspection. The seeds must meet some quality standard pre-treatment, such that Company A can guarantee a post-treatment quality level.
2. Sometimes the seeds must undergo mechanical upgrading, which is the separation of good and bad seeds based on weight.
3. Liquid separation is another way of separating good and bad seeds, using liquids with different densities.
4. Drying is necessary to preserve seed quality, by drying at the right temperature, humidity, and duration. The duration varies from half an hour up to ten hours. Drying can be used at multiple stages in the production.
5. Priming is the process of putting the seeds in a liquid for a specific time and temperature. The required time ranges from a few hours to multiple weeks, depending on the seed and treatment type. It is a proprietary technology to improve seed quality. Seeds will

germinate faster and more evenly. Furthermore, the percentage of seeds that germinate is increased. The plants from primed seeds are also less susceptible to stressful conditions, such as extreme weather.

6. X-ray upgrading is another way to separate good and bad seeds, by analyzing the embryo within the seed, using x-ray machines. The process takes one to up to four hours.
7. Coating is the process of adding powder and liquid to create a layer around the seeds. This is done by hand in a coating pan by a trained coating operator, who carefully controls the dryness and thickness of the seed coating. The process takes around 2.5 hours per batch. The advantages of a coated seed are a more efficient sowing process using mechanical planting equipment and the addition of crop protection products, nutrients and biologicals. This greatly reduces the amount of chemicals required for the farmers and protects seeds from harmful effects. A specific coating color can be added, which is branding for customers and improves visibility for farmers.
8. Sieving is a quality check, by sieving out coated seeds that are not the right size. Sieving is also done during coating. The sieving process is an extra quality check.
9. Packing and shipping is always the last step. The seeds are packed in either bags or tins of the requested size. Some customers pick up their own seeds, while others are shipped by Company A.

Of all processes, only coating and drying are in scope. As explained in Chapter 1, these processes are of strategic importance and represent the majority of the production costs. Priming is also of strategic importance, but we decide to leave it out of scope, due to the complexity and confidentiality of the process. In the remainder of Section 2.1 we describe coating and drying in detail.

2.1.2 Resources for coating and drying

The coating and drying process require several machines and trained operators. For both coating and drying operators, there are flex-operators who are available to jump in when demand exceeds available capacity. Table 2-1 lists the number of regular operators and flex-operators currently available.

The coating process is done by a coating operator with a coating pan. Coating operators require at least three months training for the basic product type. Additional training is provided depending on product demand. Drying operators are tasked with loading, configuring, and unloading the dryers. Drying operators require only two months training, no crop specific training is required.

Process	Type	Number available
Coating	Regular	11
Coating	Flex	3
Drying	Regular	1
Drying	Flex	2

Table 2-1. Number of regular and flex operators for coating and drying processes.

The coating pans vary in type and size. There are three types of coating processes: regular, rotary, and film. A treatment uses either regular or rotary coating. The optional film coat is applied after regular or rotary coating. The size of the pan determines the minimum and maximum seed quantity that can be processed. All available coating pans are listed in Table 2-2, with the number currently available. Each pan requires auxiliary equipment, such as a sieving system. Auxiliary equipment is considered part of the coating pans for the remainder of this thesis.

Technical name	Type	Diameter (cm)	Number available

PE-LR300	Rotary	300	1
PE-HS500	Rotary	500	2
PE-P055	Regular	55	3
PE-P060	Regular	60	2
PE-P090	Regular	90	1
PE-P100	Regular	100	10
PE-P160	Regular	160	2
PF-P070	Film	70	1
PF-P100	Film	100	1
PF-P120	Film	120	1
PF-RD500	Film; Rotary	500	1

Table 2-2. List of available coating machines.

There are different types of dryers, which are used at different stages in the process, for different treatments, or for different batch sizes. Table 2-3 lists the currently available dryers.

Resource name	Number available
PD-100	1
PD-101	1
PD-102	1
PD-151-2	2
PD-C0103	3
PD-C0414	11
PD-CS1-2	2
PD-S1-4	4

Table 2-3. List of available dryers.

2.1.3 Capacity decisions for coating and drying

Capacity decisions concern either an increase, decrease, or replacement of capacity. We distinguish between strategic decisions and tactical decisions. The strategic capacity decisions for Company A are solely machine procurement. Machine tooling is only applicable to priming, which we decided to leave out of scope (Section 2.1.1). Other issues, such as transportation and product allocation, are not relevant in this single site case. The tactical capacity decisions concern workforce planning.

Machine procurement is the simplest way to increase or replace machine capacity. Decreasing machine capacity through selling is not interesting for Company A. The machines are difficult to sell due to their specificity. Machine capacity can also be increased through workforce flexibility measures, such as double shifts and overtime. That way, machines can be used for more hours per week, resulting in a capacity increase. More specifically, double shifts effectively double machine capacity, because the machines are used for 15 hours a day instead of 7.5. The impact of overtime depends on how much operators work overtime.

To increase operator capacity, coating operators can be hired or the number of flex-operators can be expanded. Note that flex-operators can work at any department within Company A, such as R&D or Supply Chain. The only requirement is that they have the proper training, and are able to temporarily leave their regular work when necessary. Operator capacity can also be increased through overtime: working on Saturdays. Double shifts do not affect operator capacity. Operator capacity can be decreased through either retirement, dismissal, moving to another department, or leaving to work for a competitor.

The most recent strategic capacity decisions have been to set up an organic production line in Enkhuizen, for which a new dryer and coating pan has been procured. Other recent projects have been in new markets, such as Malaysia for rice seeds. The most important upcoming decision for Enkhuizen is the replacement of several dryers, which must be replaced before 2023. On a longer term is the replacement of coating pans, which must be replaced before 2026. These replacement decisions are the focus of this thesis, while considering the necessary capacity changes to meet capacity demand and improve delivery performance.

Recall from Section 1.3 the core problem of misaligned capacity planning. This section discussed the relevant capacity decisions that can be aligned, and how workforce planning (i.e. operators) and strategic capacity planning (i.e. machines) impact each other. Section 2.2 discusses how Company A currently makes these decisions.

2.2 CAPACITY DEMAND

In this section, we answer the following research question.

(Q2.2) How does Company A currently determine capacity demand?

Company A uses four methods to determine capacity demand. Two of the methods focus on historical data: scheduled processing times and number of orders. The other two methods focus on human judgment to make assumptions about the future: long-term and short-term judgmental forecasts.

2.2.1 Scheduled processing times

The first method Company A uses to determine capacity demand is based on scheduled processing times. Currently, Company A uses this method to make machine procurement decisions. We discuss the capacity decisions in Section 2.4.

Recall from Section 1.3 that the planning department schedules an order within two workdays of receiving an order, using a backward scheduling method. The schedule is a set of timeslots for each machine and operator that is required to fulfill the order. The schedules are stored in SAP. Company A obtains the historical capacity demand by summing the scheduled timeslot for each machine and operator. For example, Figure 2-2 shows the historical capacity demand for coating operators in 2018 and 2019, based on scheduled processing times.

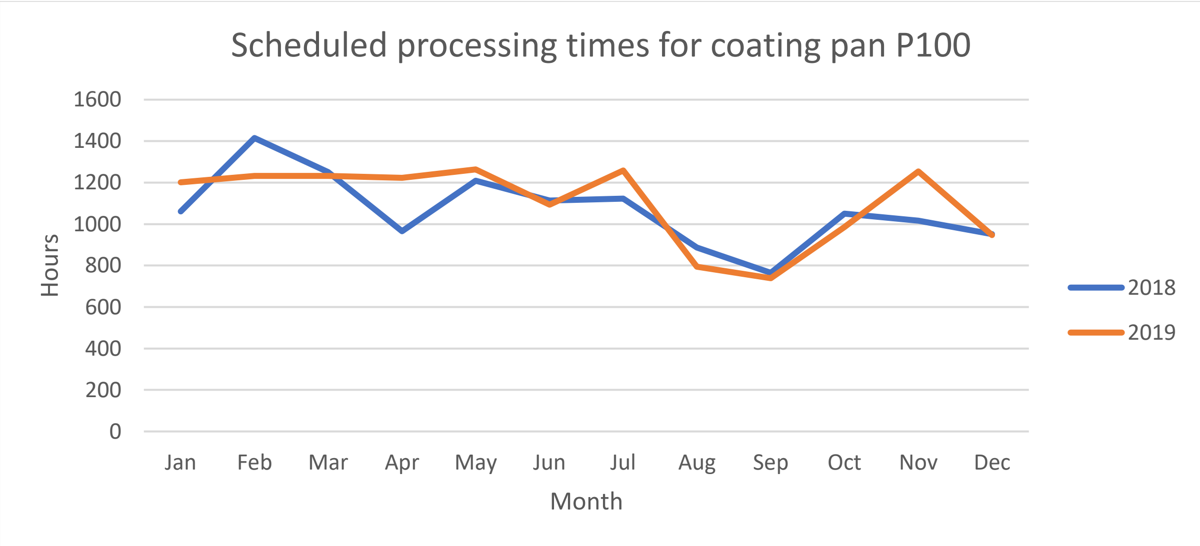

Figure 2-2. Historical capacity demand for coating pan P100.

This method has two issues. First, future demand can deviate strongly from historical demand, especially on the long-term. Figure 2-2 shows significant differences between two succeeding years. Looking 5 to 10 years in the future, these differences are likely to be larger. This method does not offer ways to model expected demand changes, such as market developments or new products, for example. Second, scheduled processing times are available since 2018, the year that SAP was integrated. Therefore, this method provides little insight in long-term developments, such as trends or historical variance.

2.2.2 Number of orders

The second method Company A uses to determine capacity demand is based on the number of orders. Currently, Company A uses this method for workforce planning.

The number of orders are an alternative for scheduled processing times. They are the basis for workforce planning, together with short-term judgmental forecasts. The number of orders are obtained by counting the number of sales orders per month for all product types. These sales orders are stored in SAP. Figure 2-3 shows the number of orders for coating. To make capacity decisions, the capacity demand is calculated using a simple formula. The capacity is on average 3 orders a day per operator. This is multiplied by the number of operators and number of days in a month. For example, there are 694 orders and 30 days in April 2019. Then Company A needs the following number of operators.

$$required\ number\ of\ operators = \frac{694\ orders}{30\ days\ *\ 3\ orders\ per\ day} = 7.7$$

So the capacity demand for operators is 7.7 in April 2019. When Company A has more operators available than required, they can be assigned to training or other activities.

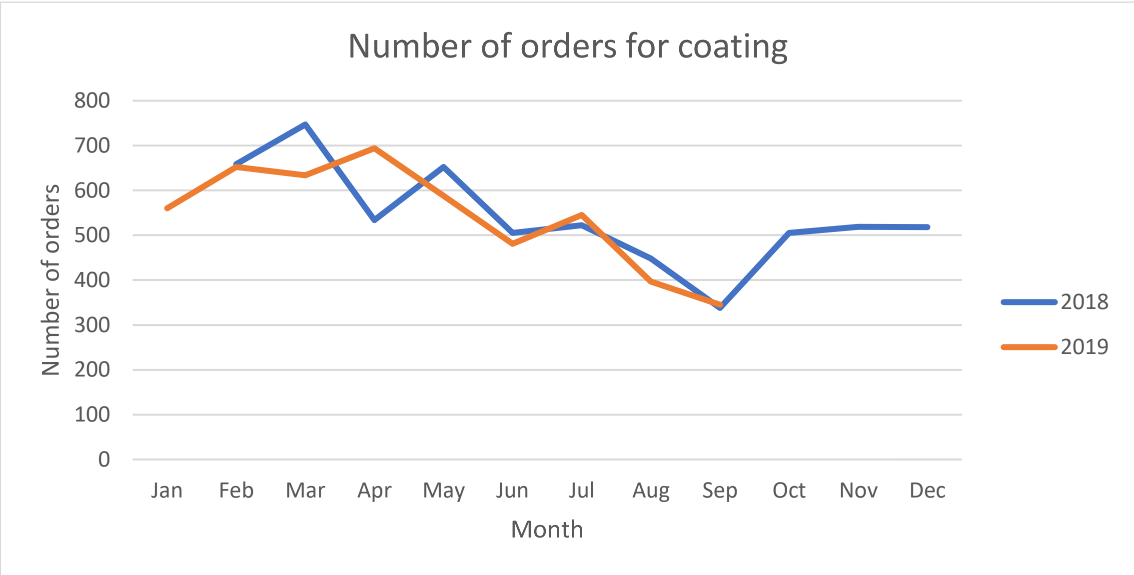

Figure 2-3. Number of orders for coating.

The advantage of this method over scheduled processing times is that it is directly related to product demand. If an additional 50 orders are expected, the planning can be adjusted accordingly. However, the main issue with this method is that the number of orders is a poor indicator of capacity demand, because the processing times and required machines vary for different products. The number of orders only provides a good indication when machines are used for every order with a fixed processing time. That is usually not the case. Figure 2-4 shows the capacity demand in hours, obtained from the scheduled processing time method, and the number of coating orders. It is immediately clear that the number of orders is not accurate for this machine type.

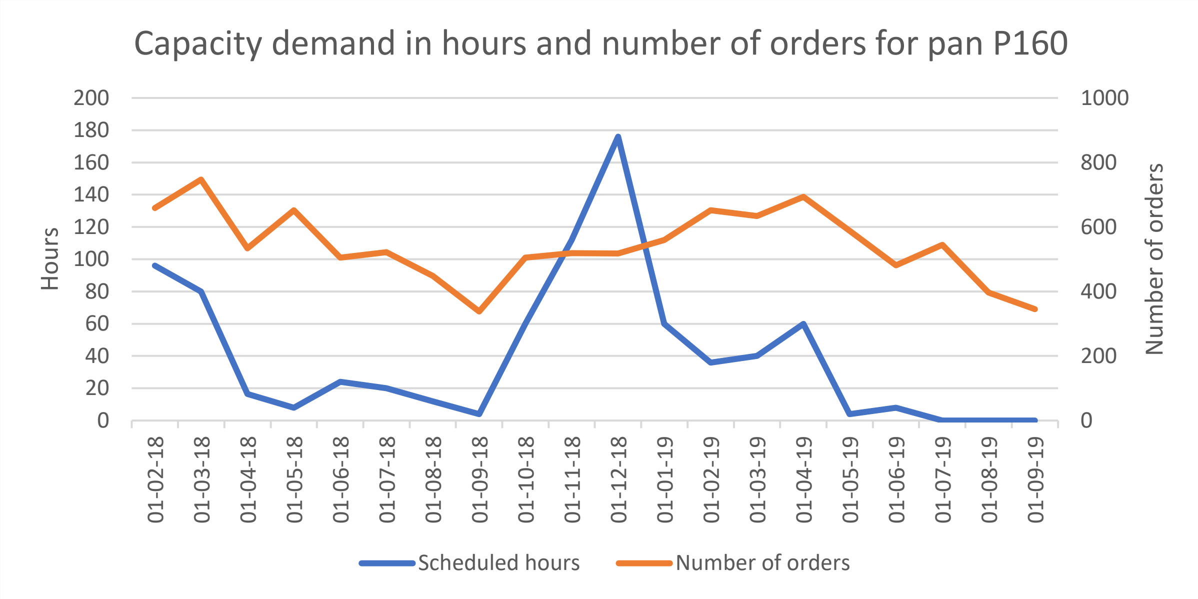

Figure 2-4. Capacity demand in hours and number of orders for pan P160.

2.2.3 Judgmental forecasts

Company A uses two types of judgmental forecasts: short-term and long-term. Short-term forecasts are provided by customers for the next 6 months each quarter. Only the top 5 customers in terms of revenue provide these forecasts for the most frequently bought products. The forecasts are expressed in total seed quantity. Currently, Company A is not able to use these forecasts for capacity planning, because the total seed quantity does not accurately translate to capacity demand. To illustrate: 1,000 seeds in one order requires less processing time and different machines than 100 seeds in 10 orders.

Long-term judgmental forecasts are provided by marketing and account managers. When setting up a new production line, marketing and account managers are asked to provide a sales prognosis. There is currently no formalized method on the contents of a sales prognosis and how it is translated to capacity demand. As discussed in Chapter 1, such capacity decisions are based on gut-feeling and experience, often resulting in overcapacity.

2.3 Current capacity planning

In this section we answer the following research question.

(Q2.3) How does Company A currently make capacity decisions?

2.3.1 Strategic capacity planning

Strategic capacity planning at Company A concerns machine procurement. The lifespan of machines is between 10 and 15 years, therefore the planning horizon is 10 years. Figure 2-5 shows the process for machine procurement decisions. For each machine procurement decisions, a project team is created, with a senior project manager overseeing the project. The decision makers are the operations manager, maintenance manager, and marketing.

The machine procurement process is initiated by a need. This need can be identified by marketing or maintenance. For example, the introduction of a production line for organic seeds was initiated by marketing, because it was market driven. The procurement of new dryers was initiated by maintenance, because the current dryers cause quality issues and maintenance costs are increasing. When initiated by marketing, they provide a long-term judgmental forecast. When initiated by maintenance, the supply chain provides historical capacity demand based on scheduled processing times.

The next step is for the project manager to select a number of alternatives, by asking the various machine builders for a quote. One of the alternatives is selected by means of a discussion with production teams, maintenance and the project team. The selection is based on various criteria, of which technical specifications and price are most important. The project team then estimates the number of machines required, based on the judgmental forecasts or historical capacity demand. Finally, the machines are purchased, manufactured, and installed.

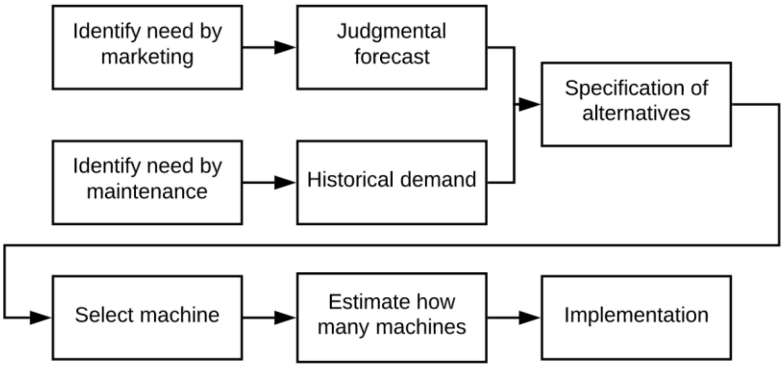

Figure 2-5. Strategic capacity planning process.

Recall from Section 1.3 that the core problems are inaccurate capacity demand calculations, unreliable demand forecasts, and misaligned capacity decisions. We observe these problems in the process for machine procurement decisions.

First the problem of misaligned capacity decisions. A new project team is created for each machine procurement decision, without overarching coordination, resulting in misaligned capacity decisions. For example, Company A does not consider the ways in which workforce flexibility can be used to increase machine capacity, thereby potentially reducing the required number of machines. Second, there is no way to accurately calculate the capacity demand from judgmental forecasts. This is especially troublesome when purchasing machines for new products, for which there is no historical data. Third and finally, capacity demand forecasts based on historical data (i.e. processing times) are unreliable. Recall from Section 2.2.1 that this data provides no way to include expected demand changes and there is little historical data available (starting 2018). Company A has no forecasting method to make use of historical data outside of the seasonal naïve method, where the demand forecast is the demand in the same period last year.

The result is that Company A struggles with estimating how many machines they need. This estimate is usually inaccurate, leading to under- or overcapacity. For example, a HS-500 pan that cost 150,000 euros was only used for two orders each year. The revenue did not come near the cost. Due to their specificity, the resale value of these machines is low.

2.3.2 Tactical capacity planning

Tactical capacity planning concerns workforce planning at Company A. The planning horizon for hiring decisions is one year, while the planning horizon for flexibility measures, is one to three months. The demand planner is responsible for workforce planning, in collaboration with the operations manager. Figure 2-6 shows the workforce planning process.

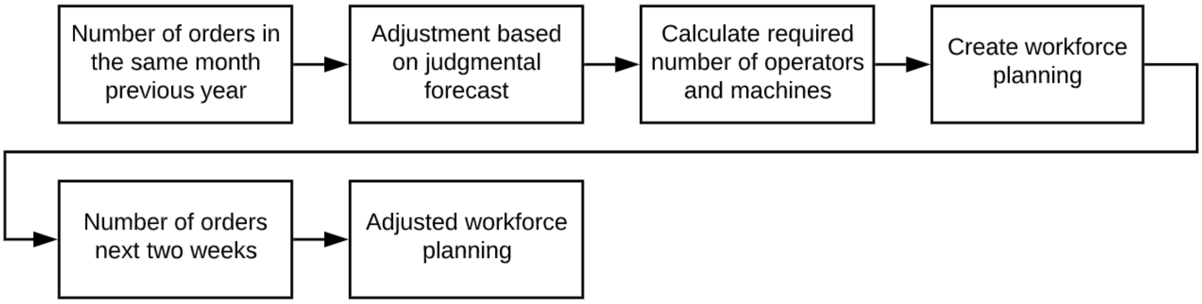

Figure 2-6. Tactical capacity planning process.

Company A uses the number of orders as the basis for workforce planning. Company A uses a seasonal naïve forecast, which means that the demand forecast for the next three months is equal to the demand the same month last year. The second step is to adjust the number of orders based on the short-term judgmental forecast. Recall from Section 2.2.3 that the forecast is in seed quantity, not the number of orders. The adjustment is a rough estimate, based on how the forecast this year deviates from last year. The third step is to calculate the number of machines and operators required from the number of orders. Section 2.2.2 discusses how this is calculated. The fourth step is to create a workforce planning. Depending on the number of operators required, they are assigned alternative activities, such as training.

During peak season, the demand planner revises the capacity plan every week, based on the actual number of orders. These are known about two weeks ahead on average. The demand planner can decide to use overtime or flex-operators to address capacity issues. Figure 2-7 shows the workforce planning for 2018. The spikes in available capacity around week 16 are caused by holidays. We observe the inaccuracies discussed in Section 1.3 and Section 2.2.2: the required capacity is sometimes higher than available capacity. However, in practice this was not the case. It comes as no surprise that the demand planner is struggling with making data-driven decisions, he must rely on experience and gut-feeling.

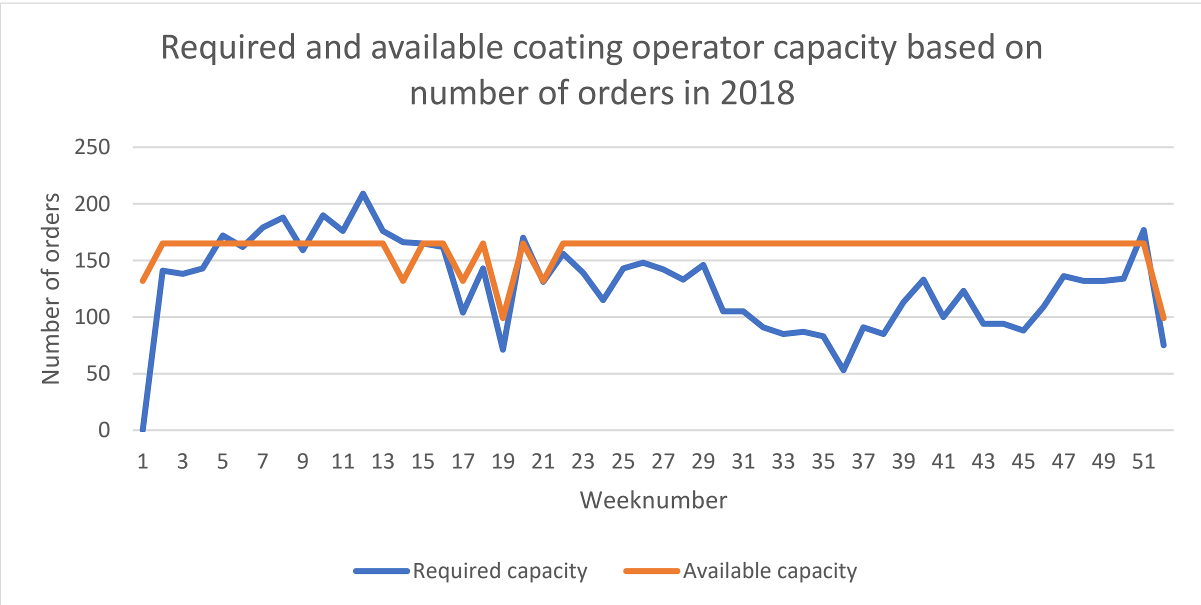

Figure 2-7. Required and available coating operator capacity

The core problems identified in Section 1.3 are observed in the tactical capacity planning. Double shifts are part of workforce planning, but impact machine capacity, not operator capacity. However, machine capacity is not considered in workforce planning, therefore these decisions are simply the same each year: double shifts from November till April. Thus, operator and machine capacity decisions are not aligned. Furthermore, the number of orders is an inaccurate measure. In Figure 2-7, we observe that the required operator capacity is higher than the available operator capacity in weeks 7 through 15. Figure 2-8 shows the required and available capacity for coating operators, but based on scheduled processing time, which is more accurate. Surprisingly, we observe that there was no overcapacity in 2018.

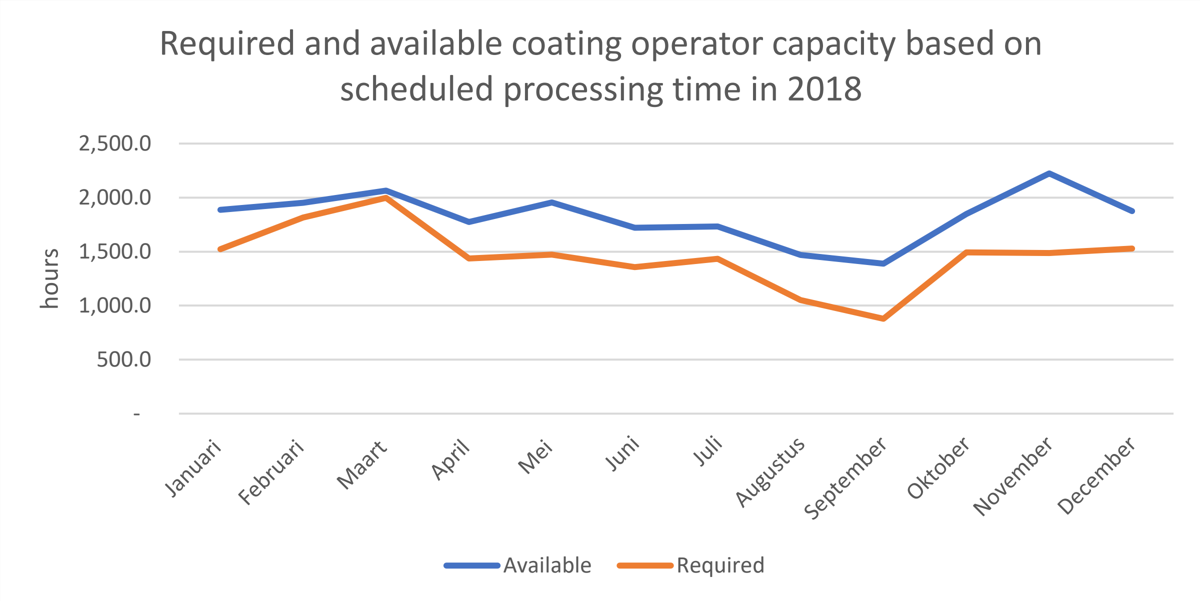

Figure 2-8. Required and available coating operator capacity based on scheduled processing time.

2.4 CONCLUSION

Company A's production processes are currently organized in a clear flow, where coating, drying, and priming are the most important processes. Drying and coating processes each use a variety of machine types, with a total of 19 machine types. The workforce planning concerns primarily the number of operators, flex-operators, single or double shifts, and overtime. (Q2.1)

Company A uses scheduled processing times and the number of orders to determine the capacity demand. The problems with scheduled processing times are a lack of data and the inability to take judgmental forecasts into account. While the number of orders solves the problems, it is much less accurate estimate of the capacity demand than the scheduled processing times. Company A uses the seasonal naïve forecasting method, which is why both measures do not reflect uncertainty in capacity demand. Furthermore, there is currently no structured and accurate way to include judgmental forecasts in capacity demand calculations. (Q2.2)

The core problem of misaligned capacity decisions is visible from the current decisions-making processes: each decision for machine procurement is taken without considering other machine procurement decisions or workforce planning. (Q2.3)

3 LITERATURE REVIEW

We identified three core problems in Chapter 1: inaccurate calculation of capacity demand, unreliable demand forecasts, and misaligned capacity decisions. In Chapter 2, we identified the shortcomings of the current forecasting method and capacity planning method. The seasonal naïve method, based on scheduled processing times, results in a forecast that does not reflect uncertainty. Furthermore, it does not make use of judgmental forecasts, which potentially have a large impact on capacity demand.

In this chapter, we identify the available literature on these subjects and that address these issues. Section 3.1 discusses the forecasting procedure, forecasting models that use historical data, models that use human judgment, and forecasting performance measures. Section 3.2 discusses frameworks for capacity planning literature and relevant strategic capacity models.

3.1 FORECASTING

In this section, we answer the following three research questions.

(Q3.1) What are the top performing forecasting models from literature that use historical data to model uncertainty and seasonality?

(Q3.2) What forecasting methods are available in literature that use human judgment?

(Q3.2) How should forecasting performance be measured according to literature?

We answer research question Q3.1 in Section 3.1.2, where we focus on forecasts based on historical data. In Section 3.1.3 we answer research question Q3.2, concerning forecasts based on human judgment. Finally, we answer research question Q3.3 in Section 3.1.4, by describing forecasting performance measures.

3.1.1 Forecasting procedure

One of the simplest forecasting procedures consists of five steps: problem definition; gathering data; preliminary analysis; choosing and fitting models; using and evaluating the models. (Hyndman & Athanasopoulos, 2018) This is a generic approach, useful in most situations. However, it does not take into account the specific contexts in which the forecast will be used. Chopra and Meindl describe a forecasting procedure for capacity planning, which is more useful for forecasting in the context of this research. The steps in their procedure are the following. (Chopra & Meindl, 2013)

1. Understand the objective of the forecast. A good definition includes how the forecast will be used, who will use it, and the role of the forecast in the decision-making process.
2. Integrate forecasting throughout the supply chain, the forecasts used should be consistent with each other. For example, the forecast for strategic capacity planning must be consistent with the forecast used for workforce planning.
3. Identify demand characteristics. Demand can show seasonality and trends. Demand can also depend on external factors, such as promotional activities.
4. Decide the level of aggregation. A higher aggregation lowers forecast error, but the level of aggregation must be detailed enough to make accurate decisions. For example, a

company can decide to use product groups, when a group of products show similar demand patterns or rely on the same external factors.
5. Performance measurement. Finally, the forecast performance must be evaluated using measures that are relevant for the objective of the forecast.

In Table 3-1, we show the similarities and differences between Hyndman's and Chopra's procedure. Most notably, choosing and fitting models is not an explicit step in Chopra's procedure, while it should be part of the forecasting procedure according to Hyndman.

Step	Hyndman	Chopra & Meindl
1	Problem definition	Understand the forecast objective
2		Integrate forecasting throughout supply chain
3	Preliminary analysis	Identify demand characteristics
4		Decide level of aggregation
5	Choosing and fitting models	
6	Using and evaluating models	Performance measurement

Table 3-1. Similarities and differences between Hyndman's procedure and Chopra's procedure.

3.1.2 Forecasting based on historical data

3.1.2.1 Classification of forecasting methods

There is a variety of forecasting methods that make use of historical data. Forecasting methods can be categorized in the following six categories. (Hyndman & Athanasopoulos, 2018)

1. Simple methods. Most notable examples of simple methods are average, naïve, seasonal naïve, and drift. Average takes the mean of all historical data as forecast. Naïve sets the forecast to be equal to the value of the last observation. Seasonal naïve sets the forecast to be equal to the last observed value from the same season of the year. Finally, the drift method allows the forecast to increase or decrease over time. This change is the 'drift', which is the average change observed in historical data.
2. Time series regression models. The basic idea is that the time series to forecast has a linear relationship with another known time series. For example, to forecast the monthly sales, the advertising spend can be used as a predictor.
3. Exponential smoothing models. Forecasts are obtained through weighted averages of past observations. The weights are decreased exponentially as observations are further from the present, such that more recent observations have a higher weight. Some of the most well-known models are Holt's method (Holt, 1957) to model trends, which was extended to become the Holt-Winter's method (Winters, 1960) to include seasonality.
4. ARIMA models aim to describe autocorrelations in data. Box and Jenkins popularized these models in 1970. Its most recent edition remains the main reference for ARIMA modelling. (2015)
5. Advanced methods is a collection of methods that build on the aforementioned methods, while making use of recent advancements in other fields. Some examples are neural network models, bootstrapping and bagging, and machine learning. (Bergmeir, 2016)

3.1.2.2 Comparing forecasting methods

We are interested in the top performing forecasting models from literature. To separate the good from the bad forecasting models, the M-Competition was introduced. (Makridakis, et al., 1982) The fourth edition, the M4 Competition, took place in 2018. Over 250 universities and companies have enrolled in this competition to submit their forecasting models. The competition used 100,000 time series to compare the forecasting performance of the models. Table 3-2 denotes the number of series per data frequency and domain. Company A would be the 'Industry' category, with monthly time intervals. There are 10,017 of such time series included, making the results relevant for our research.

Time interval between successive observations	Micro	Industry	Macro	Finance	Demographic	Other	Total
Yearly	6,538	3,716	3,903	6,519	1,088	1,236	**23,000**
Quarterly	6,020	4,637	5,315	5,305	1,858	865	**24,000**
Monthly	10,975	10,017	10,016	10,987	5,728	277	**48,000**
Weekly	112	6	41	164	24	12	**359**
Daily	1,476	422	127	1,559	10	633	**4,227**
Hourly	0	0	0	0	0	414	**414**
Total	**25,121**	**18,798**	**19,402**	**24,534**	**8,708**	**3,437**	**100,000**

Table 3-2. Number of M4 series per data frequency and domain. (Makridakis, et al., 2020)

Statistical benchmarks were used to compare the submitted methods. Two aspects of a forecast were measured: Point Forecasts (PFs) and Prediction Intervals (PIs). A PF is the best educated guess of the actual value, which is usually the main focus of a forecast and therefore used for the main ranking. PI gives an interval within which the actual value is expected with a specified probability. A prediction interval can be written as

$$\hat{y}_{T+h|T} \pm c \cdot \hat{\sigma}_h$$

where $\hat{y}_{T+h|T}$ is the point forecast and $\hat{\sigma}_h$ is the standard deviation of the forecast distribution. The multiplier c depends on the coverage probability. Assuming normally distributed forecast errors, the value of c can be obtained from the standard normal distribution for the specified coverage probability. When forecasting one step ahead, the standard deviation of the forecast distribution is almost equal to the standard deviation of the residuals. However, as the forecast horizon h increases, $\hat{\sigma}_h$ increases as well. The uncertainty becomes larger as forecasts are made further in the future. The main difference between confidence intervals and prediction intervals is that prediction intervals must account for both the uncertainty in knowing the value of the population mean and data scatter, so the prediction interval is always wider. (Hyndman, 2013) The M4 includes a second ranking for prediction intervals.

The winner of the M4 Competition was Smyl from Uber Technologies, who mixed exponential smoothing methods with Recurrent Neural Networks. (2019) Another high-ranking method uses a combination of 7 statistical models, while using machine learning to assign weights for the averaging of these methods (Montero-Manso, et al., 2020). These top performing methods are a combination of advanced techniques and exponential smoothing, according to Hyndman's aforementioned classification.

Next up in ranking are ETS (Error, Trend, Seasonality), ARIMA (AutoRegressive Integrated Moving Average), and Theta. The implementation of ETS used in the competition is the state space approach. (Hyndman, et al., 2008) These three were the top performers in the previous M3 competition, and were used as benchmarks in the M4 competition. ETS is based on exponential smoothing. ARIMA describes autocorrelation in the data. Theta is a decomposition model, which has been shown to be equivalent to a specific exponential smoothing model. (Hyndman & Billah, 2003)

An advantage of ETS and ARIMA over Theta is that these can be used to calculate prediction intervals. ETS has been shown to produce more accurate prediction interval than ARIMA. (Makridakis, et al., 2020) Prediction intervals are very useful for capacity planning. Recall from Chapter 2 that the main issue is that the current forecasts do not reflect uncertainty. Company A is not interested in having capacity to meet the expected demand, but to have sufficient capacity with a higher probability. A one-tailed upper prediction interval of capacity demand specifies how much capacity is required to fulfill all demand with a specified probability. In Chapter 4 we motivate our decision to use the ETS model for our research. In the remainder of this section we discuss the ETS model in more detail.

3.1.2.3 ETS state space approach

Exponential smoothing models are the basis of the ETS state space approach. Table 3-3 shows a classification of exponential smoothing models, by trend and seasonality component. Gardner discerns two trend types: additive is a constant trend and additive damped adds a parameter that "dampens" the trend to a flat line some time in the future. (Gardner, 1985) He also discerns two seasonality types: additive is a constant seasonality component and multiplicative scales with the forecast level.

Trend Component	Seasonal Component		
	N (None)	**A (Additive)**	**M (Multiplicative)**
N (None)	(N,N)	(N,A)	(N,M)
A (Additive)	(A,N)	(A,A)	(A,M)
Ad (Additive damped)	(Ad,N)	(Ad,A)	(Ad,M)

Table 3-3. A two-way classification of exponential smoothing models. (Gardner, 1985)

Some of these models are well-known by other terms. For example, (N,N) is the simple exponential smoothing model, and (A,A) is the Additive Holt-Winter's model. For each model, there is a set of forecast equations and smoothing equations. The state space approach introduces an underlying statistical model for forecasts, such that prediction intervals can be calculated, which involves a third component: the error. The classification in Table 3-3 is extended to include the error component. Errors can be additive (A) or multiplicative (M), thereby doubling the number of models to 18. Each model can be described by state space equations, consisting of a measurement equation and a set of state equations. The measurement equation describes the observed data. The set of state equations describe how the level, trend, and seasonality change over time. These state space equations exist for all 18 models.

For all 18 models, it is assumed that the residuals are normally and independently distributed with mean 0 and variance σ^2. Or in short: $e_t = \varepsilon_t \sim NID(0, \sigma^2)$. The prediction interval is calculated using
$\hat{y}_{T+h|T} \pm c \cdot \hat{\sigma}_h$ for most models. The forecast variance formulas are known for additive models and several multiplicative models. For some ETS models, there are no known formulas. In these cases, Monte Carlo can be used to simulate future sample paths and calculate prediction intervals from percentiles of these sample paths.

The ETS state space approach uses maximum likelihood methods to estimate smoothing parameters and initial states. The restrictions for the smoothing parameters are: $0 < \alpha < 1$; $0 < \beta < \alpha$; $0 < \gamma < 1 - \alpha$. The best model is the one with the highest predictive accuracy. Several measures of predictive accuracy exist. Maybe the most well-known are R-squared and adjusted R-squared. However, the ETS state space approach uses Akaike's Information Criterion (AIC), which is based on maximum likelihood. AIC is defined as

$$AIC = -2\log(L) + 2k$$

where L is the likelihood of the model, and k is the number of parameters and initial states in the model. The unknown parameters are collected in a vector, for which there is a 'prediction error decomposition' of the likelihood function, which is maximized with respect to the parameter vector. Further mathematical details are too extensive to discuss here, but the main idea is clear: maximum likelihood and AIC are used to select the model and estimate parameters.

To conclude, the ETS state space approach is an accurate and complete forecasting method. It has top-performing accuracy of prediction intervals, rivaled only by Smyl's forecasting method.

3.1.3 Forecasting based on human judgment

3.1.3.1 When to use judgmental forecasts

In Chapter 2, we described how Company A creates and uses forecasts based on human judgment. A variety of methods are available in literature that define how such judgmental forecasts are to be created and used. In their review of progress in judgmental forecasting the past 25 years, Lawrence et al. (2006) highlight the importance and added value of human involvement in forecasting, especially due to knowledge people have that is not reflected in historical data. This contextual knowledge is defined as any information relevant to the forecasting task, other than the time series. (Lawrence, et al., 2006) Basic contextual knowledge can be the underlying meaning of time series data. For example, that an increase in sales revenue is not due to selling more products, but due to a price hike. More advanced domain knowledge concerns causal information. For example, how a year with little rain leads to lower crop yields, and thus lower demand for seed treatments.

Judgmental forecasts without contextual knowledge showcase lower accuracy than statistical models (Carbone & Gorr, 1983) Judgmental forecasting performance is especially lower when estimating trends and seasonality, most notably by dampening up-and down trends. (O'Connor, et al., 1997) Judgmental forecasting improves forecasts in three cases. First, when contextual knowledge represents a component that cannot be modelled by statistical methods, such as promotions and new product launches. Second, when contextual knowledge contains recent information not reflected in historical data. Third, when forecasters can exert control over the demand, such as when sales are increased through promotions or discounts to reach targets.

3.1.3.2 Ways to create judgmental forecasts

Judgmental forecasts can be created in various ways, which generally fall within the following categories. (Hyndman & Athanasopoulos, 2018)

1. The Delphi method assumes that forecasts from a group are more accurate than from individuals. It aims to achieve consensus amongst experts in a structured and iterative way. There are five stages: assemble an expert panel; individual expert forecasts are collected; individual forecasts are combined and summarized; feedback is provided to the experts, who adjust their forecasts accordingly; final forecasts are obtained by aggregating the individual expert forecasts. The forecast-feedback loop is usually iterated several times, until sufficient consensus is reached. (Rowe & Wright, 1999) The main issue with the Delphi method is that it is time intensive, requiring between 5 to 20 experts over several iterations. Group meetings address this issue, but the resulting forecasts are often optimistic and overconfident, due to group dynamics. (Buehler, et al., 2005)
2. Forecasting by analogy is based on comparison to similar cases, such as pricing real estate. A structured approach similar to Delphi has been defined, using an expert panel, identifying as many analogies as possible, then aggregating the analogy forecasts. (Green & Armstrong, 2007)
3. Scenario forecasting provides a method to move beyond point forecasts. Instead, several forecasts are created, each with a probability. For example, a realistic, optimistic and pessimistic scenario can be defined. This method is frequently used in combination with

stochastic optimization for capacity planning, since Eppen et al. introduced it in 1989. (Eppen, et al., 1989)

4. New product forecasting is usually only based on judgmental methods, because historical data is unavailable. The aforementioned three methods can all be used to forecast new product demand. In addition, several forecasting methods have been designed specifically for new products. Sales force composite uses forecasts by each branch or store, which are then aggregated. However, such forecasts are often biased, because the user (e.g. sales manager) is the one generating forecasts, due to self-serving bias by generating low forecasts or optimistic salespeople. Instead of low-level forecasts, the executive opinion method relies on top-level forecasts. Often established in group meetings, the aforementioned bias is often a problem. Finally, customer intentions can be collected through surveys. However, the main problem is that intentions often deviate from actual purchases, depending on the industry and timing of data collection. (Randall & Wolff, 1994)
5. Judgmental adjustments are used when historical data is available, but external factors not reflected in historical data have a significant impact. Adjustments are most effective when there is strong evidence of the need for large adjustments. Small adjustments have been found to decrease accuracy, especially when forecasters read systematic patterns in the noise associated with a series. (O'Connor, et al., 1993) One way is to restrict adjustments to cases where naïve forecasts perform best, or where specific contextual knowledge heavily influences future demand, such as promotional activity. (Goodwin & Fildes, 1999) The adjustment process consists of two stages: first deciding whether a statistical forecast needs adjustments, then estimating the size of the required adjustment. Most research has focused on the first stage. There are currently no methods in literature that have proven to be successful for estimating the adjustment size.

3.1.3.3 Ways to improve judgmental forecasts

The following approaches are used to improve judgmental forecasts. (Lawrence, et al., 2006)

1. Three types of feedback. Outcome feedback informs the forecasters of the latest observations. Performance feedback includes accuracy and bias of past forecasts. Task properties feedback provides statistical information about the task, it is usually given before producing the forecast. Feedback has shown to improve forecast accuracy. (Goodwin & Fildes, 1999) Outcome feedback is least effective, because it is hard to discern random errors from systematic inaccuracies. (Klayman, 1988) Task properties feedback is most effective, because it helps identify incorrect hypotheses of the forecaster. (Balzer, et al., 1989)
2. Decomposition. Decomposition methods split tasks in smaller sub-tasks for forecasters, aggregating the forecasts afterwards. Decomposition does not necessarily improve accuracy, especially when the sub-tasks are more complex or less familiar for the forecaster. (Goodwin & Wright, 1993)
3. Taking advice. Advice is most useful when it comes from independent sources, to prevent bias. Forecasters weigh advice depending on the reputation of the source. (Yaniv, 2004)
4. Bootstrapping and correction. Statistical forecasting methods can be used to find systematic biases in judgmental forecast, then adjust the forecasts for this bias. Theil's method (Theil, 1971) has shown to be effective, by regressing the historic outcomes onto the forecast, then using this to adjust future forecasts. (Goodwin, 2000)
5. Combining forecasts. Combining statistical with judgmental forecasts can be done in several ways. A simple combination method is to take a weighted average of the statistical and judgmental forecast. When combining, it is important that judgmental forecasts are

based on contextual knowledge (Sanders & Ritzman, 1995) and that the person who weighs and combines the forecast must not forecast himself, due to bias. (Harvey & Harries, 2004)

To summarize, the main advantage of qualitative methods is that they do not rely solely on historical data, but make use of contextual data. The main disadvantages are the danger of overconfidence and the dependency on expertise.

3.1.4 Forecasting performance measures

Forecast error measurements can be classified in the following categories. (Scherbakov, et al., 2013)

1. Absolute forecasting errors are expressed in the same unit as the forecast, such as sales revenue. The main advantage it is easy to interpret the measure.
2. Percentage-based errors are relative to the forecasted quantity. The advantage over absolute errors is that these measures can compare accuracy between forecasts with different units.
3. Symmetric errors address the issue that negative errors are often weighed heavier than positive errors, or vice versa.
4. Scaled errors scale the error relative to the error of another method.
5. There are several other measures that cannot be placed in any of the categories above. Such measures are usually created for a specific purpose.

In the absolute category, one of the simplest measures is the Mean Absolute Error (MAE). MAE is intuitive and easy to compute. However, it is scale-dependent, so it cannot be used to compare between series of different units. The Mean Percentage Error (MPE) is useful to compare bias between series of different units. MAE and MPE are calculated as follows, where y_t is the forecast and x_t the actual value.

$$MAE = \frac{\sum_t |x_t - y_t|}{n}, \qquad MPE = \frac{1}{n}\sum\nolimits_t \frac{x_t - y_t}{x_t}$$

The Mean Absolute Percentage Error is one of the most widely used performance measures. However, it puts a heavier penalty on negative errors than on positive errors. This is addressed by the symmetric MAPE (sMAPE) (Makridakis, 1993). The disadvantage is that it takes on extreme values if the actuals are close to zero, and indefinite or infinite if zero. However, it is intuitive and useful for time series where all values are larger than 10.

$$sMAPE = \frac{1}{n}\sum\nolimits_t \frac{|x_t - y_t|}{(x_t + y_t)}$$

To address the issues with actuals close to zero, Hyndman and Koehler (2006) introduce the Mean Absolute Scaled Error (MASE). This measure scales the Mean Absolute Error (MAE) of the evaluated forecast with the MAE of the naïve forecasting method, by dividing the two. The scaled error is less than one if the forecast is better than the naïve forecast, and greater than one if it is worse than the naïve forecast. For seasonal timeseries, the seasonal naïve forecast is used. MASE is calculated as follows, where m is the length of one seasonal cycle. The seasonal naïve forecast is x_{t-m}.

$$MASE = \frac{MAE\ of\ forecast}{MAE\ of\ seasonal\ naive\ forecast} = \left[\frac{\sum_t |x_t - y_t|}{n}\right] / \left[\frac{\sum_t |x_t - x_{t-m}|}{n}\right]$$

A last interesting measure is specifically for aggregate production planning: the Cumulative Absolute Forecast Error (CAFE). (Ha, et al., 2018) It is the product sum of cumulative forecasting error. This measure is optimized for total cost by adding weight factors for backorders and inventory costs. It has been tested and validated with data from the M3 competition.

3.2 CAPACITY PLANNING

In this section, we answer the following research questions.

(Q3.4) What frameworks are available in literature to classify capacity planning models?

(Q3.5) What capacity planning models are available in literature for strategic capacity planning with workforce flexibility?

3.2.1 Classification of capacity planning models

Capacity planning is usually divided in three levels: strategic, tactical, and operational. Table 3-4 lists the timescale and key decision associated with each level of capacity planning. (Slack & Lewis, 2011) Recall from Chapter 2 that at Company A the ability to keep or fluctuate capacity levels (i.e. tactical level) has a large impact on how much capacity is required (i.e. strategic). In Chapter 4 we motivate our decision to focus on strategic capacity planning, while taking the impact of strategic decisions on tactical capacity planning into account. Recall from Section 1.5 that operational capacity planning (e.g. scheduling) is out of scope.

Level	Timescale	Key question
Strategic	Years-Months	How much capacity do we need and where should it be located?
Tactical	Months-Weeks	To what extend do we keep or fluctuate capacity levels?
Operational	Weeks-Hours-Minutes	Which resources are allocated to which tasks and when are they loaded?

Table 3-4. Three levels of capacity planning. (Slack & Lewis, 2011)

Martinez-Costa et al. have created a conceptual framework for capacity planning, based on their review of 57 strategic capacity planning models, with a focus on the manufacturing industry. (2014) They highlight the challenges of long term capacity planning, because such planning aims to integrate traditionally isolated areas, such as new product development and technology selection, as well as strategic and tactical decisions. (Levis, 2004) Figure 3-1 provides an overview of this framework, which consists of three phases: problem definition, model design, and solution procedures. The output is always a capacity plan, sometimes accompanied by a financial plan or product development plan. This framework is useful to identify relevant capacity planning models, based on the decisions and factors considered, and to structure the model design discussed in Chapter 4.

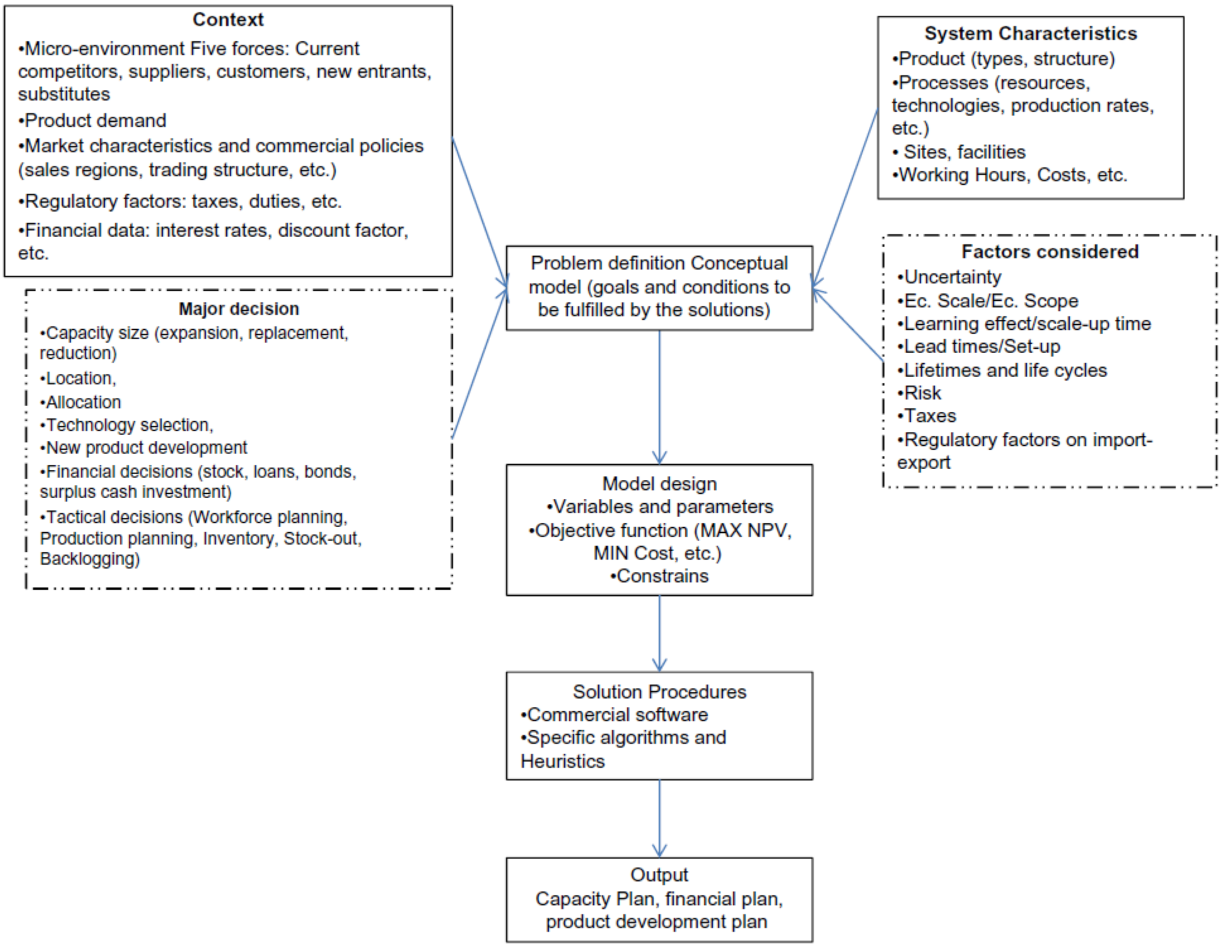

Figure 3-1. A conceptual framework for strategic capacity planning. (Martínez-Costa, et al., 2014)

3.2.1.1 Phase one – problem definition

The problem definition phase starts with identifying external- and internal information (i.e. context and system characteristics). Examples of context are market characteristics and regulatory factors. System characteristics describe the product, process, and facilities involved. We discussed the context and system characteristics in Chapter 1 and Chapter 2. Second, the decisions to make and factors to consider must be described. Finally, these inputs are used to define the problem, which states the goals and conditions of the model. The following decisions are most common in research concerning strategic capacity planning. (Martínez-Costa, et al., 2014)

1. Capacity size. The expansion, reduction, and replacement of capacity size are usually the most important decision. Most models only consider expansion, but recent developments, such as shorter life cycles, have increased interest in reduction. Replacement is especially relevant when physical deterioration has a high impact on capacity, due to inefficiencies, breakdowns, or high operation costs. By considering replacement and expansion simultaneously, scale advantages can be obtained. (Chand, et al., 2000)
2. Location. For multi-site capacity planning problems, location is crucial. There are usually two options for capacity expansion: either with or without new site installation. Then, the main goal is to optimize the transportation costs versus capacity expansion costs. (Chen, et al., 2013)

3. Allocation. When a company has resources that can produce multiple products, the model must allocate resources to products and determine capacity of which resource to expand. For multi-site problems, product quantities must be allocated to production sites as well.
4. Technology selection. When there are alternative technologies available, which can be switched between sites, it is beneficial to jointly optimize the capacity expansion and capacity configuration (i.e. technology mix). (Karabuk & Wu, 2003)
5. New product development. When new products require either the modification of existing or introduction of new resources, it is beneficial to jointly optimize the product portfolio and capacity. (Levis & Papageorgiou, 2004)

Strategic capacity planning and tactical decisions can be closely related. When that is the case, the impact of strategic decisions on tactical decisions should be taken into account. Note that the tactical decisions itself are made at a later stage, but these can be anticipated in strategic capacity planning. The following are two common examples of this interaction.

1. Inventory and backlogging. In sectors where inventory is built up to deal with demand fluctuations, capacity and inventory can be considered substitutes. Joint optimization can result in significant cost savings. (Bradley & Glynn, 2002) Backlogging can be included as well, however this issue is covered in one paper only. (Wang, et al., 2007)
2. Workforce planning. Usually tactical workforce planning options are cheaper than strategic capacity expansion. Workforce flexibility can be used to deal with seasonal demand, for example through one or more shifts, thereby reducing the need for equipment acquisition. To exploit the impact of these tactical options, they must be anticipated in the strategic capacity planning. (Bihlmaier, et al., 2009)

Once the decisions are defined, the factors to consider must be identified. In practice, there can be other relevant factors than the ones listed below. These must be identified for each situation.

1. Uncertainty. Demand uncertainty is most commonly included in strategic capacity plans, due to uncertainty in long-term demand. Other sources of uncertainty can be: capacity, throughput, technology evolution, government policies, and prices. Uncertainty is usually modelled by defining a number of scenarios, which are optimized using stochastic models. These are often two-stage models, where first-stage decisions are typically capacity levels. In the second-stage, recourse actions are taken after uncertainty is realized, such as outsourcing. Models cannot include all source of uncertainty and all possible scenarios, so assumptions must be made carefully. (Rastogi, et al., 2011)
2. Economies of scale. In some industries, such as electrical power, economies of scale are of major importance. To incorporate these, concave or fixed-charge cost functions are used in literature. (Ahmed & Sahinidis, 2003)
3. Lead-times and set-up times. In case of short product life cycles, long investment lead-times are important to consider for the timing of capacity expansion. Set-up times are relevant when they impact manufacturing efficiency. (Rajasekharan & Peters, 2000)

The decisions to make and factors to consider has been used to categorize strategic capacity planning models. (Martínez-Costa, et al., 2014) It is especially useful to identify similar models and gaps in research. In Chapter 4 we use this categorization to select models that serve as a basis for our model design.

3.2.1.2 Phase two – model design

Based on the problem description, a model can be designed. At the most basic level, three types of models can be distinguished: analytical approaches, simulation approaches, and hybrid approaches that combine analytical and simulation approaches. Models based on fuzzy set theory also exist, but are uncommon. (Peidro, et al., 2009) Alternatively, by specifying the type of analytical approach, the model types can be categorized as follows. (Geng & Jiang, 2009)

1. A static capacity model is widely used in practice, due to its ease of use. It is often implemented in Excel. Capacity is calculated using simple formulas, such as the one below. Such methods However, this measure is highly aggregated, thus the method lacks accuracy.

$$machines\ to\ be\ procured = \frac{capacity\ to\ meet\ demand}{one\ machine\ capacity} - original\ number\ of\ machines$$

2. Simulation-based methods are used to optimize capacity levels by evaluating performance for certain capacity levels. That way, capacity uncertainty can be modelled. The accuracy of a simulation model depends on the assumptions and estimates related to arrival distributions and processing time. Detailed and reliable data is required to build an accurate simulation model.
3. Queueing models are an alternative for simulation methods to evaluate performance. They require less data than simulation, but are often mathematically complex models.
4. Linear programming (LP) is simple and can be optimized, but it does not include capacity uncertainty. Deterministic LP models are not robust regarding both capacity and demand uncertainty.
5. Stochastic programming does consider demand uncertainty. However, these models can be difficult to optimize, depending on the problem size. The accuracy of scenarios and associated probabilities is crucial in stochastic models.

Note that this categorization does not include the hybrid approach mentioned by Peidro et al. Mathematical programming methods are by far the most common in literature. (Martínez-Costa, et al., 2014) We hypothesize that this is due to the aggregated nature of strategic capacity planning, while simulation models require more detailed information. However, simulation models are indispensable when capacity uncertainty plays a major role.

3.2.1.3 Phase three – solution procedures

Based on the model design, a solution procedure can be chosen. Static capacity models do not require a solution procedure, only a simple Excel formula.

Simulation-based methods can use a neighborhood search heuristic to determine capacity levels in a trial-and-error way, for example. Simulation is used to evaluate performance for each capacity configuration generated by the neighborhood search heuristic. Another possible procedure is the Genetic Algorithm. Queueing models can use similar trial-and-error procedures, except that they usually require less resources, so more trials can be evaluated in the same computation time.

Deterministic linear programming models can often be solved analytically or with an algorithm. Coin-or CLP is an academic and frequently used open source solver. Commercial alternatives are AIMMS, Gurobi and GAMS, for example. Stochastic programming models require more advanced procedures. Coin-or SMI is an open source solver for these types of models. Commercial alternatives are again provided by AIMMS and GAMS.

3.2.2 Strategic capacity planning models with workforce planning

We motivate our model selection decision in Chapter 4. For now, recall from Chapter 2 that strategic capacity planning and workforce planning are strongly related at Company A. Therefore, we discuss the strategic capacity planning models that anticipate workforce planning in this section.

From the total of 57 mathematical programming models (Martínez-Costa, et al., 2014), only two papers consider workforce planning. Both papers are applied to the automotive industry, which is a capital intensive industry where workforce planning is strongly related to capacity. Furthermore, demand uncertainty and product lifecycles play an important role in capacity decisions in this industry.

Fleischmann's model (Fleischmann, et al., 2006) concerns strategic investments in three departments: body assembly, paint shop, and final assembly. Fleischmann does not use a stochastic model for demand uncertainty, because they "could make no serious assumptions about the probability distribution of the future demand of new products over a 12-year planning horizon." Instead, Fleischmann uses flexibility reserves and demand scenarios to compare different strategies.

Flexibility reserves are modelled by defining a disposable capacity level, which is the maximal capacity minus a flexibility reserve (Figure 3-2). Note that workforce planning is simply modelled as overtime. Demand scenarios are compared for three capacity strategies: a reference strategy, by restricting capacity decisions to the current strategy, and two improved strategies. Improved strategies are obtained by removing some restrictions to capacity decisions, based on expert opinions. The three strategies are optimized for each demand scenario. These results are used by the company to make the final capacity decisions.

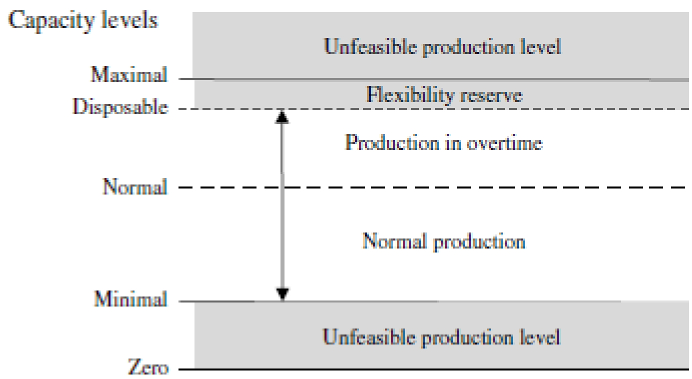

Figure 3-2. Capacity levels in Fleischmann's model. (Fleischmann, et al., 2006)

Bihlmaier's model (Bihlmaier, et al., 2009) concerns strategic flexibility and capacity planning in production networks. Bihlmaier models workforce planning in a more comprehensive way than Fleischmann's model, which only considers overtime. To do this, Bihlmaier uses shift models. Each shift model has three parameters: available capacity, number of employees required, and a cost factor. For each time period, a binary decision variable indicates which shift model is deployed. The shift model costs are added in the objective function. The following constraints are defined for the shift model.

1. Select a shift model with sufficient available capacity to fulfil demand.
2. Ensure that only one shift model is selected in each time period.
3. Ensure that a shift model can only be selected if there are sufficient employees.
4. Determine how many employees must be hired or dismissed.

To model demand uncertainty, Bihlmaier uses a two-stage stochastic model with recourse. This method is used in most strategic capacity models that consider demand uncertainty. (Martínez-Costa, et al., 2014)

UNIVERSITY OF TWENTE.

3.3 CONCLUSION

The top performing forecasting models from literature were hybrid methods, using a combination of statistical methods and neural networks. The runners up were more traditional methods: ARIMA and ETS. To reflect uncertainty, some models are able to generate prediction intervals for a certain coverage probability. (Q3.1)

The forecasting methods that use human judgment are often highly structured, involving multiple forecasts, to address the errors in human judgment. The main advantage and use of human judgment is when external factors not reflected in historical data have a major impact or when historical data is not available. Judgmental methods are used to generate point forecasts, for scenarios, new products, or adjusting statistical forecasts. It is important to use one or more ways to improve judgmental forecasts, of which the most successful are task properties feedback and Theil's method. (Q3.2)

Several measures are available to assess forecasting performance. Scherbakov identified the advantages of the various types of measures. Absolute measures are intuitive, while percentage-based measures are useful to compare different series. Scaled errors improves on percentage-based measures by allowing for low values. (Q3.3)

To classify capacity planning models, several useful frameworks are available. Slack and Lewis provide insight in the differences between strategic, tactical, and operational capacity planning. Martínez-Costa et al. provide insight in the range of decisions and factors that can be considered in strategic capacity planning. Geng and Jiang provide insight in the types of models that are used, which are usually mathematical programming models, but also simulation-based models. (Q3.4)

There are only two strategic capacity planning models available to anticipate workforce planning: Bihlmaier et al. and Fleischmann et al. Anticipating workforce planning is most useful for capital-intensive companies. Bihlmaier et al. model workforce planning through shift models. (Q3.5)

4 MODEL DESIGN

The research goal is to create a machine- and operator capacity plan that deals with seasonal and uncertain demand in a cost-efficient way. In Chapter 1, we identified three core problems: inaccurate calculation of required and available capacity; unreliable demand forecasts; capacity decisions are misaligned. In Chapter 2, we further discussed how the currently used methods lead to the core problems. In our model design we improve on the weaknesses of the current methods identified in Chapter 2, such that we can address the core problems identified in Chapter 1. Additionally, Chapter 2 provides information on the current production process and planning method, which we use to model the reality accurately, such that our model can be used in practice. In Chapter 3, we identified relevant literature, from which we select models to (partially) use in our model. In this chapter we answer the following research questions.

(Q4.1) How can the capacity demand be calculated more accurately?
(Q4.2) How can the forecasting models from literature be applied to Company A?
(Q4.3) How can the capacity planning models from literature be applied to Company A?

In Section 4.1 we provide an overview of the model and motivate our design decisions. In Section 4.2, we describe the steps of the capacity planning model in more detail by providing small examples. In Section 4.3 and 4.5 we explain the calculation model and optimization model, which are used in various steps of our complete model.

4.1 MODEL OVERVIEW AND MOTIVATION

Figure 4-1 visualizes the model design, which consists of seven steps. Our model consists of two parts: first generate capacity demand scenarios (steps 1 through 5), then determine the capacity plan that satisfies capacity demand at the lowest cost for each scenario (step 6 and 7). Recall from Chapter 2 that capacity demand is the processing time, measured in monthly hours per machine- and operator type, required to satisfy product demand. The steps of our model are the following.

1. Calculate historical capacity demand from sales orders
2. Forecast future capacity demand scenarios from historical capacity demand
3. Create judgmental sales forecasts
4. Calculate adjusted capacity demand from judgmental sales forecasts, for which the same model part is used as step 1.
5. Forecast adjusted capacity demand scenarios, for which the same model part is used as step 2.
6. Determine optimal capacity strategy for each scenario using a MIP model
7. Evaluate each capacity strategy for various scenarios

The planning horizon for our strategic capacity model is ten years in periods of one month. Section 4.1.6 motivates this decision. Recall from Chapter 1 and 2 that we decide to forecast capacity demand instead of product demand, because there is too little data to forecast the number of orders and seed quantity of each product type reliably. This information is needed to accurately calculate capacity demand. Therefore, Step 1 is to calculate historical capacity demand, which we use to forecast future capacity demand in Step 2.

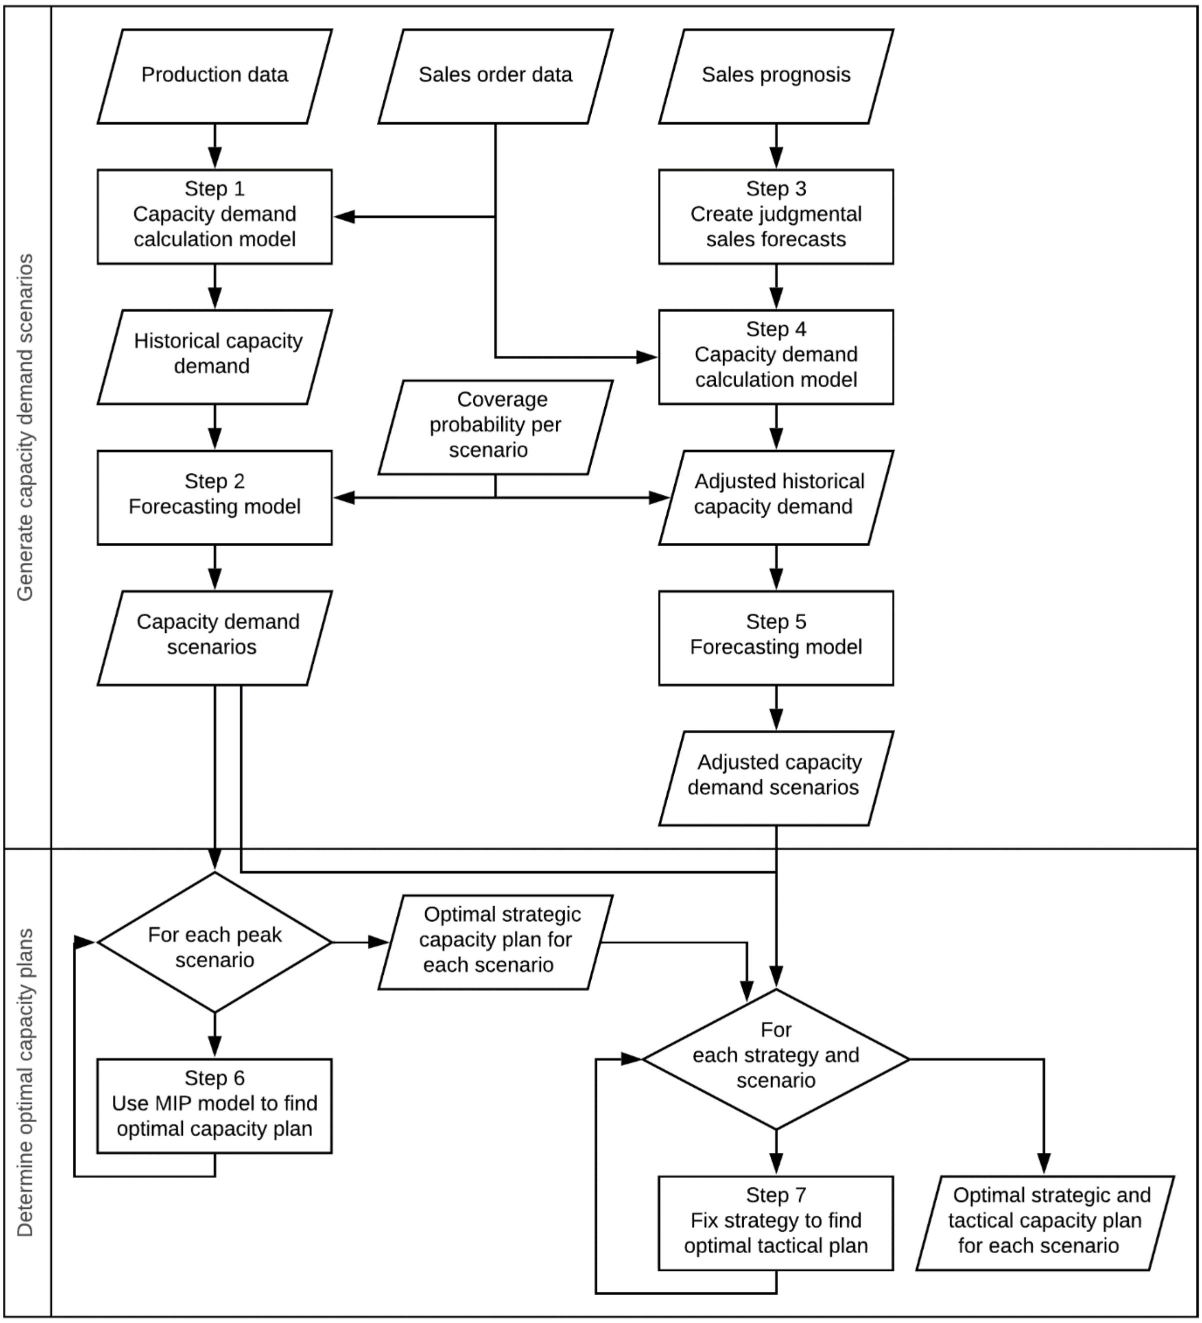

Figure 4-1. Overview of the capacity planning model.

4.1.1 Step 1: Calculate historical capacity demand

The first step is to calculate historical capacity demand, using sales orders and production data as input. We decide to use sales orders instead of Company A's current methods (i.e. scheduled processing times and number of orders, Section 2.2) for four reasons. First, sales order data is available since 2015, while scheduled processing times since 2018. These three years of additional data will result in more accurate forecasts. Second, this calculation method is more accurate than the number of orders, thereby addressing the first core problem of inaccurate capacity demand calculations. Third, this method enables the use of judgmental forecasts, which is discussed in step 3 and 4. Finally, this method enables Company A to retrospectively consider changes in the production process. When Company A introduces a new machine, the model is able to calculate the historical capacity demand as if this new machine would have been used instead of the old one. By updating the production data, the changed process is automatically considered in the calculation model.

We assume that the production data from SAP is up-to-date. Company A uses the production data for various critical business processes, such as scheduling. Company A has strict updating policies when changes are made to products or processes. Therefore, we can make this assumption. We also assume that scheduled processing times in SAP are consistent with actual processing times. Variations in processing times carefully controlled, because deviations endanger product quality. Furthermore, the schedules have been proven to be realistic over the past years. Therefore, we can also safely make this assumption.

Recall from Chapter 1 that piecewise linear transformations cannot be used to accurately calculate capacity demand. Instead, we design a calculation model that considers all relevant variables to calculate historical capacity demand from sales orders. We do not use a model from literature, because the calculations are company-specific.

4.1.2 Step 2: Forecast capacity demand scenarios

The second step is to generate capacity demand forecasts, using historical capacity demand as input. The main issue with the current seasonal naïve forecasting method is that it does not reflect uncertainty. Recall from Chapter 3 that the most commonly used method to include demand uncertainty is through demand scenarios with probabilities, based on data and human judgment, which are input for a stochastic optimization model. (Martínez-Costa, et al., 2014) However, Company A is not able to make such assumptions about the probability distribution of future demand through scenarios. Furthermore, they desire to compare capacity strategy alternatives, such that a trade-off can be made between certainty (of having sufficient capacity) and costs. For these reasons, we design a new method to include demand uncertainty.

We decide to use one-tailed upper prediction intervals (PI) as capacity demand scenarios. Each scenario corresponds with a coverage probability of the prediction interval. For example, a one-tailed upper PI with a coverage probability of 70% means that there is a 70% probability that the actual capacity demand will be less than the PI. This is the type of certainty that Company A is looking for, the coverage probability is an intuitive measure for capacity planning. We assume forecast errors are normally distributed when calculating prediction intervals.

We select a model based on our literature review in Section 3.1. We decide to use the ETS forecasting model, instead of the currently used seasonal naïve forecast, to further address the second core problem of unreliable demand forecasts. We choose ETS over the slightly better performing hybrid forecasting methods, because the underlying exponential smoothing models are widely proven in practice. We choose ETS over ARIMA because ETS generates more accurate prediction intervals. We do not use the Theta method, as it is not able to generate prediction intervals. (Makridakis, et al., 2020)

The uncertainty can become unrealistically large due to being far in the future. To address this issue, we decide to keep the forecast variance constant after the third year, such that the scenarios remain realistic. We choose three years, because Company A wants to purchase machines of one type at most once every three years, due to economies of scale. Therefore, the uncertainty within three years must be considered. We recommend that Company A uses the capacity planning model each year to reevaluate the capacity plan based on the most recent data, to make adjustments where needed.

4.1.3 Step 3: Create judgmental sales forecasts

The third step is to create judgmental sales forecasts, based on the sales prognoses provided by sales, marketing and customers. Recall from Chapter 2 that judgmental forecasts can currently not be accurately translated to capacity demand, which contributes to the core problem of unreliable demand forecasts. To address this issue, we design two methods to include judgmental forecasts: adjustment factors and future orders. These methods are implemented through step 3, 4 and 5.

Adjustment factors are the expected percentage changes from historical demand for specific customers, products, or a combination thereof. Fleischmann et al. (2006) use a similar method, by increasing demand with a percentage for the optimistic scenario. Recall from Section 3.1 that these adjustment factors should only be used for exceptions caused by external factors that are not reflected in historical data. Recall from Section 3.1 that judgmental forecasts are especially important when historical data is not available. For Company A, this is the case for new product demand that requires new machinery. Demand for new products can be expressed as future orders.

Information for both adjustment factors and future orders can be obtained through discussion with sales, marketing, and customers. The accuracy of such judgmental forecasts is crucial. Currently, Company A solely uses outcome feedback to improve accuracy of this information. However, we recommend Company A to use task properties feedback, which literature has proven to be the most effective (Section 3.1).

4.1.4 Step 4: Calculate adjusted capacity demand

The fourth step is to calculate adjusted capacity demand, using sales orders and adjustment factors as input. The calculation model is the same one used in step 2. The capacity demand for specific customers and products is multiplied by the respective adjustment factors and summed to obtain the adjusted capacity demand. By doing this, we assume a percentage change in capacity demand for judgmental adjustments. However, a demand decrease can result in either smaller orders or fewer orders. Both impact capacity demand differently. This is a simplification that must be considered by decision makers when using adjustment factors. The calculation model is also used to separately calculate future capacity demand by using future sales orders as input.

4.1.5 Step 5: Forecast adjusted capacity demand scenarios

The fifth step is to generate adjusted capacity demand scenarios, using the adjusted capacity demand as input. There are two differences with Step 2. First, the model uses adjusted capacity demand as input, which is output of step 4. Second, the model adds the future capacity demand from step 2 to the capacity demand forecasts, to obtain the adjusted capacity demand scenarios.

4.1.6 Step 6: Determine optimal capacity strategies

The sixth step is to determine the optimal capacity strategy given a set of parameters (i.e. input). The most important parameters are capacity demand, available capacity, and cost parameters. Most parameters are based on historical data and actuals, and therefore accurate. We apply a sensitivity analysis for the uncertain parameters in Chapter 5.

We decide to use a mathematical programming model for strategic capacity planning. Alternatively, simulation-based models are sometimes used for strategic capacity planning. Simulation models are especially important when there is capacity uncertainty, such as variable processing times. (Geng & Jiang, 2009) At Company A, processing times are stable, because the processes are predictable. Processing times are prescribed by R&D and closely monitored to achieve consistent quality. Therefore, there is no need to create a detailed simulation model. Furthermore, we would have to make assumptions about the arrival distribution of over 100 different products, which is difficult due to the lack of data and underlying variables. For example, the weather conditions impact each crop in a different way.

4.1.6.1 Anticipating workforce planning

To address the core problem of misaligned capacity decisions, we decide to anticipate tactical decisions (i.e. workforce planning) in our optimization model for strategic capacity decisions. Recall from Chapter 2 that the workforce planning and machine capacity are related. For example, double shifts effectively double machine capacity. In literature, only two strategic capacity models consider workforce planning. (Martínez-Costa, et al., 2014) First, Bihlmaier et al. (2009) consider workforce planning in a two-stage stochastic model, by using shift models. Second, Fleischmann et al. (2006) consider workforce planning in a deterministic model, by using overtime decisions. Their research is the main foundation for our mathematical programming model.

Recall from Section 3.2 that Bihlmaier et al. model workforce planning using shift models, while Fleischmann et al. use overtime. To model workforce planning, we use the number of operators and flexibility measures such as flex-operators, single or double shifts, and overtime.

4.1.6.2 Including demand uncertainty

Our mathematical programming must deal with demand uncertainty. To do this, Bihlmaier et al. define demand scenarios, each with a quantity and probability. They use a two-stage stochastic model to optimize the capacity plan for these demand scenarios. Fleischmann et al. state that they can make no reasonable assumption about the probability distribution of the future demand of products over a 12-year horizon. Therefore, they use a simple scenario technique by defining the expected scenario and a scenario with increased demand (i.e. 30% increase). They use a discrete model to optimize the capacity plan for the expected scenario. They fix the product allocation decisions and optimizes the capacity plan for the increased demand scenario.

To model demand uncertainty, we use the capacity demand scenarios generated in step 2, where each scenario corresponds to a coverage probability. The optimization model generates the most cost-efficient decisions such that the available capacity is sufficient to meet the capacity demand scenario. The output of this step is an optimal capacity strategy for each scenario.

Note that a scenario with 70% coverage probability means that the actual capacity demand will be at most equal to this level, but is most likely less. Therefore, it would be incorrect to optimize the strategic and tactical capacity decisions for this scenario. Instead, we optimize the strategic and tactical capacity decisions for the point forecast (i.e. most likely scenario), and set a constraint that the maximum capacity including workforce flexibility must be sufficient to deal with the capacity demand scenario. Company A is not interested in a capacity strategy with a coverage probability lower than 30%, because it is their business strategy to be a reliable and flexible partner for their customers.

4.1.6.3 Planning horizon and period

Fleischmann et al. consider a 12-year planning horizon, which is used by BMW, the company on which the paper focuses. Bihlmaier et al. state that the planning horizon should cover at least two product life cycles, because they deal with product succession, for which demand overlaps. Therefore, they use a planning horizon of 14 years, as the product life cycle of cars is 5 to 7 years. For Company A, product life cycles are not as important. Innovations revolve around formulation or machine settings. For a new coating, the formulation of the coating powder changes, but the same coating pan can be used. Therefore, we use machine life cycles to determine the planning horizon. At Company A, machines are replaced after about 10 years. Therefore, we decide to use a planning horizon of 10 years.

Bihlmaier et al. use periods of one year to reduce the problem size, as it is a multi-site problem, including transportation and product allocation. Fleischmann et al. also use one-year periods. However, both papers focus on the car industry, where the production rate is relatively smooth throughout the year, to increase efficiency. Recall from Chapter 1 that at Company A, and the agriculture industry in general, production leveling is limited. Treatments for most crops are limited to a small time window, in which Company A must provide sufficient capacity. Therefore, we decide to use periods of one month. The problem size of our model is sufficiently small to use monthly periods over 10 years, because it is a single-site problem.

4.1.7 Step 7: Evaluate capacity strategies

In step 7, we assess how the capacity strategies from step 6 perform in different scenarios. Based on this performance, management can decide between a more expensive and save strategy with a higher coverage probability, or a cheaper strategy. We use the optimization model used in step 6, but we fix the strategic decisions and optimize for the tactical decisions only. That way, we can evaluate how a capacity strategy performs on a tactical level in various scenarios. This method of comparing alternative capacity strategies is similar to Fleischmann et al., except they relax constraints to obtain alternative capacity strategies.

4.2 Model description

4.2.1 Step 1: Calculate historical capacity demand

The first step is to calculate the historical capacity demand from historical sales orders. We measure historical capacity demand per month and for each machine- and operator type. Table 4-1 shows an example of a sales order, which data is retrieved from SAP. Note that seed quantity is given in thousands (t). The model also uses production data from SAP, which consists of ten different data exports, including the bill of materials and processing times.

Customer	Treatment	Requested delivery date	Seed quantity (t)
SeedCompany	Split pill 3.5 let	20-08-19	505

Table 4-1. Example of sales order data as input for the first step of the model.

To calculate historical capacity demand, the model first uses the bill of materials to find all sub-processes that are part of the treatment. Then, for each sub-process, the sales orders are separated into multiple process orders, to deal with the maximum seed quantity constraints of each process. For each process order, the required resources and processing times are obtained. Finally, these processing times are summed by requested delivery date, to obtain the historical capacity demand for each resource type. We explain this calculation in detail in Section 4.3.

The output of the first step is the historical capacity demand. Table 4-2 shows this output, for the example of input in Table 4-1.

Resource	Month	Processing time (h)
Coating operators	01-08-19	2.5
Coating pan P100	01-08-19	2.5

Table 4-2. Example of historical capacity demand as output for the first step and input for the second step.

4.2.2 Step 2: Forecast capacity demand scenarios

The second step is to forecast various capacity demand scenarios, based on the historical capacity demand calculated in the previous step (Table 4-2). We use the ETS forecasting model, as discussed in Section 4.1, to generate scenarios.

The first step of ETS is to select a model type from the Error, Trend, Seasonality taxonomy. The model with the lowest AIC (Akaike's Information Criterion) is selected. Next, the model parameters are estimated using a likelihood function for the vector of the unknown model parameters, based on one-step-ahead prediction distributions. After selecting the model and estimating the model parameters, the point forecasts are calculated for the planning horizon. Finally, the forecast variance and prediction intervals are calculated, based on the coverage probability. The coverage probability for each scenario is provided as input (example in Table 4-3). The output of the forecasting model are the capacity demand scenarios: the forecasted monthly hours per machine- and operator type that are required to satisfy demand (example in Table 4-4).

Recall from Section 4.1 that each scenario corresponds to a coverage probability. Table 4-3 shows the four scenarios we use. These coverage probabilities are intuitive and meaningful for Company A. A company can decide to use different coverage probabilities, depending on the risks they are willing to take and the options they have to deal with capacity shortages.

Scenario	Coverage probability
Pessimistic	30%
Realistic	50%
Optimistic	70%
Very optimistic	90%

Table 4-3. Coverage probabilities for each capacity demand scenario as input for the second step.

Table 4-4 shows the scenarios for April 2020 for coating operators and coating pan P100. The complete forecast covers the entire planning horizon and all machine- and operator types. Recall from Section 4.1 that we keep the forecast variance constant after the third year.

Resource	Month	30	50	70	90
Coating operators	01-01-20	1380.43	1481.74	1583.05	1729.33
Coating pan P100	01-01-20	1048.32	1114.53	1180.75	1276.35

Table 4-4. Example of capacity demand scenarios for one month.

4.2.3 Step 3: Create judgmental sales forecasts

Recall from Section 4.1 that we design two ways to include judgmental forecasts: adjustment factors and future sales orders. The long-term sales prognosis is provided by sales and marketing, which is the starting point for discussion. The adjustment factors are determined in discussion with sales, marketing and supply chain. The adjustment factors can be specified for each product-customer combination, but also for an entire product-range. Table 4-5 shows an example of adjustment factors, where the demand for treatment Split pill 3.5 let of SeedCompany is expected to halve. In this example, the demand for Therm 3 tom is expected to increase by 20%.

Customer	Treatment	Adjustment factor
SeedCompany	Split pill 3.5 let	0.5
All	Therm 3 tom	1.2

Table 4-5. Example of adjustment factor provided by company experts.

Recall from Section 4.1 that future sales orders are useful when there is no historical demand available, such as for new products. These future sales orders are provided by sales and marketing, based on customer intentions. Table 4-6 shows an example of future sales orders. In this example, Company A expects 20 large orders (average seed quantity of 500,000) and 40 small orders (average seed quantity of 100,000) in the first month of January 2020. Company experts, such as account managers, must provide this information for the 10-year planning horizon. Such new product demand forecasts are usually rough estimates, but they do provide insight in the impact on capacity demand.

Treatment	Requested delivery date	Seed quantity t	Frequency
Split Let New	01-01-20	500	20
Split Let New	01-01-20	100	40

Table 4-6. Example of judgmental sales forecast as input for the third step.

4.2.4 Step 4: Calculate adjusted capacity demand

The adjustment factor and future sales orders are both used in a different way to calculate adjusted capacity demand. For adjustment factors, the same calculation model from Step 1 is used, but an extra step is added. Before summing capacity demand, the following formula is applied, where the judgmental factor is 1 if no adjustments are made.

$$adjusted\ capacity\ demand\ = historical\ capacity\ demand\ *\ judgmental\ factor$$

For future sales orders, we provide them as input for the same calculation model from Step 1 instead of historical sales orders. To do this, we duplicate the future sales orders into individual sales orders, based on the frequency. The result is the capacity demand for future months, which we add to the forecast in Step 5.

4.2.5 Step 5: Forecast adjusted capacity demand scenarios

In this step, the same model used in Step 2 is used, except that the input is the adjusted capacity demand from Step 4. The same coverage probabilities are used to generate the respective scenarios. After forecasting the scenarios, the future capacity demand calculated from future sales orders (Step 4) is added to the forecasts.

4.2.6 Step 6: Determine optimal capacity strategies

The sixth step is to determine the optimal capacity plan for each capacity demand scenario. We designed a MIP model that jointly optimizes the strategic and tactical decisions. The most important parameters for the MIP model are the capacity demand (scenario), available capacity parameters, and cost parameters. We discuss this MIP model in detail in Section 4.4. Recall from Section 4.1 that we decide to optimize the strategic and tactical decisions for the realistic scenario (50%), while the maximum capacity including workforce flexibility must be sufficient to satisfy either the realistic (50%), optimistic (70%), or very optimistic (90%) scenario. As the coverage probability increases, the resulting capacity strategy will be safer (a higher probability of having sufficient capacity) but more expensive.

We define the following formulas to calculate available capacity for machines and operators. The Operating Equipment Effectiveness (OEE) is often used for capacity planning, which is a multiplication of the availability, performance, and quality. The availability (or uptime) is the percentage of scheduled time that the resource is available to operate. The performance is the speed at which an operation runs compared to the designed speed. The quality is the percentage of products that are processed correctly the first time round, so without rework or scrap. These concepts can be applied to operators as well.

$$\begin{aligned} &available\ capacity\ of\ machine\ type\ i \\ &\qquad = number\ of\ machines * monthly\ workdays *\ daily\ hours * OEE \end{aligned}$$

$$\begin{aligned} &available\ capacity\ of\ operator\ type\ j \\ &\qquad = number\ of\ operators * monthly\ workdays * daily\ hours * OEE \end{aligned}$$

The output of this step is the optimal strategic capacity plan for each scenario. Table 4-7 shows an example of this output for one machine type and one month. It tells us that the optimal strategy for the realistic scenario (50%) is to have 11 coating pan P100 available, and to purchase 1 coating pan P100 in January 2020.

Resource	Month	Scenario	Max capacity (h)	Number of machines	Purchased
Coating pan P100	01-01-20	50%	1100	11	1
Coating pan P100	01-01-20	70%	1200	12	2
Coating pan P100	01-01-20	90%	1400	14	4

Table 4-7. Example of output of fourth step for one resource and one month.

4.2.7 Step 7: Evaluate capacity strategies

The seventh and final step is to evaluate the strategic capacity plans determined in the Step 6. To do this, we fix the strategic decisions and optimize the tactical decisions, while changing the capacity demand scenario. If the problem is infeasible, the capacity strategy is not able to satisfy the capacity demand of the scenario. This can be the case when the strategy is optimized for a scenario with a lower coverage probability. If the problem is feasible, the result is a tactical capacity plan that is able to fulfill the capacity demand, given the strategic capacity plan.

4.3 Capacity Demand Calculation Model

In this section, we answer the following research question.

(Q4.1) How can the capacity demand be calculated more accurately?

Recall from Section 4.1 and 4.2 that the calculation model is used in steps 1 and 4 of our model. Figure 4-2 visualizes the steps of this calculation model. Note that this description is simplified for the sake of brevity, we focus on the most important steps. We store the input and output in Excel. The calculation model is implemented in R.

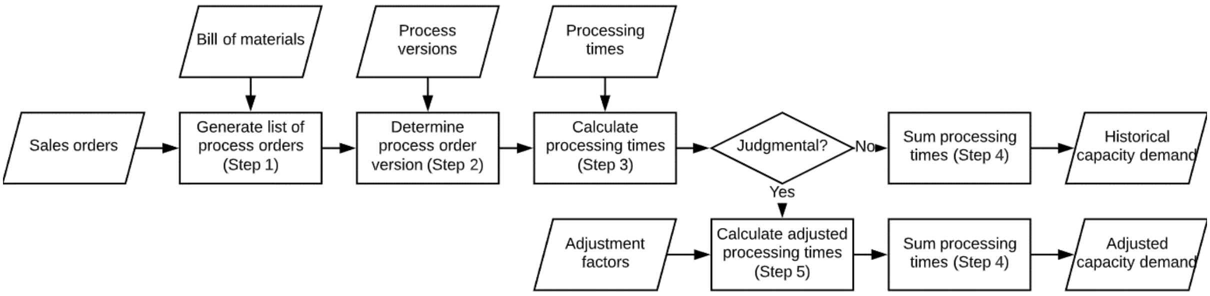

Figure 4-2. Overview of the calculation model.

4.3.1 Step 1: From sales order to process order

Recall from Section 4.2 that the input is a list of the historical sales orders. The calculation model also uses the bill of materials from the production data, which states the processes that must be performed for each product. For each sales order, the bill of materials is used to generate the list of processes to perform. The result is a list of process orders. Table 4-8 and Table 4-9 show an example of one sales order and the resulting two process orders.

Customer	Treatment	Requested delivery date	Seed quantity (t)
SeedCompany	Split pill 3.5 let	20-08-19	505

Table 4-8. Example of a historical sales order.

Customer	Treatment	Requested delivery date	Seed quantity t	Process order
SeedCompany	Split pill 3.5 let	20-08-19	505	Coating 3.5 let
SeedCompany	Split pill 3.5 let	20-08-19	505	Drying 3.5 let

Table 4-9. Example of process orders for the sales order example.

4.3.2 Step 2: Determine process order version

In the second step the model determines which version of the process must be used, for each process order in step one (Table 4-9). Company A has resources with varying seed capacity ranges for each process. The simplest example is the coating process, for which Company A has coating pans with varying diameters. A larger coating pan is able to process larger seed quantities. To deal with this, Company A has defined separate versions of each process. Each version corresponds to a seed quantity range. For the coating process, there is a version with a small coating pan and a version with a large coating pan, each with predefined seed quantity ranges. These versions with seed quantity ranges are obtained from production data.

For the example in Table 4-9, Table 4-10 shows the output: the process order versions. Note that the coating process order is split in two process orders, because the maximum seed quantity is 300. This is not necessary for the drying process order, because the maximum seed quantity is sufficiently large. Company A always splits the seed quantity into equal batches, to ensure a uniform output.

Customer	Treatment	Requested delivery date	Seed quantity	Process order	Version
SeedCompany	Split pill 3.5 let	20-08-19	252.5	Coating 3.5 let	100-300
SeedCompany	Split pill 3.5 let	20-08-19	252.5	Coating 3.5 let	100-300
SeedCompany	Split pill 3.5 let	20-08-19	505	Drying 3.5 let	400-1000

Table 4-10. Example of process orders with version.

4.3.3 Step 3: From process orders to processing times

The third step is to generate a list of resources and processing times for each process order. The production task list from production data is input. The production task list contains the resources and processing times required for each process order and version. Starting with input from Table 4-10, Table 4-11 shows the output of this step.

Requested delivery date	Seed quantity	Process order	Version	Resource	Time h
20-08-19	252.5	Coating 3.5 let	100-300	Operator coa	2.5
20-08-19	252.5	Coating 3.5 let	100-300	Coating pan L	2.5
20-08-19	252.5	Coating 3.5 let	100-300	Operator coa	2.5
20-08-19	252.5	Coating 3.5 let	100-300	Coating pan L	2.5
20-08-19	505	Drying 3.5 let	400-1000	Operator dry	0.15
20-08-19	505	Drying 3.5 let	400-1000	Dryer L	6

Table 4-11. Example of production tasks for process orders and versions.

4.3.4 Step 4: Sum processing times

The fourth step is to sum the processing time of the production tasks by resource and month to obtain the historical capacity demand. Thereby, the demand is aggregated for all customers and products. Recall from Section 4.1 that we aggregate to forecast the capacity demand more accurately; there is insufficient data otherwise. Also recall from Section 4.1 that we aggregate per month to include seasonality and workforce planning. We use the requested delivery date to aggregate per month. By doing that, we assume that the processing takes place in the month of the requested delivery date. For example, if the requested delivery date is 02-08-19, the order is likely to be processed in July instead of August, while we assume the order is processed in August. We discuss this simplification in Chapter 5.

Starting with input from Table 4-11, Table 4-12 shows the resulting historical capacity demand. When starting with all sales orders as input instead, we obtain a meaningful historical capacity demand with seasonal patterns. The complete results are discussed in Chapter 5.

Requested delivery date	Resource	Time h
01-08-19	Operator coa	5
01-08-19	Coating pan L	5
01-08-19	Operator dry	0.15
01-08-19	Dryer L	6

Table 4-12. Example of historical capacity demand.

4.3.5 Step 5: Calculate adjusted processing times

Recall from Section 4.2 that to calculate adjusted capacity demand, we add a step to the calculation model. The input is adjustment factors, which is simply a multiplier for the processing time for specific products, customers, or a combination thereof. The output has the same format as Table 4-11, except with adjusted processing time, depending on the adjustment factors. For example, Table 4-13 shows the processing times if the adjustment factor is 2 for SeedCompany.

Requested delivery date	Seed quantity	Process order	Version	Resource	Time h
20-08-19	252.5	Coating 3.5 let	100-300	Operator coa	5
20-08-19	252.5	Coating 3.5 let	100-300	Coating pan L	5
20-08-19	252.5	Coating 3.5 let	100-300	Operator coa	5
20-08-19	252.5	Coating 3.5 let	100-300	Coating pan L	5
20-08-19	505	Drying 3.5 let	400-1000	Operator dry	.3
20-08-19	505	Drying 3.5 let	400-1000	Dryer L	12

Table 4-13. Example of production tasks for process orders and versions, after adjustment factor calculation.

After applying the adjustment factors, the same summation from Step 4 is applied to obtain the monthly capacity demand.

4.4 STRATEGIC CAPACITY PLANNING MODEL

In this section, we answer the following research question.

(Q4.3) How can the capacity planning model from literature be applied to Company A?

4.4.1 Strategic capacity planning with workforce planning and demand uncertainty

The purpose of our MIP model is to determine the capacity plan that satisfies demand at the lowest cost, while considering demand uncertainty. Recall from Section 4.1 that we use the models from Bihlmaier et al. (2009) and Fleischmann et al. (2006) as basis for our MIP model. Both models are unique in considering workforce planning in their strategic capacity plan. Recall from Section 4.1 that we anticipate tactical workforce planning in our strategic capacity plan, because the workforce planning decisions impact the machine capacity. To find an optimal strategic capacity plan, a trade-off be made between purchasing more machines or increasing workforce flexibility, for example. Our MIP model is a simplified version of Bihlmaier's multi-site model, because we deal with a single-site problem. We do not consider product allocation nor transportation problems.

Bihlmaier extends his strategic capacity planning model by integrating workforce planning through shift models. He defines a shift model as an amount of available capacity for which a number of employees is required. In our model, we break the shift model down into regular time, flexible time, overtime, and number of shifts. Each scales with the number of employees. Regular time is always available. Overtime and the number of shifts are both a binary monthly decision. One or two shifts. With or without overtime.

Bihlmaier extends his model by defining demand scenarios with a certain probability and finds the optimal capacity plan through stochastic optimization. Recall from Section 4.1 that we use scenarios based on coverage probability to generate capacity strategy alternatives.

4.4.2 Indices, parameters and decision variables

Let $i \in I$ be the machine types, which are operated by operator types $j \in J$. The planning horizon consists of months $t \in T$.

Symbol	Definition
$i \in I$	Set of machine types
$j \in J$	Set of operator types
$t \in T$	Set of months in planning horizon

Table 4-14. List of indices

The cost-related parameters are listed in Table 4-15. All other parameters are listed in Table 4-16. The capacity demand of machines and operators (c_{it}^{m}, c_{jt}^{o}) is the capacity demand scenario, which is the 50% scenario when generating capacity strategies. Recall from Section 4.1 that to generate capacity strategies, we add a constraint that the maximum machine capacity including workforce flexibility should be equal to a peak scenario (p_{it}). The reason is that the trade-off between machine capacity and workforce planning is made based on realistic demand, while having a higher maximum capacity to deal with extremes. The peak scenario represents these extremes, and has a coverage probability at least as high as the realistic scenario (50%).

Symbol	Definition	Unit
cm_i	Maintenance cost of machine i	Euros per month
cp_i	Procurement cost of machine i	Euros
cw_j	Wage in regular shift of operator j	Euros per month
ch_j	Hiring cost of operator j	Euros
cl_j	Leaving cost of operator j	Euros
ct_j	Training cost of (flex-)operator j	Euros
cd_j	Deployment cost of flex-operator j	Euros per month
co_j	Overtime cost, as multiplier for the wage in regular shift	% of wage
cs_j	Double shift cost, as multiplier for the wage in regular shift	% of wage

Table 4-15. List of cost parameters.

Symbol	Definition	Unit
c_{it}^{m}	Capacity demand for machine i in month t	Hours
c_{jt}^{o}	Capacity demand for operator j in month t	Hours
p_{it}	Peak capacity demand for machine i in month t	Hours
d_t^{w}	Number of workdays in month t	Days per month
d_t^{s}	Number of Saturdays in month t	Days per month
e_i^{m}	Overall equipment effectiveness of machine i	%
e_j^{o}	Overall equipment effectiveness of operator j	%
h_i^{m}	Daily number of hours available for machine i	Hours per day

h_j^o	Daily number of hours available for operator j	Hours per day
a_j	Fraction of operators available of type j (e.g. not sick)	%
xy_{ij}	1 if operator j is required for machine i to operate, 0 if not	Binary
m	Arbitrary large number	

Table 4-16. List of miscellaneous parameters.

The strategic capacity decisions concern machines. Capacity is increased by purchasing new machines (P_{it}), which results in a number of available machines (X_{it}). Capacity replacement is not included explicitly. However, capacity replacement can be modeled by setting the number of machines at time 0 ($X_{j,t=0}$) at 0. Reducing the number of machines is only interesting if maintenance cost is high, or machines can be sold, which is not the case for Company A.

The tactical capacity decisions concern operators. Capacity is increased through hiring (H_{jt}) regular operators (Y_{jt}) and training (T_{jt}) flex operators (F_{jt}). Flex operators can be deployed for a fraction of a month (D_{jt}), which is why this is the only decisions variable that is a positive real number. Operators can leave (L_{jt}) for various reasons, such as pension, moving to another department or leaving to work for a competitor.

The strategic and tactical capacity decisions are integrated using overtime (O_{jt}) and shifts (S_{jt}). When an operator type works in double shifts, the capacity of the machines operated by this operator type is effectively doubled, at the expense of an increase in wages by the factor cs_j. The reason is that by using double shifts (which are non-overlapping), the machines can be used for 15 hours instead of 7.5 hours each day. When an operator type works overtime, the capacity of machines operated by this operator type increases by 12% on average. The reason is that in overtime, each operator works for an additional 5 hours on Saturdays, which results in an average time increase of 12% per month. This usually only happens in case of emergencies, which is on average one month each year, due to peak demand.

Symbol	Definition	Unit
X_{it}	Number of machines of type i available in month t	Integer
P_{it}	Number of machines of type i purchased in month t	Integer
Y_{jt}	Number of operators of type j available in month t	FTE
H_{jt}	Number of operators of type j hired in month t	FTE
L_{jt}	Number of operators of type j leaving in month t	FTE
F_{jt}	Number of flex-operators of type j available in month t	FTE
D_{jt}	Number of flex-operators of type j deployed in month t	FTE
T_{jt}	Number of flex-operators of type j trained in month t	FTE
O_{jt}	0 for no overtime; 1 for overtime for operator type j in month t	Binary
S_{jt}	0 for regular shift and 1 for double shift for operator j in month t	Binary
Z_{jt}	Wage factor. Y_{jt} if regular shift; $(1 + cs_j) * Y_{jt}$ if double shift; $(1 + co_j) * Y_{jt}$ if overtime; $(1 + cs_j + co_j) * Y_{jt}$ if double shift and overtime	Double

Table 4-17. List of decision variables.

4.4.3 Objective and constraints

The MILP model that determines the optimal capacity plan is shown below. The objective minimizes the costs associated with the capacity plan, which are: machine maintenance cost, cost of purchasing machines, shift- and overtime-dependent wages, deployment of flex-operators, cost of hiring and training new operators, and the training of flex-operators.

$$min \sum_i \sum_t cm_i \cdot X_{it} + \sum_i \sum_t cp_i \cdot P_{it} + \sum_j \sum_t cw_j \cdot Z_{jt} + \sum_j \sum_t cd_j \cdot D_{jt} + \sum_j \sum_t (ch_j + ct_j) \cdot H_{jt} + \sum_j \sum_t cl_j \cdot L_{jt} + \sum_j \sum_t ct_j \cdot T_{jt}$$

subject to

1) $X_{it} \cdot d_t^w \cdot h_i^m \cdot e_i^m \geq c_{it}^m - \sum_j m \cdot (S_{jt} + O_{jt}) \cdot xy_{ij} \quad \forall\, i, t$
2) $2 \cdot X_{it} \cdot d_t^w \cdot h_i^m \cdot e_i^m \geq c_{it}^m - \sum_j m \cdot O_{jt} \cdot xy_{ij} \quad \forall\, i, t$
3) $X_{it} \cdot d_t^w \cdot h_i^m \cdot e_i^m + 0.12 \cdot X_{it} \cdot d_t^s \cdot h_i^m \cdot e_i^m \geq c_{it}^m - \sum_j m \cdot S_{jt} \cdot xy_{ij} \quad \forall\, i, t$
4) $2 \cdot X_{it} \cdot d_t^w \cdot h_i^m \cdot e_i^m + 0.12 \cdot X_{it} \cdot d_t^s \cdot h_i^m \cdot e_i^m \geq p_{it} \quad \forall\, i, t$
5) $(Y_{jt} + D_{jt}) \cdot d_t^w \cdot h_j^y \cdot e_j^y \geq c_{jt}^y - m \cdot O_{jt} \quad \forall\, j, t$
6) $(Y_{jt} + D_{jt}) \cdot d_t^w \cdot h_j^y \cdot e_j^y \cdot a_j + 0.12 \cdot Y_{jt} \cdot d_t^s \cdot h_j^y \cdot e_j^y \cdot a_j \geq c_{jt}^y \quad \forall\, j, t$
7) $Z_{jt} \geq Y_{jt} \quad \forall\, j, t$
8) $Z_{jt} \geq (1 + cs_j) \cdot Y_{jt} - m \cdot (1 - S_{jt}) \quad \forall\, j, t$
9) $Z_{jt} \geq (1 + co_j) \cdot Y_{jt} - m \cdot (1 - O_{jt}) \quad \forall\, j, t$
10) $Z_{jt} \geq (1 + cs_j + co_j) \cdot Y_{jt} - m \cdot (2 - S_{jt} - O_{jt}) \quad \forall\, j, t$
11) $X_{it} + X_{i,t-1} - P_{i,t} = 0 \quad \forall\, i, t$
12) $Y_{jt} + Y_{j,t-1} - H_{jt} + L_{jt} = 0 \quad \forall\, j, t$
13) $F_{jt} + F_{j,t-1} - T_{jt} = 0 \quad \forall\, j, t$
14) $D_{jt} \leq F_{jt} \;\forall\, j, t$
15) $X_{it}, P_{it}, Y_{jt}, H_{jt}, L_{jt}, F_{jt}, T_{jt}, Z_{jt} \in \mathbb{N}^0$
16) $D_{jt} \in \mathbb{R}_{\geq 0}$
17) $O_{jt}, S_{jt} \in \{0, 1\}$

Constraints 1, 2, 3 and 4 together ensure available machine capacity is sufficient to fulfill the required machine capacity. To obtain the available capacity, the number of machines is multiplied by the workdays in a month, the hours available per day, and the machine OEE. Constraint 1 reflects the machine capacity when the operator who operates the machine, denoted by xy_{ij}, has regular shifts. It is unconstrained when this operator has double shifts or uses overtime. Constraint 2 reflects the machine capacity when the operator has double shifts, which is double the normal capacity. Constraint 3 reflects the machine capacity when the operator works overtime, which is 0.12 times the daily capacity, multiplied by the number of Saturdays. Constraint 4 reflects the machine capacity when the operators work both double shifts and overtime, which is the maximum capacity. For constraint 4, the peak capacity scenario is used instead, because the maximum capacity must be able to deal with the peaks in demand.

Constraint 5 and 6 ensure available operator capacity is sufficient to fulfill the required operator capacity. The number of regular and flexible operators is multiplied by the workdays in a month, the hours available per day, operator OEE, and the fraction of operators who are available (e.g. not sick or on leave). Constraint 5 reflects the operator capacity in case of no overtime, which is unconstrained when overtime is used. Constraint 6 reflects the operator capacity when they work overtime.

Constraints 7, 8, 9, 10 reflect the wages for regular shifts, double shifts, overtime, or both double shifts and overtime, respectively. In case of regular shifts, $Z_{jt}$is equal to the number of operators (Y_{jt}), because Z_{jt} is minimized in the objective function. In case of double shift, Z_{jt} is equal to the number of operators, multiplied by the cost factor for double shifts ($1 + cs_j$). In case of overtime, Z_{jt} is equal to the number of operators, multiplied by the cost factor for overtime

$(1 + co_j)$. In case of overtime and double shifts, Z_{jt} is equal to the number of operators, multiplied by the cost factor for double shifts and overtime $(1 + cs_j + co_j)$.

Constraint 11 is a balance equation for the number of machines, considering machine procurement. Constraint 12 is a balance equation for the regular operators, considering hiring and leaving. Constraint 13 is a balance equation for flex-operators, considering training.

Constraint 14 limits the deployed flex-operators to the number of flex-operators available. When flex-operators are not deployed, they work at their regular function within the company.

Constraint 15, 16, and 17 define the decision variables, either non-negative integers, non-negative real numbers, or binary.

Note that we do not consider lead-time in our model. Company A expressed their preference for not considering lead-times, because it simplifies the interpretation of the model. Decision makers can simply calculate back in time to decide when they need to start the hiring or purchasing process.

The capacity optimization model has been implemented in Excel, using the OpenSolver plugin with the CBC solver engine. This plugin and solver engine is part of the academic COIN-OR open source library for operations research.

4.5 DISAGGREGATING CAPACITY DEMAND

We assume that Company A is able to smooth capacity demand in each month, such that a capacity plan with monthly periods is sufficiently detailed. While deviations are expected on a weekly level, due to customer requested due dates, these deviations must not be too large. If this assumption does not hold, the capacity plan can be feasible on a monthly level, while being infeasible on a weekly level. In that case, the capacity plan is not very useful in practice.

To validate the assumption, we disaggregate the demand forecast from months to weeks, while simulating the variance as observed in historical data. By comparing the disaggregated demand forecast to the available capacity, the result is the number of weeks with sufficient capacity in a year. We choose a year because it covers an entire season. Our method consists of the following steps. Figure 1-1 visualizes this method.

1. The first step is to deseasonalize the historical capacity demand, using the seasonal factors as calculated by the forecasting model. Variance between weeks within each month is of interest, not variance caused by seasonality.
2. The second step is to calculate the standard deviation of weekly capacity demand. This is calculated from deseasonalized historical capacity demand that is grouped by week.
3. The third step is to split the monthly point forecasts evenly into weeks, such that we obtain a weekly forecasts.
4. The fourth step is to calculate the weekly available capacity, based on the capacity decisions and parameters from our capacity planning model.
5. The fifth step is to generate a random number from the standard normal distribution.
6. This random number is multiplied by the standard deviation from Step 2 to obtain a deviation for each week. Note that by doing this we assume normally distributed errors.
7. The deviation is added to the weekly forecasts from Step 3 to obtain a weekly forecast with a random component, based on variance in historical data.
8. The final step is to calculate the fraction of weeks in which there is sufficient capacity, by comparing the available capacity from Step 4 to the weekly forecast from Step 7.

Because there is a random component, we decide to do multiple replications of Steps 5 through 8. The number of replications is discussed in Chapter 5. To compare the disaggregated (weekly) to the aggregated (monthly) capacity demand, we follow the same method for months instead of weeks. The results are discussed in Chapter 5.

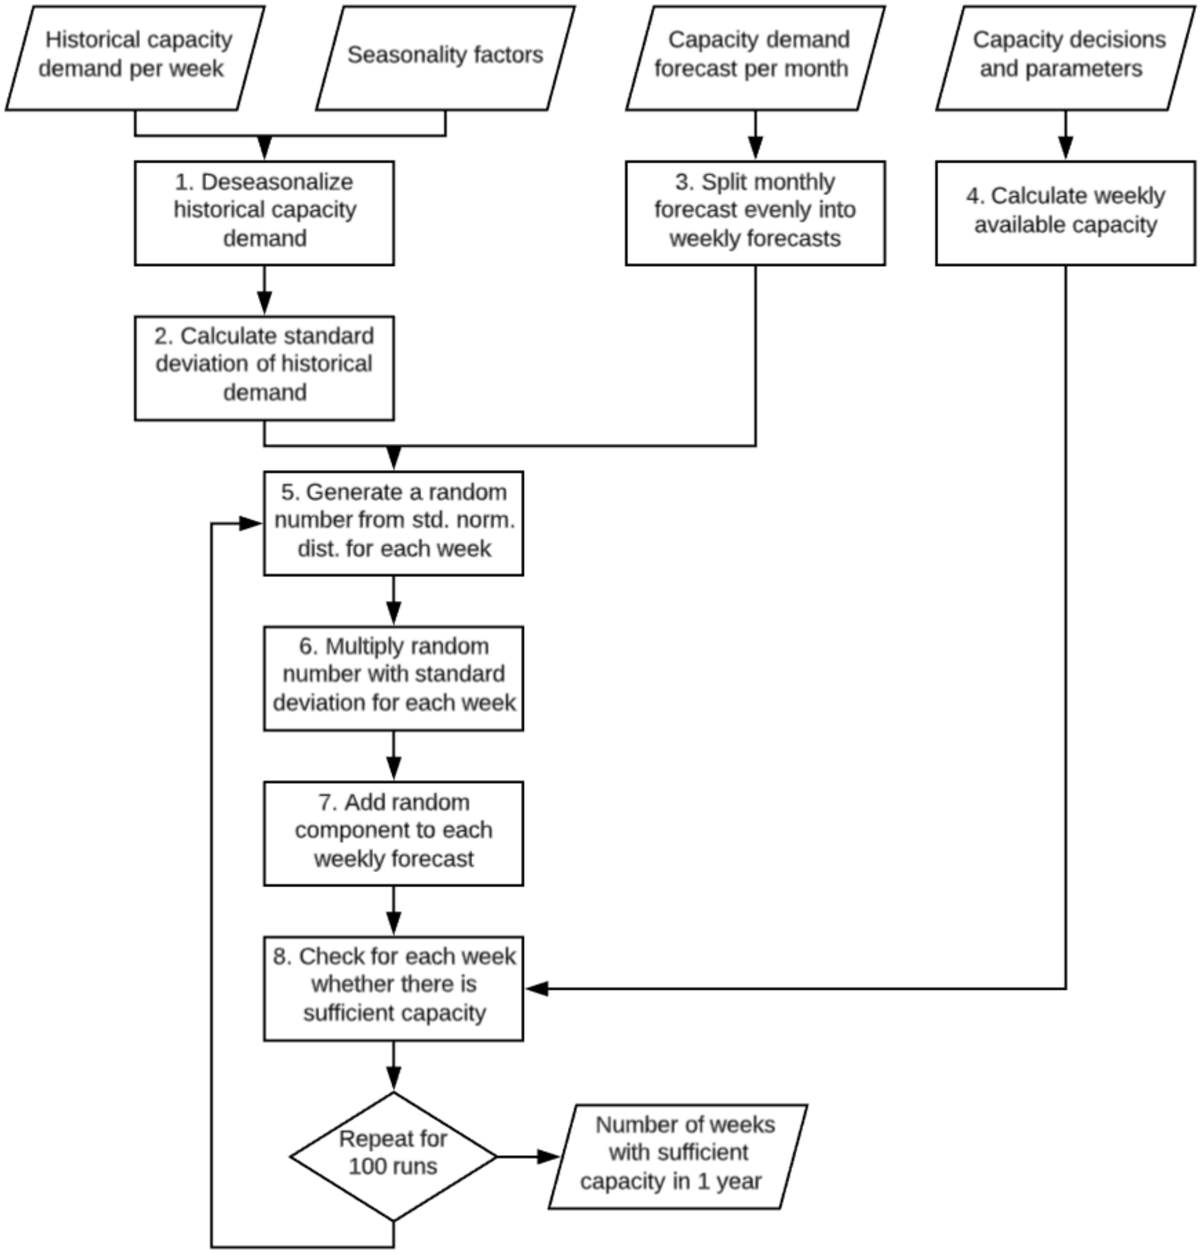

Figure 4-3. Method to disaggregate capacity demand forecast.

4.6 CONCLUSION

In this Chapter, we designed a capacity planning model that addresses the three core problems identified in Chapter 1. We calculate capacity demand more accurately, by using historical sales orders and production data. (Q4.1) We generate more accurate capacity demand forecasts, by starting with historical capacity demand. We use ETS to model seasonality and uncertainty, through prediction intervals. (Q4.2) Continuing with these capacity demand forecast scenarios, we determine a set of optimal capacity plans, considering both strategic decisions and workforce planning. These capacity plans enable company management to make the trade-off between uncertainty and capacity costs, thereby dealing with demand uncertainty. (Q4.3)

The model contributes to theory in three ways. First, for companies where it is more complex to calculate capacity requirements from product demand, our model provides a much more accurate solution. A highly custom calculation can be created to calculate capacity demand, instead of using product demand in the capacity optimization model. Second, we have designed two ways to use judgmental forecasts with statistical forecasts. It results in more accurate forecasts in cases where contextual information, such as legislation and new products, heavily impacts future demand. Finally, we have designed a more intuitive way to include demand uncertainty. We use prediction intervals to determine how much capacity is needed to be able to fulfill demand with a specific certainty.

5 MODEL RESULTS

We formulated the main research question in Chapter 1 as follows: how can machine- and operator capacity planning deal with seasonal and variable demand to improve on-time delivery in a cost-efficient way? We answer that question in this chapter. Based on the current situation analysis (Chapter 2), and the models from literature (Chapter 3), we designed the model in Chapter 4. Recall from Section 4.1 that the model consists of the following steps.

1. Calculate historical capacity demand from sales orders
2. Forecast future capacity demand scenarios from historical capacity demand
3. Create judgmental sales forecasts
4. Calculate adjusted capacity demand from judgmental sales forecasts
5. Forecast adjusted capacity demand scenarios
6. Determine optimal capacity strategy for each scenario using a MIP model
7. Evaluate each capacity strategy for various scenarios

Section 5.1 discusses the results and validation of the first two steps; the calculation and forecasting of capacity demand. Section 5.2 discusses Step 6; the optimal capacity strategies for each scenario. Section 5.3 discusses the impact of judgmental forecasts through adjusted capacity demand scenarios, the results of Step 3 through 5. Section 5.4 discusses Step 7; evaluating each capacity strategy. Finally, in Section 5.5, a sensitivity analysis of the model is performed.

5.1 MODEL VALIDATION

The optimization model described in Chapter 4 determines the optimal capacity strategy based on a set of parameters and decisions. This capacity strategy is only as good as the formulation of the model and the accuracy of the input parameters. The model formulation is described in Chapter 4. The most important input parameters for a capacity plan are the capacity demand and available capacity. We assess the accuracy of capacity demand in this section by answering the following research questions. Section 5.5 assesses the accuracy of the available capacity parameters, by performing a sensitivity analysis of the parameters subject to uncertainty and change.

(Q5.1) How accurate can our model calculate capacity demand?
(Q5.2) How accurate can our model forecast capacity demand?

5.1.1 Historical capacity demand calculation

The first step of our model is to calculate historical capacity demand from sales orders, using production data. To validate the accuracy of this calculation, we compare the calculated capacity demand with the actuals. The realized processing times are stored in SAP for each machine type, starting February 2018. We sum the processing times per month and machine to compare them with the calculated capacity demand. We choose to focus this section on dryer C0414, coating pan P100, and coating operators, because these are the top 3 when it comes to the amount of costs and capacity demand.

Figure 5-1, Figure 5-2, and Figure 5-3 show the calculated and actual capacity demand for coating operators, coating pan P100, and dryer C0414, respectively. We expect that the actuals are smoother than the calculated values, because we sum capacity demand by requested delivery date, while Company A is able to smooth demand between weeks. However, we do not expect a large effect, as Company A is not able to smooth over longer periods, such as several months, because demand is tied to a timing window of about two to three weeks. We observe this smoothing effect in Figure 5-1 between November 2018 and February 2019. In Figure 5-2, this effect is visible around May 2019. Finally, in Figure 5-3 we see smoothing in November and December 2018.

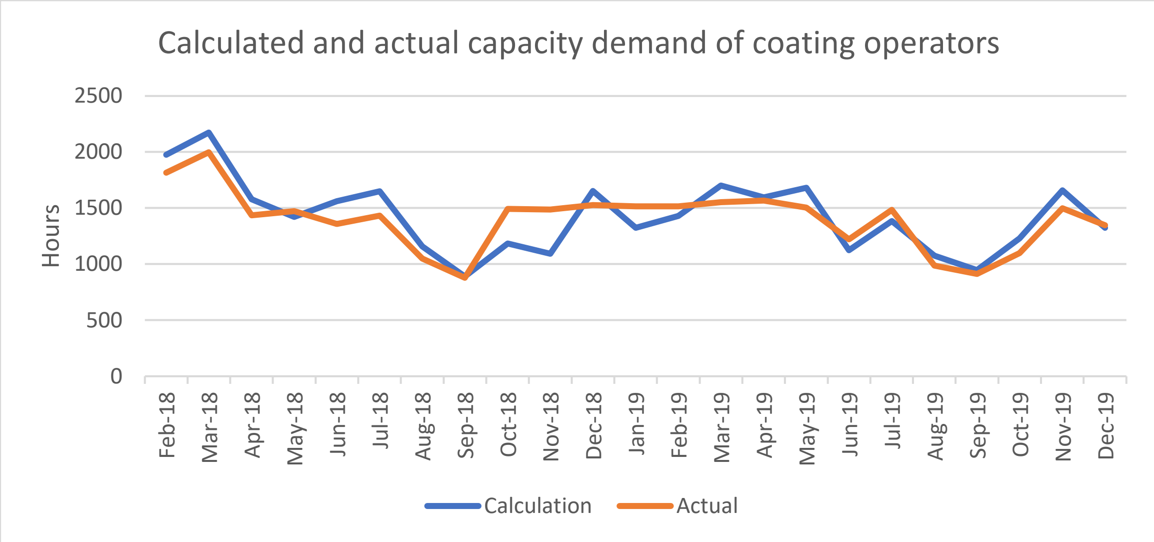

Figure 5-1. Calculated vs actual capacity requirements of coating operators.

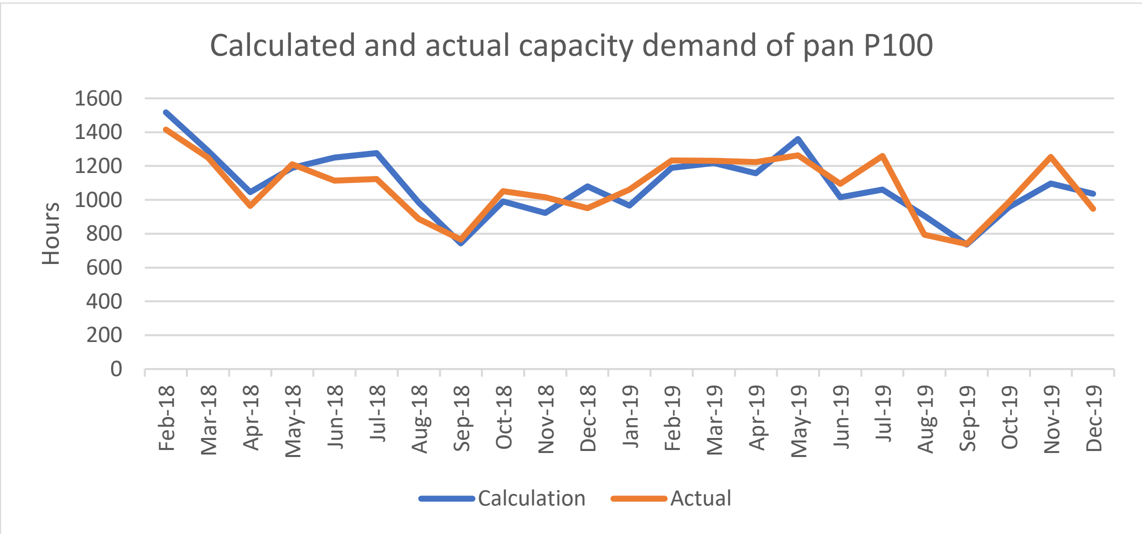

Figure 5-2. Calculated vs actual capacity requirements of coating pan P100.

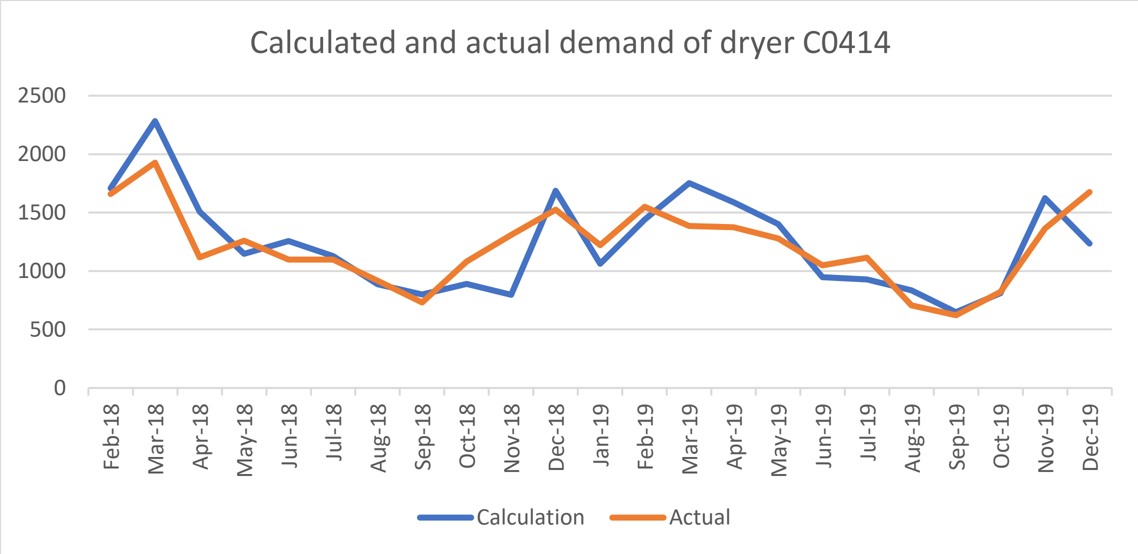

Figure 5-3. Calculated vs actual capacity requirements of dryer C0414.

To evaluate the accuracy of our calculated capacity demand, we use performance measures that are usually used for forecasting. To do this, we define the error as the actual values minus the calculated values. We use three performance measures from Chapter 3, the literature review. We use the Mean Absolute Error to measure the mean error in hours, which helps to understand the accuracy from a practical point of view. We use the symmetric Mean Absolute Percentage Error to compare the accuracy between resources. And lastly, we use the Mean Percentage Error to assess the bias. We do not use the Mean Absolute Scaled error, because it makes no sense to scale the calculation error to the naïve forecast error.

Table 5-1 shows the performance measures for coating operators, coating pans P100 and dryers C0414. The MPE indicates a small bias; the calculation overestimates the capacity demand by 1% to 2.2% on average. We expect that the reason lies with scheduling. In SAP, there are hard limits on maximum seed quantities, while in practice these are not as strict. Therefore, the planners split sales orders in fewer process orders, by violating the maximum seed quantity limit. The bias will probably not have a significant impact on the capacity strategies, because 2% of the 14 dryers currently available is only 0.28 machine. Company A has less machines of each other type.

MAE is highest for dryer C0414, with 182 hours. The available monthly capacity for one dryer is about 120 machine hours, so the error is about one and a machine month. The sMAPE for coating opeators and pan P100 is lower than dryer C0414. The errors are not the same for the same month between years, which is consistent with the small bias. For example, while the demand for dryer C0414 is calculated too high in November 2018, it is calculated too low in November 2018. Therefore, because we have data starting in 2015, we expect that the error will be averaged between the years, resulting in a lower forecast error. This is discussed in Section 5.1.2.

Performance measure	Coating operator	Pan P100	Dryer C0414
Mean Percentage Error (MPE)	-0.022	-0.010	-0.019
Mean Absolute Error (MAE) in hours	137	83	182
Symmetric Mean Absolute Percentage Error (sMAPE)	0.048	0.038	0.072

Table 5-1. Performance measures for calculated vs actual capacity requirements.

The current calculation method measures capacity demand in number of orders, which is used to determine the capacity plan relative to previous year. If there are more orders, more operators and workforce flexibility is required, and vice versa. The advantage of our calculation method is that we measure capacity demand in hours, which can be used to determine a capacity plan directly and more accurately, instead of relative to previous year. Figure 5-4 shows the number of orders as compared to the calculated and actual capacity demand for dryer C0414, with the number of orders on the secondary axis. We observe that the number of orders is a much more inaccurate method to determine capacity demand than our calculation model that uses sales orders and prdouction data. In addition to improved accuracy, the advantages of our calculation method over the current method are as follows.

1. We measure capacity demand in hours instead of number of orders. Hours can directly be translated to number of machines and operators.
2. By calculating from sales orders, we have data starting from February 2015, instead of February 2018, resulting in more accurate forecasts.
3. We can calculate capacity demand for judgmental forecasts. For example, by forecasting the number and seed quantity of future orders for a new product or a new machine type.

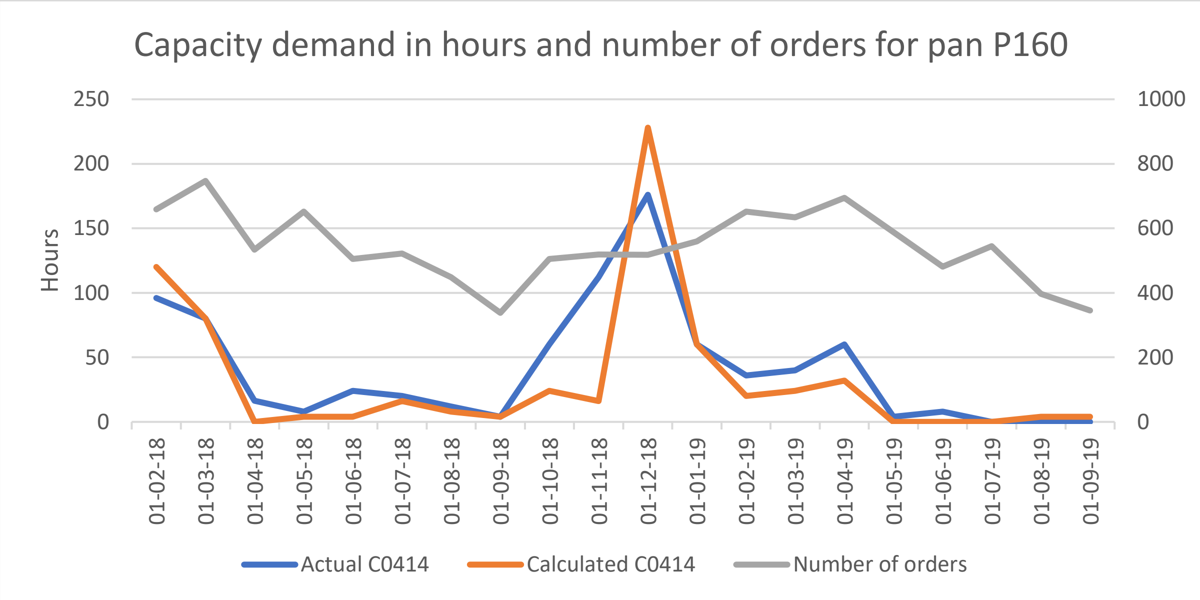

Figure 5-4. Number of orders compared to the calculated and actual capacity demand for pan P160.

5.1.2 Capacity demand forecast

The second step in our model is to forecast capacity demand, based on historical capacity demand, using the ETS forecasting model. To validate the forecasts, we compare the point forecast to the actuals. We need at least four seasonal cycles to generate accurate seasonal forecasts. Since we have data available since 2015, we only have 2019 to evaluate forecasts. As more data comes available, Company A should update the performance evaluation of the forecasts.

Figure 5-5, Figure 5-6 and Figure 5-7 show the forecasted and actual capacity requirements for coating operators, coating pan P100, and dryer C0414, respectively. We observe the largest over-forecast in March 2019, which is not a surprise, as demand was lower than expected that month. The number of orders was about 15% lower in March 2019 than in March 2018. The reason is that in 2018 there were severe quality issues with one crop, effectively halving demand for that product in 2019. P100 is an exception; it does not show an over-forecast in March 2019. This specific crop is processed on a different coating pan type, so that is in line with expectations.

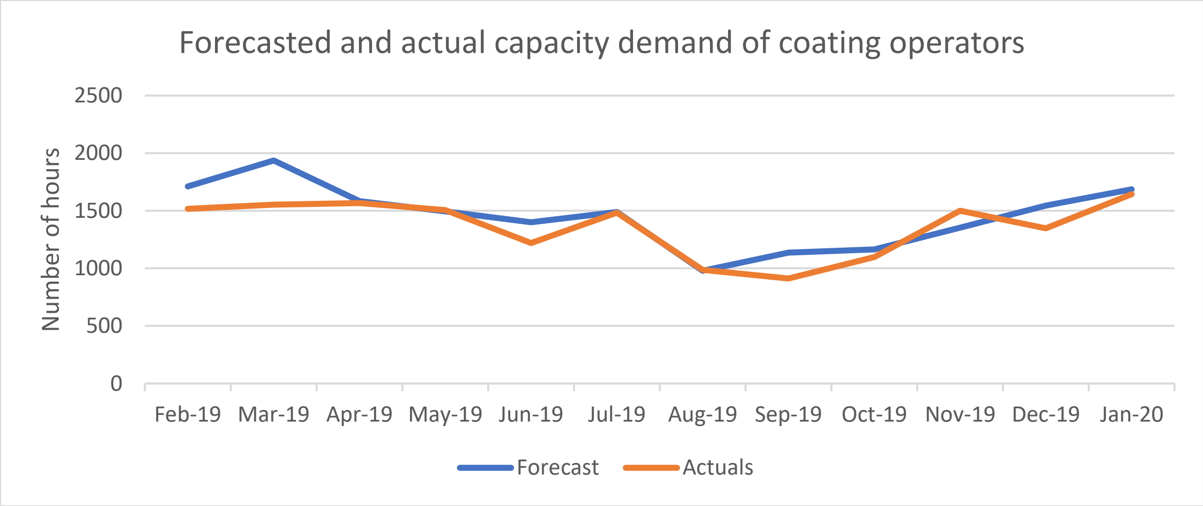

Figure 5-5. Point forecast and actual capacity demand for coating operators.

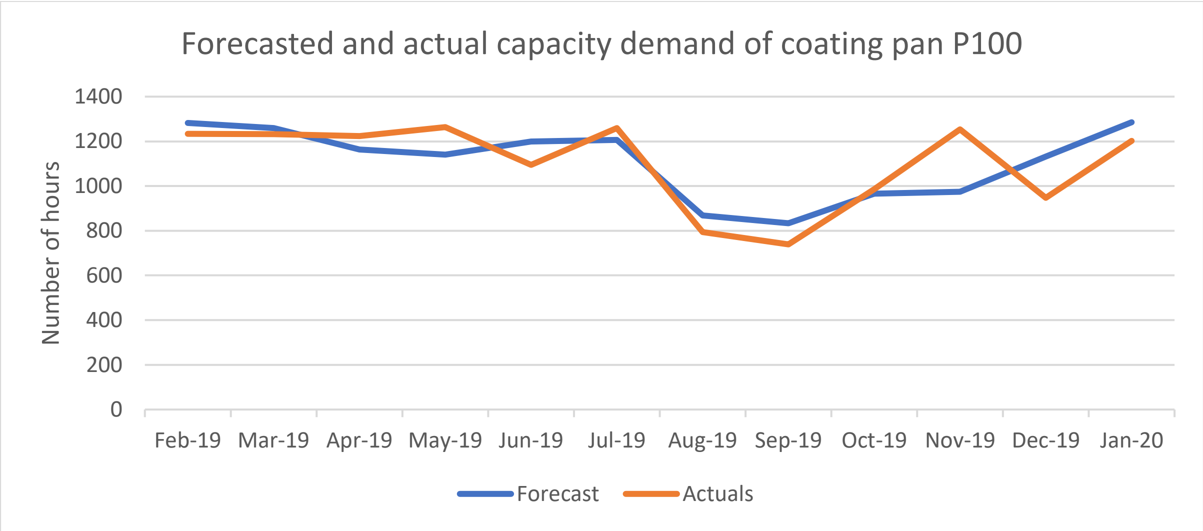

Figure 5-6. Point forecast and actual capacity demand for pan P100.

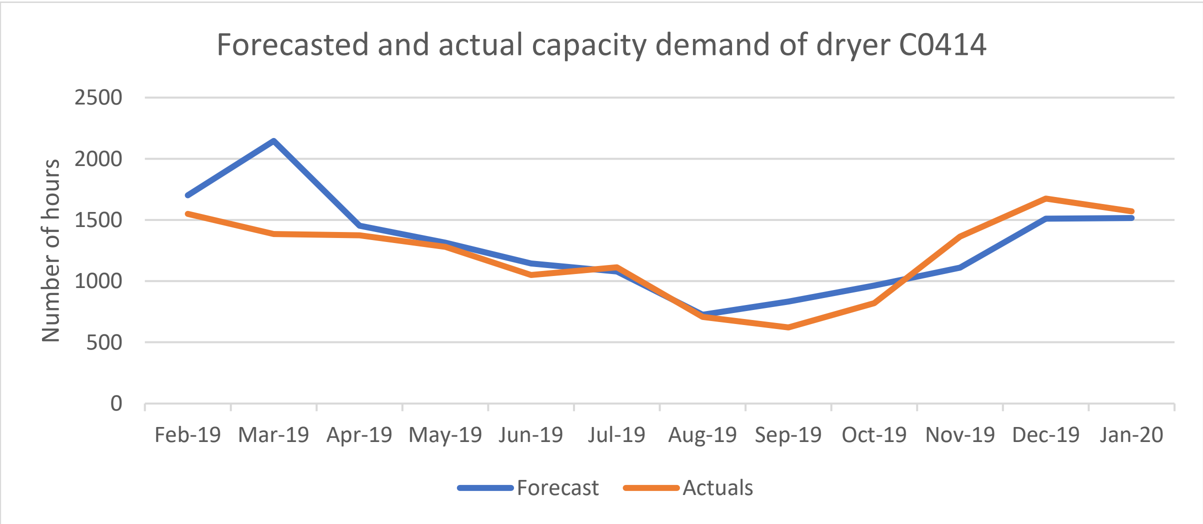

Figure 5-7. Point forecast and actual capacity demand for dryer C0414.

To calculate the performance measures, we define forecast error as the actual values minus the forecasted capacity demand. Table 5-2 shows the performance measures. The Mean Percentage Error shows a negative bias, which is over-forecasting. The bias is much higher for coating operators and dryer C0414, which is explained by the low demand in March 2019 that we previously discussed. Taking this outlier into account, the mean percentage error is about -5% for coating operators and -4% for dryers. This is less than half of one operator and dryer in terms of capacity, which is an acceptable bias for Company A. This bias should be re-evaluated as more data becomes available.

The mean absolute error of coating operators is almost equal to one operator, which has a capacity of about 130 hours per month. The mean absolute error for pan P100 is about two-third of one machine (i.e. 150 hours) and for dryer C0414 it equals nearly one and a half machine (i.e. 120 hours). Even though the mean absolute error is inflated because of low demand in March 2019, the error is still substantial. When looking at the symmetric mean percentage error, we see that the forecast accuracy is similar for coating operators and pan P100, while dryer C0414 is worse by two percent.

The mean absolute scaled error is below 1 for each resource, which means that the forecasting method performs better than the seasonal naïve forecast. The forecast for coating operators shows the most improvement, it is twice as accurate as the seasonal naïve forecast.

Performance measure	**Coating operator**	**Pan P100**	**Dryer C0414**
Mean Percentage Error (MPE)	-0.075	-0.018	-0.084
Mean Absolute Error (MAE)	123	96	166
Symmetric Mean Percentage Error (sMAPE)	0.044	0.045	0.064
Mean Absolute Scaled Error (MASE)	0.49	0.69	0.65

Table 5-2. Forecast performance measures for ETS forecasts of capacity requirements.

We conclude that the forecast is a large improvement over the current seasonal naïve forecast. However, we must remain aware that the mean absolute error equals around one operator or machine when interpreting the capacity strategies in the next sections.

The result of the second step is not just a point forecast, but a set of capacity demand scenarios. Recall from Chapter 4 that a capacity demand scenario is the one-tailed upper prediction interval for a coverage probability. We name each scenario by this coverage probability. We define four scenarios: 30, 50, 70, and 90. Note that scenario 50 is equal to the point forecast, because we assume a standard normal distribution. Figure 5-8, Figure 5-9, and Figure 5-10 show the four scenarios for coating operators, pan P100 and dryer C0414, respectively.

In addition to scenarios based on historical capacity demand only, we generate adjusted scenarios based on both historical data and judgmental forecasts in the third and fourth step of our model. We refer to these adjusted scenarios as 30-A, 50-A, 70-A, and 90-A. We discuss the adjusted scenarios in Section 5.3.

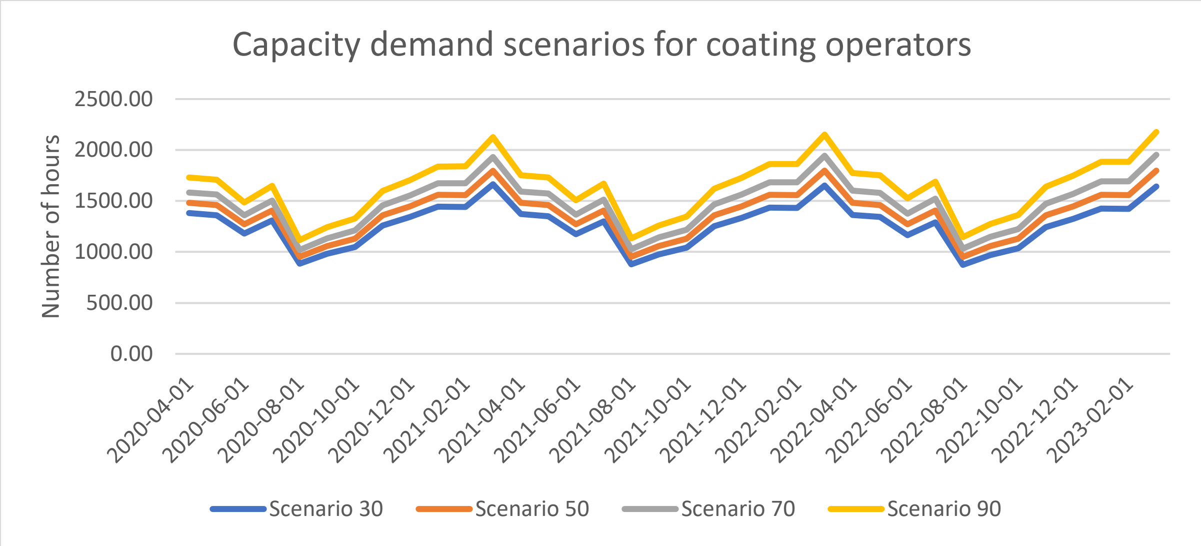

Figure 5-8. Capacity demand scenarios for coating operators.

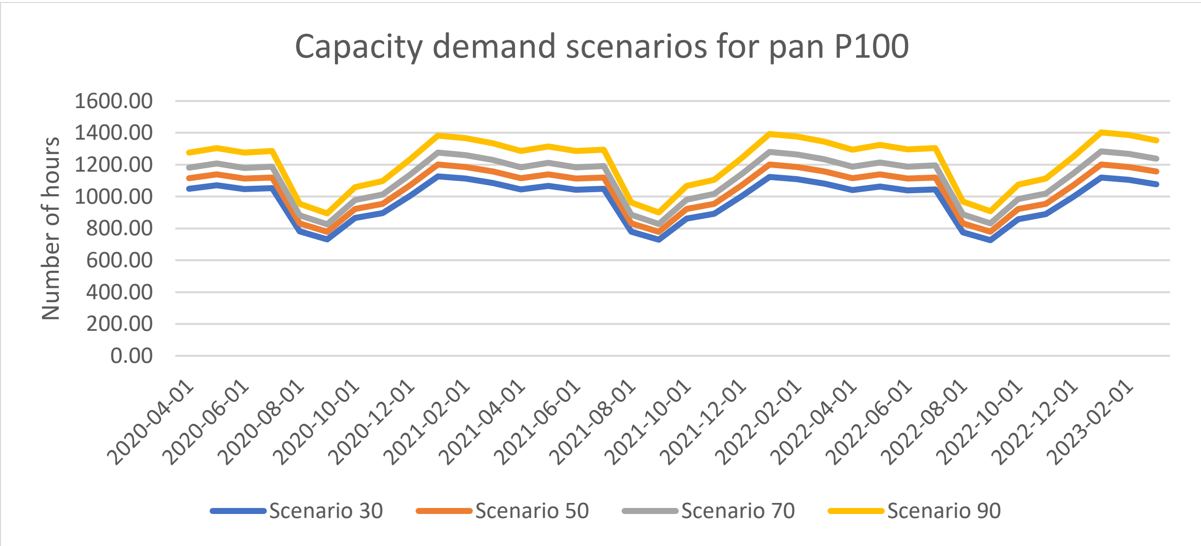

Figure 5-9. Capacity demand scenarios for pan P100.

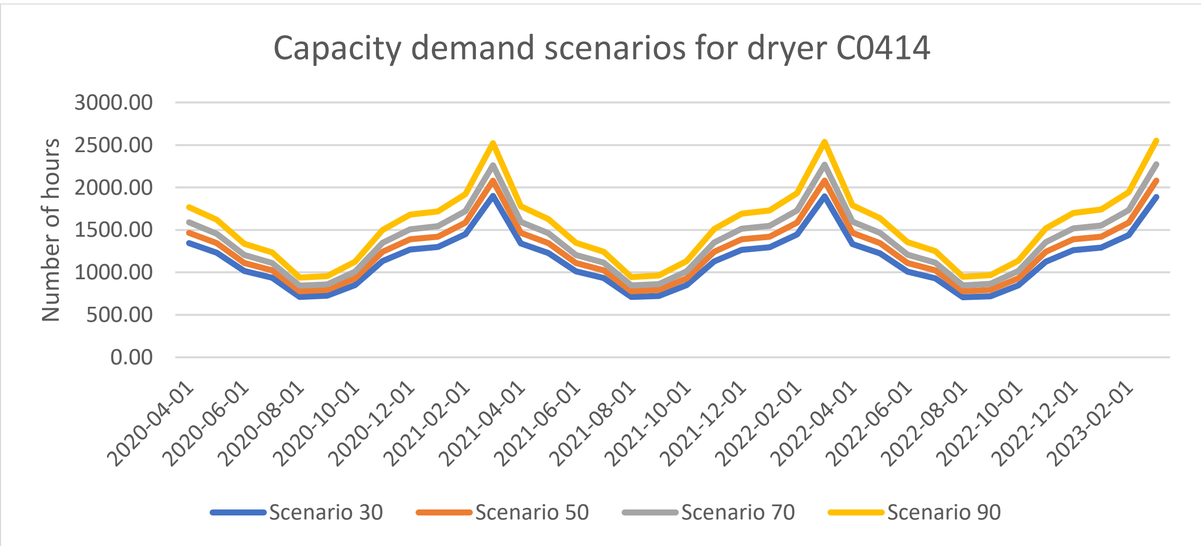

Figure 5-10. Capacity demand scenarios for dryer C0414.

5.2 CAPACITY STRATEGIES FOR VARIOUS SCENARIOS

We discussed the capacity demand scenarios in Section 5.1. In this section, we use these scenarios to generate capacity strategies and answer the following research question.

(Q5.3) How does the capacity strategy from our model compare to the current capacity plan?

5.2.1 Generating the capacity strategies

We discussed the results of the first two model steps, the capacity demand scenarios, in Section 5.1. Section 5.3 discusses the results of Step 3, 4, and 5. Step 6 is to determine the optimal capacity strategy for various scenarios. We define the capacity strategy as a set of machine procurement decisions, because these are the strategic capacity decisions for Company A. Recall from Section 4.5 that we obtain a capacity strategy by optimizing the capacity plan for two scenarios simultaneously: the realistic scenario (50%) and a peak scenario. We optimize the capacity plan for the realistic scenario, while we increase the capacity demand to the peak scenario for the maximum machine capacity constraint. That way, the number of machines is sufficient to deal with capacity demand peaks through workforce flexibility measures only.

We use three peak scenarios, with coverage probabilities of 50, 70, and 90. Company A is not interested in a strategy with a confidence lower than 50%, because it is their business strategy to be a reliable and flexible partner for their customers. In the remainder of this thesis, we refer to each capacity strategy by the coverage probability of the peak scenario (e.g. Strategy 70 is generated with Scenario 70 as peak scenario). To summarize, the process for generating the capacity strategies is visualized in Figure 5-11. Note that the order in generating capacity plans is not relevant. The input is a set of capacity scenarios. The output is a set of capacity strategies, which we obtain by running the optimization model.

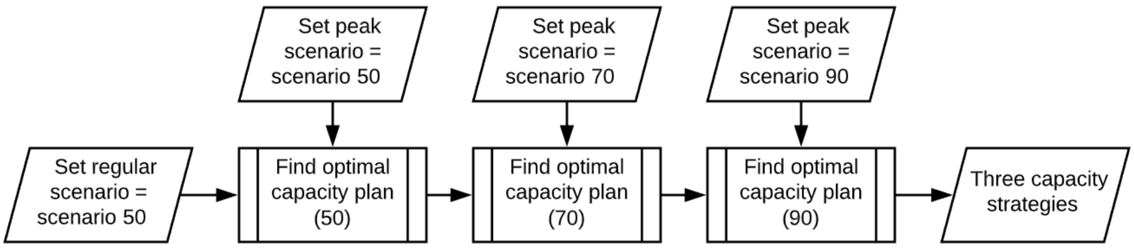

Figure 5-11. Generate capacity strategies for three peak scenarios.

Recall from Chapter 2 that Company A's interest is the replacement of dryers and coating pans. Dryers will be replaced in the coming 3 years and coating pans in the coming 6 years. To make the trade-off between workforce flexibility and machine procurement, replacement of both dryers and pans must be considered simultaneously. Recall from Section 4.5 that we model replacement by setting the initial number of machines at 0 for the machine types that need to be replaced. We obtain three capacity strategies, the machine procurement decisions, by following the process in Figure 5-11. To compare these strategies with the current situation and answer the research sub-question, we define the 'current' strategy, where we purchase the same number of machines as currently in use.

5.2.2 Comparing capacity levels

Figure 5-12 shows the capacity strategies that we obtain from running our optimization model. To compare the strategies, we show the number of machines purchased for each machine type. We observe that the current number of machines is equal to or higher than the number of machines purchased in Strategy 90. Therefore, the current strategy already provides sufficient capacity with at least 90% certainty. Recall from the problem cluster in Chapter 1 that we hypothesized that Company A experiences decreased delivery performance during peak season due to a lack of capacity. While that may still be the case on an operational (i.e. weekly or daily) level, we have shown that this is not the case on a tactical (i.e. monthly) level. More specifically, even though a 90% confidence still leaves 10% chance of undercapacity (in case the available capacity is equal to the capacity demand scenario), undercapacity is not a structural problem on a tactical level.

Takeaway 1
With current capacity levels, capacity shortage is not a structural problem on a monthly level.

The hypothesis that decreased delivery performance is caused by capacity shortage was based on interviews with Company A's management as well as the production team. Unfortunately, the data was not sufficient to verify this claim. Moreover, we cannot simply discard their experiences and opinions as invalid. Instead, we think that capacity shortages are on an operational level. Company A uses the backward scheduling method, based on customer requested date. An order is scheduled within two days of receiving it, no matter how far the requested date lies in the future. Schedules are barely changed, except in case of an emergency by a simple swapping of orders. The reason is that moving the schedule is a time-intensive task. Note that for some products, the time between processing steps must be at most a few minutes. The result is that there are no sequential timeslots available between processing steps, while there are free timeslots for both processes at different times. Such issues occur frequently, but cannot be addressed, because the schedule can barely be changed. In the meanwhile, there is sufficient capacity on an aggregated level. We suggest a future research direction for Company A that addresses this issue in Chapter 6, on the subject of scheduling.

From Figure 5-12, we also observe that there is currently an overcapacity, even at the 90% coverage probability, of four coating pans: P055, P060, P100, and P160. Apparently, Company A can achieve the same capacity certainty with less costs by purchasing less machines than currently available. Recall from Chapter 2 that the reason for introducing double shifts was a shortage of dryers. The coating pans have been purchased with regular shifts in mind. By optimizing machine procurement decisions and workforce planning simultaneously, we find that Company A is able to achieve the same capacity certainty with lower costs, by purchasing less dryers and using more double shifts (i.e. workforce flexibility). To confirm this conclusion, we optimize the strategic capacity plan for Strategy 90 (which is closest to the current strategy) while fixing the workforce planning to the default plan for double shifts. We find that the number of P060 increases from one to two and the number of P100 increases from 6 to 8, showing that we indeed are able to save on procurement costs by jointly optimizing procurement decisions with workforce flexibility.

Takeaway 2
Company A is able to save costs by jointly optimizing workforce flexibility and machine procurement decisions. They can reduce the number of coating pans by using more double shifts.

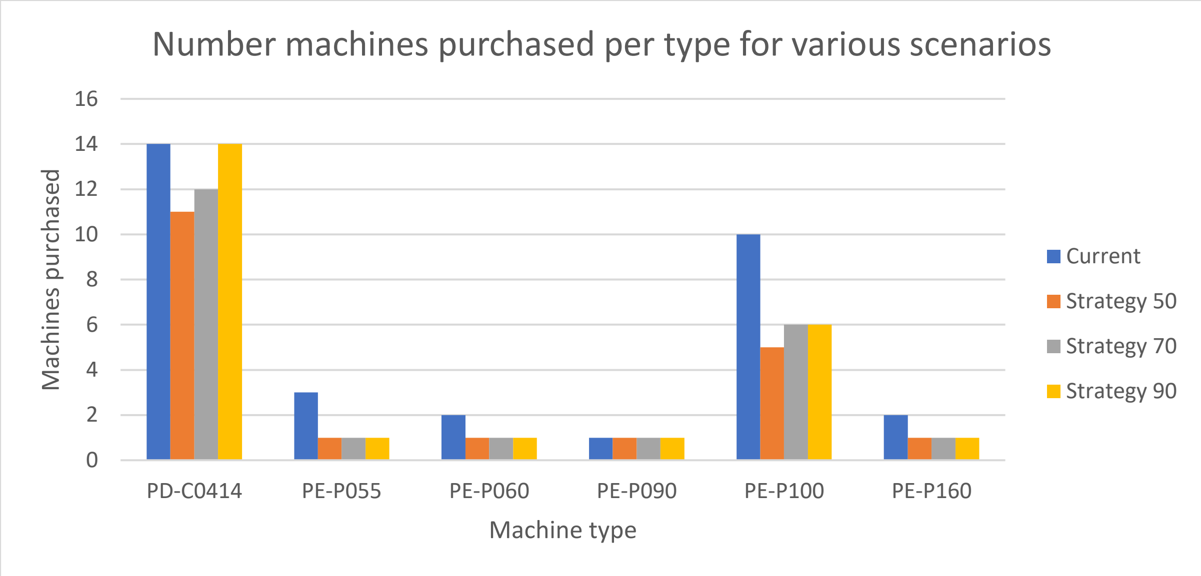

Figure 5-12. Number of machines purchased per type for various scenarios.

Figure 5-13, Figure 5-14, and Figure 5-15 shows the capacity demand and available capacity for capacity strategy 50, 70, and 90, respectively. The information is shown for each month, the months on the axis are not completely shown due to the limited space. We observe a larger capacity reserve for capacity strategies with a higher coverage probability, because the number of machines increase. We also observe strong demand seasonality and no trend. Company A expects a downward trend in the future, but currently demand has been stable, so the results are consistent with expectations. Demand increases to a peak 2080 hours in March from a low 777 hours in August. Capacity strategy 50 has a 50% chance to result in undercapacity when the maximum available capacity is equal to Scenario 50. This is the case in March for dryers C0414, therefore it is not in line with Company A's business strategy to be a reliable and flexible partner. Instead, we suggest Company A to use either strategy 70 or strategy 90. We discuss capacity costs in Section 5.2.3.

Takeaway 3

We suggest that Company A does not use strategy 50, because it too often results in capacity shortages. Instead they should use strategy 70 or 90.

The peak in March at 2080 hours is significantly higher than capacity demand in other months, with February at 1585 hours and April at 1466 hours as second and third. This makes a strong case for demand smoothing. Recall from Chapter 1 that demand smoothing is difficult for Company A, because of time windows for specific crops. However, we suggest Company A to further investigate this opportunity, as the model provides more information about the timing and potential financial benefits of demand smoothing. Furthermore, demand smoothing between a few weeks at the end of February and start of April may already make a big difference. We discuss this future research direction in Chapter 6.

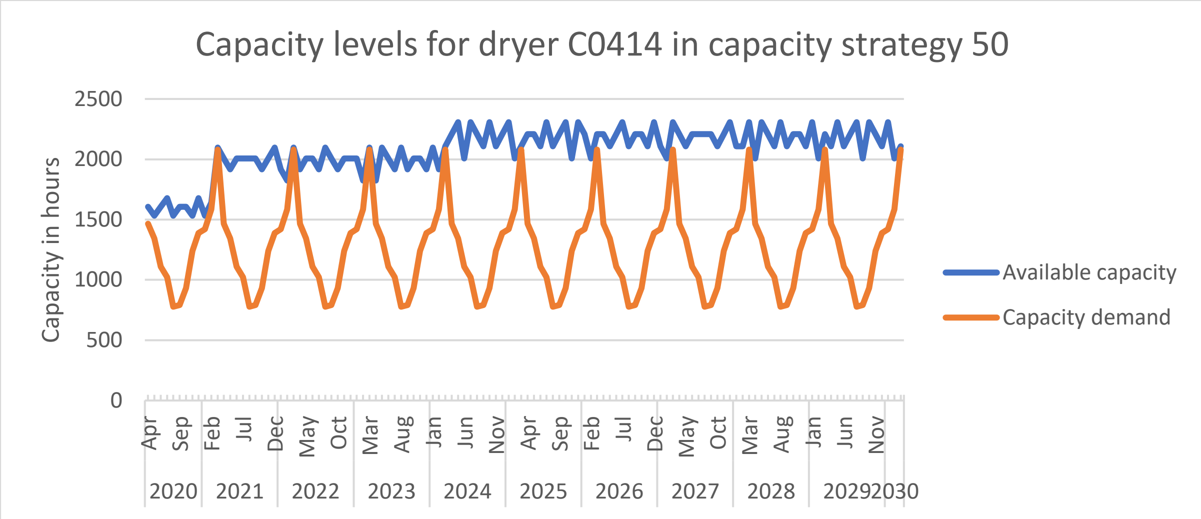

Figure 5-13. Capacity demand and available capacity for capacity strategy 50.

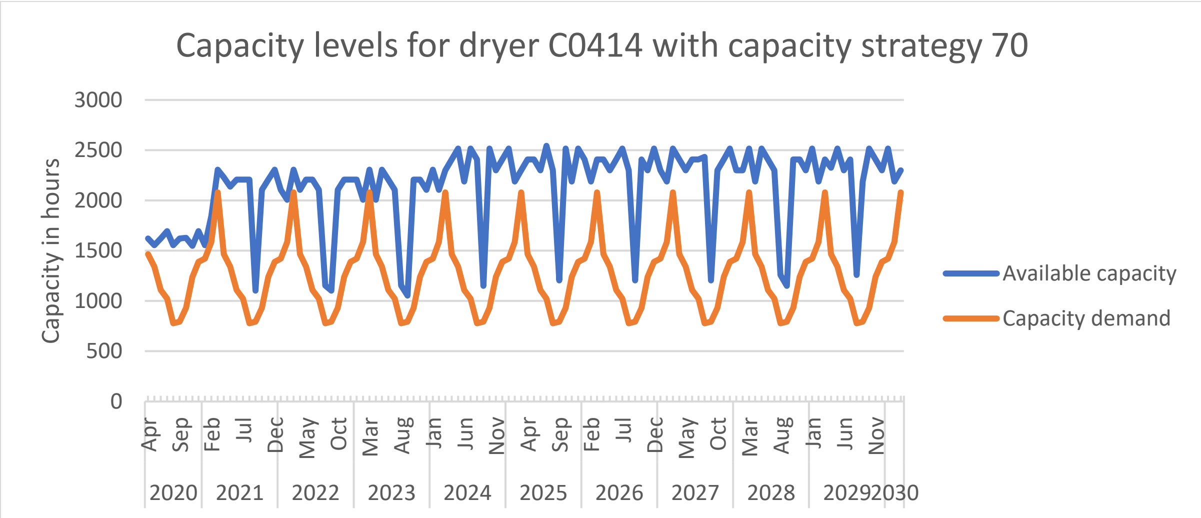

Figure 5-14. Capacity demand and available capacity for capacity strategy 70.

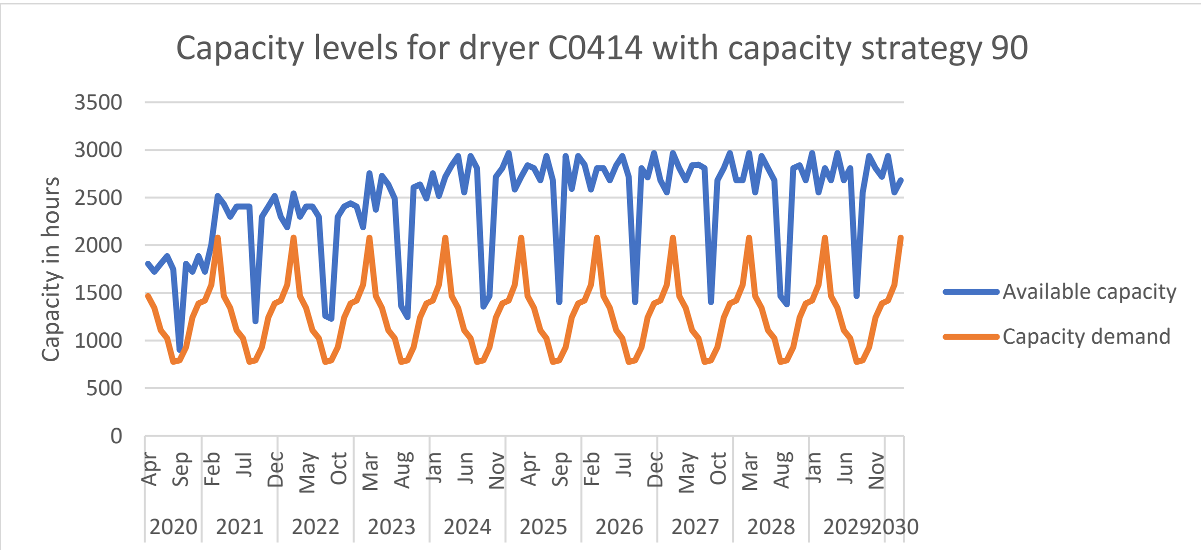

Figure 5-15. Capacity demand and available capacity for capacity strategy 90.

In Figure 5-14 and Figure 5-15 we observe large downward spikes in available capacity, which are not visible in Figure 5-13. These are downward spikes are the use of single shifts, while double shifts are used in all other months. Recall from Chapter 2 that double shifts effectively double machine capacity. Capacity strategy 50 requires the use of double shifts in each month, while the other strategies use single shifts in months with low capacity demand. The latter is possible because the number of machines is higher for capacity strategy 70 and 90. We discuss the cost implications of this interaction in Section 5.2.3. For now, we conclude that it is more cost-efficient to use workforce flexibility (double shifts) instead of purchasing more machines in general.

Takeaway 4
It is more cost-efficient to use double shifts than to purchase more machines to satisfy capacity demand.

5.2.3 Comparing costs

Capacity strategies with a higher coverage probability are more expensive, which we observe in Figure 5-16. As the coverage probability increases, the one-tailed prediction interval of capacity demand increases at an increasing rate. Therefore, we expect that each additional percent of confidence comes at an increasing cost. However, because machine purchases are an integer decision, the costs increase in jumps. Coincidentally, the increase from strategy 50 to strategy 70 is 423k euros, while the increase from strategy 70 to strategy 90 is 383k euros. We also generated results for strategy 95, which saw an 247k cost increase. The latter confirms our expectation that costs increase at an increasing rate for higher coverage probabilities.

Takeaway 5
Cost increases in jumps with an increasing coverage probability. Costs increase at an increasing rate for higher coverage probabilities.

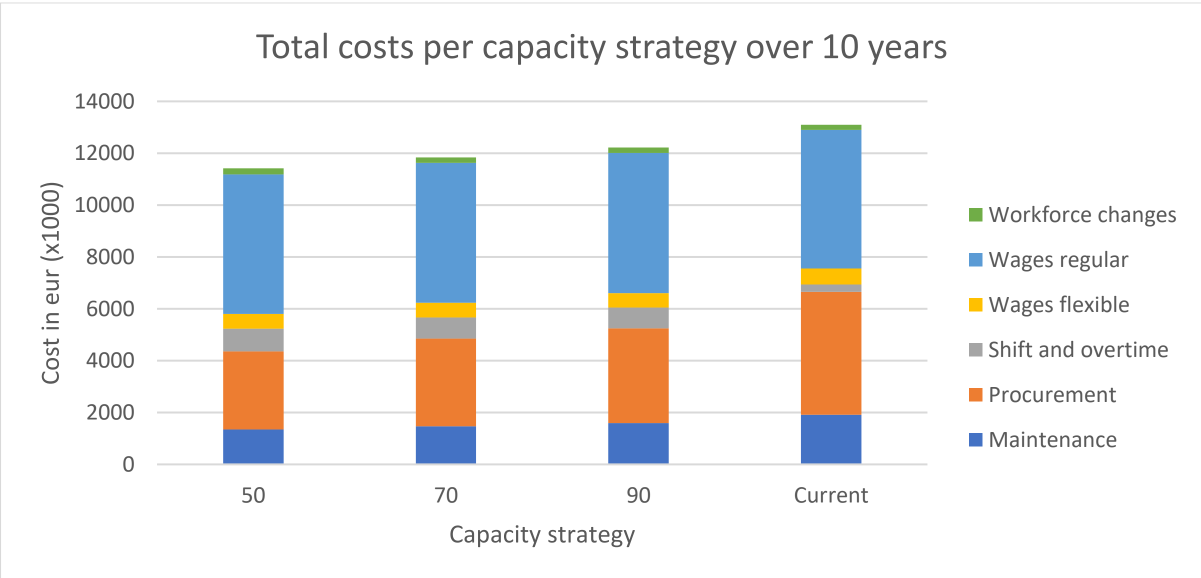

Figure 5-16. Total costs per capacity strategy.

A cost increase of less than 4% to move from 50% to 70% confidence does not seem very large. The reason lies with the trade-off between strategic decisions and workforce flexibility. We compare the costs for the strategic decisions (i.e. procurement and maintenance) with the costs for workforce in Figure 5-17. We observe that the increase in procurement and maintenance costs

UNIVERSITY
OF TWENTE.

is partially offset by a decrease in workforce costs. This decrease is due to requiring less workforce flexibility measures, such as double shifts and overtime, because there are more machines available.

Takeaway 6
Increasing the coverage probability for the capacity strategy is relatively inexpensive, because increases in machine procurement are partially offset by decreases in workforce flexibility costs.

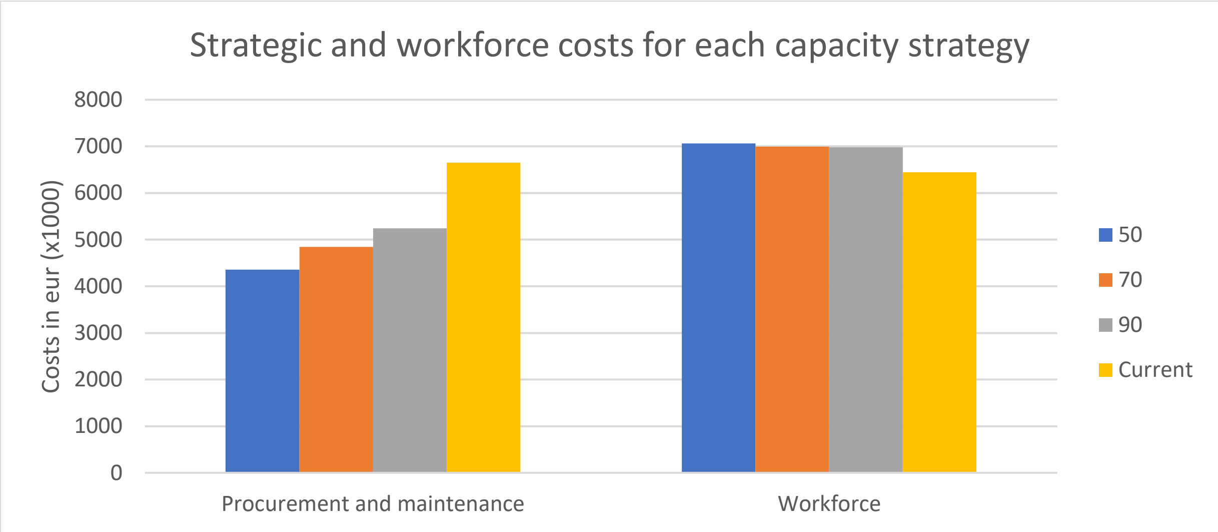

Figure 5-17. Costs for strategic decisions and workforce planning between capacity strategies.

5.3 IMPACT OF JUDGMENTAL ADJUSTMENTS

In this section, we explore how the capacity strategies change when considering judgmental sales forecasts. The adjustment factors are the result of an interview with a senior account manager and a market analyst. Recall from Chapter 4 that these adjustment factors are a multiplier for the historical capacity demand for a certain product and/or customer. In general, Company A expects a 10% demand decrease in Enkhuizen for the four largest customers, due to two reasons. First, customers are decentralizing their seed processing operations. Second, customers are pursuing vertical integration, moving towards in-house production. Medium and small customers do not have the resources and scale to do this, so demand for these customers is expected to remain the same.

We denote an adjusted scenario, and the respective capacity strategy, with a coverage probability of 50 as “50-A”, for example. Figure 5-18 shows the impact of judgmental adjustments. We observe that a 10% reduction for the top four customers impacts each machine and month differently. For example, the demand for April for dryer C0414 is reduced by only 3.8%, while the demand for the same dryer for May is reduced by 9.2%.

Takeaway 7
The impact of judgmental adjustments on capacity demand varies for each machine and month.

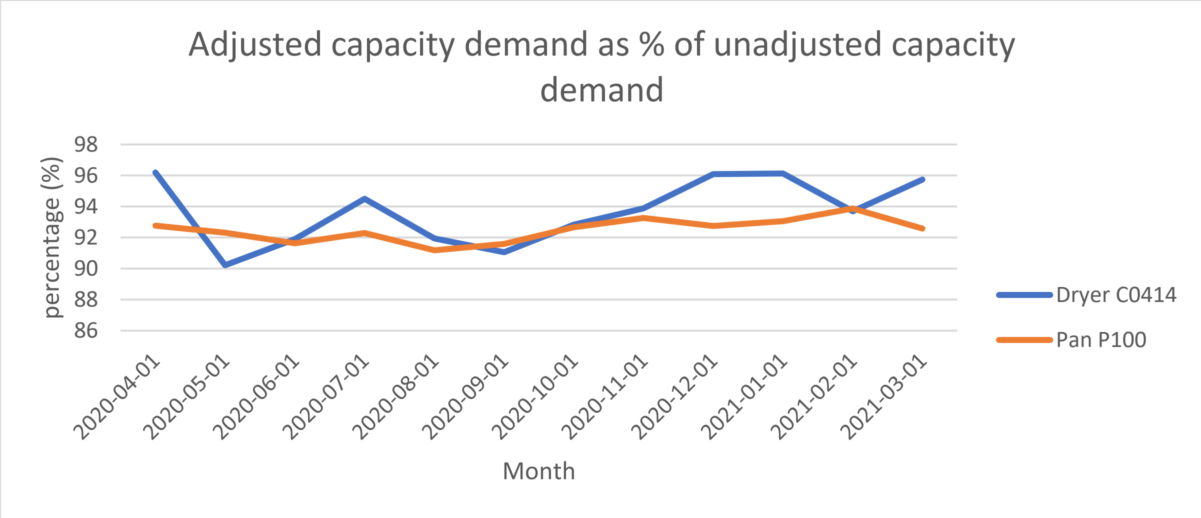

Figure 5-18. Impact of judgmental forecasts on capacity demand scenarios.

To generate capacity strategies, we follow the same approach as in Section 5.2, using the same coverage probabilities, except that we use adjusted scenarios. The capacity strategies (i.e. set of machine procurement decisions) based on adjusted scenarios are identical to those in Figure 5-12. We observe that there are no changes to the capacity strategies. The reason is because the impact of the judgmental forecasts is not large enough to require a change in the capacity strategy. In case of more drastic judgmental forecasts (e.g. adjustment factors further from one), the strategies are likely to change.

Takeaway 8
The optimal capacity strategy for each coverage probability do not change when using adjusted scenarios, based on judgmental sales forecasts.

Recall from Chapter 4 that we designed a second way to include judgmental forecasts, by means of future sales orders. The purpose of this method is to determine the impact of new products. Company A has launched a new organic product range in January 2020. They purchased one dryer and coating pan specifically for these organic products. However, it is too early to provide sales forecasts for this product range, because customers are currently doing test runs with Company A. According to the senior account manager, this method will be very useful next year, once the testing has finished and commercial orders will arrive.

5.4 Compare capacity strategies

In this section, we evaluate the capacity strategies discussed in Section 5.2 and Section 5.3. This is the sixth and final step of the model. To do this, we fix the machine procurement decisions according to the capacity strategy and optimize the workforce planning for the pessimistic (30%), realistic (50%), and optimistic (70%) scenario. We do this for each capacity strategy alternative, corresponding to the coverage probabilities 50%, 70% and 90%.

Figure 5-19 shows the costs of each capacity strategy, where we group the three capacity strategies (50, 70, 90) per scenario, such that we can identify the differences between the capacity strategies for each scenario. As expected, the capacity strategy with a lower coverage probability is always less expensive than with a higher coverage probability. Machine procurement cost is always higher than workforce measures such as double shifts.

Takeaway 9
Capacity costs decrease for scenarios with lower coverage probability, because less workforce flexibility is required to satisfy demand.

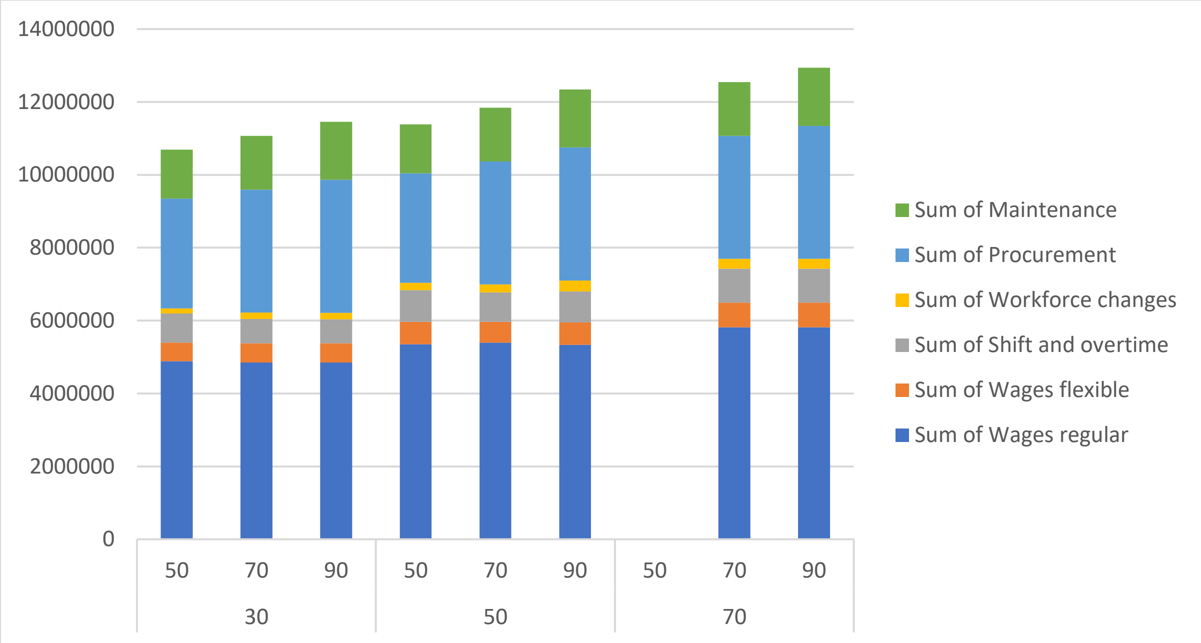

Figure 5-19. Total costs of each capacity strategy in each scenario.

The most important takeaway from this is that there is little flexibility for optimal capacity strategies, because all workforce flexibility is already used to reduce the procurement costs. Strategy 50 is infeasible in scenario 70, which is why there is no data displayed in the graph for this instance. Therefore, it is important for Company A to go for a strategy with sufficient certainty.

Takeaway 10
Flexibility is low for optimal strategies, as it is optimal to use as much workforce flexibility as possible to reduce machine procurement cost.

5.5 Sensitivity analysis

In this section, we answer the following research question.

(Q 5.3) What is the sensitivity of the model regarding parameters subject to uncertainty or change?

Our model for strategic capacity planning is only valid when the input parameters are accurate. Section 5.1 verified the accuracy of the capacity demand scenarios. There are two parameters that are subject to uncertainty and change: machine lifespan and maximum Overall Equipment Effectiveness (OEE). We perform a sensitivity analysis on these two parameters.

The machine lifespan is at least 10 years. Once a machine has been used for 10 years, Company A must make the decision to either continue with maintenance or purchase a new machine. Maintenance costs increase significantly past 10 years, because around that time, the more expensive components must be replaced. We do not take this increase in maintenance cost into account in this model, because Company A is not able to make an estimate for new machines on the maintenance cost curve over the years. Instead, we test the sensitivity of the capacity strategies to an increased lifespan. The current machines are pushed to a maximum lifespan of 16 years, so we test the sensitivity with 16 years. An increase in lifespan can also be modeled as a decrease in procurement costs, while keeping the planning horizon constant. We adjust the machine procurement costs by a factor of 0.625 (10 year lifespan /16 year lifespan = 0.625). We obtain capacity strategies by following the same steps as in Section 5.2.

As lifespan increases, the cost per year of use becomes lower, which makes machine capacity relatively less expensive. Therefore, we expect either an increase or no change in the number of machines purchased within one machine lifespan. It might be more cost-efficient to use less workforce flexibility and more machine capacity. Figure 5-20 shows the resulting capacity strategies. When we compare these strategies with the original capacity strategies (Section 5.2), we see one additional machine procured of coating pan P060 and P100 for strategies 70 and 90. Capacity strategy 50 does not change. The longer lifespan does not impact C0414, because its procurement costs outweigh the costs of workforce flexibility.

Takeaway 9
As machine lifespan increases, the number of machines for lower coverage probabilities increases only for P100, such that less double shifts are required. Strategy for level 90 does not change.

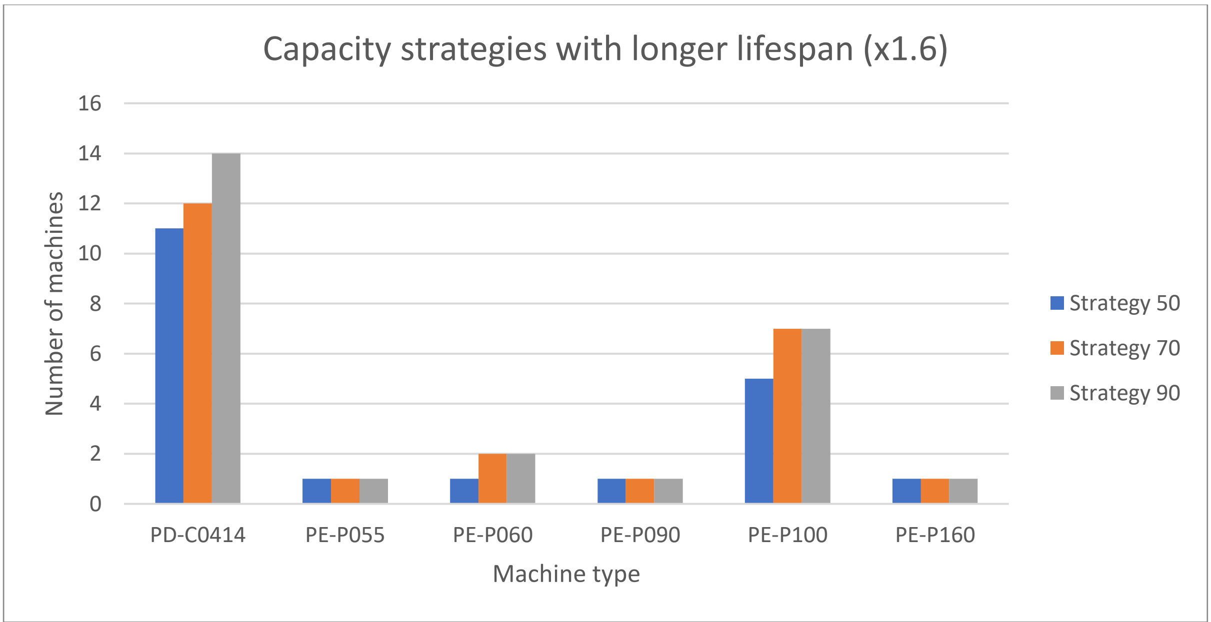

Figure 5-20. Capacity strategies with a longer lifespan (or: lower procurement cost).

The next parameter in our sensitivity analysis is maximum OEE. The maximum OEE is the only parameter in calculating available capacity that is subject to variability and change. We determined the maximum OEE based on monthly processing times over the past two years. We found that the maximum OEE is different every year, most likely depending on the variety in product demand. Furthermore, maximum OEE can change through improvement efforts. Several lean projects have improved OEE in the past. We test for both an OEE increase and decrease of 10%. This covers the variability observed the past two years.

As OEE increases, we expect a decrease in machine procurement, because less machines are required to do the same work.

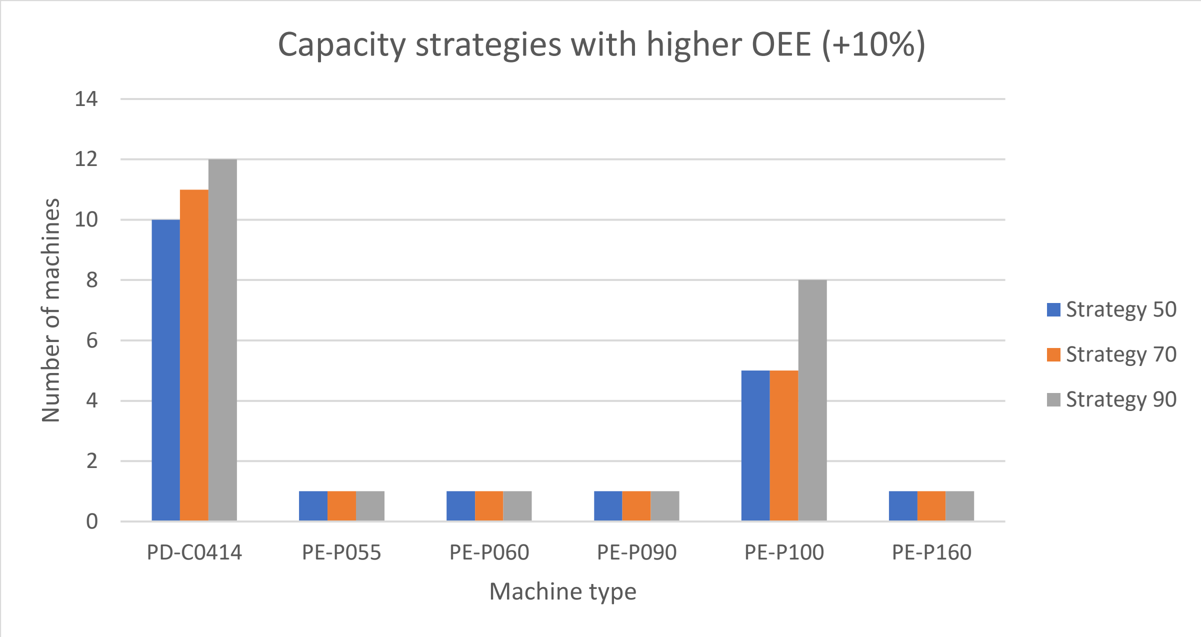

Figure 5-22 shows the capacity strategies with higher OEE, and Figure 5-21 shows the changes of these compared to the default capacity strategies. We observe a decrease for dryer C0414 for each strategy. With a higher OEE, less dryers are required to do produce the same output. There is one exception, for coating pan P100 with strategy 90. Because Company A already needs a certain number of pan P100 to deal with peak demand at a 90 coverage probability, it is cost-efficient to increase the number of pans even more, such that Company A can meet capacity demand with less double shifts.

Takeaway 10
An increase in OEE generally leads to a reduction in the number of machines. However, at a 90 coverage probability, it is cost-efficient to purchase more P100's, to be able to work with less double shifts.

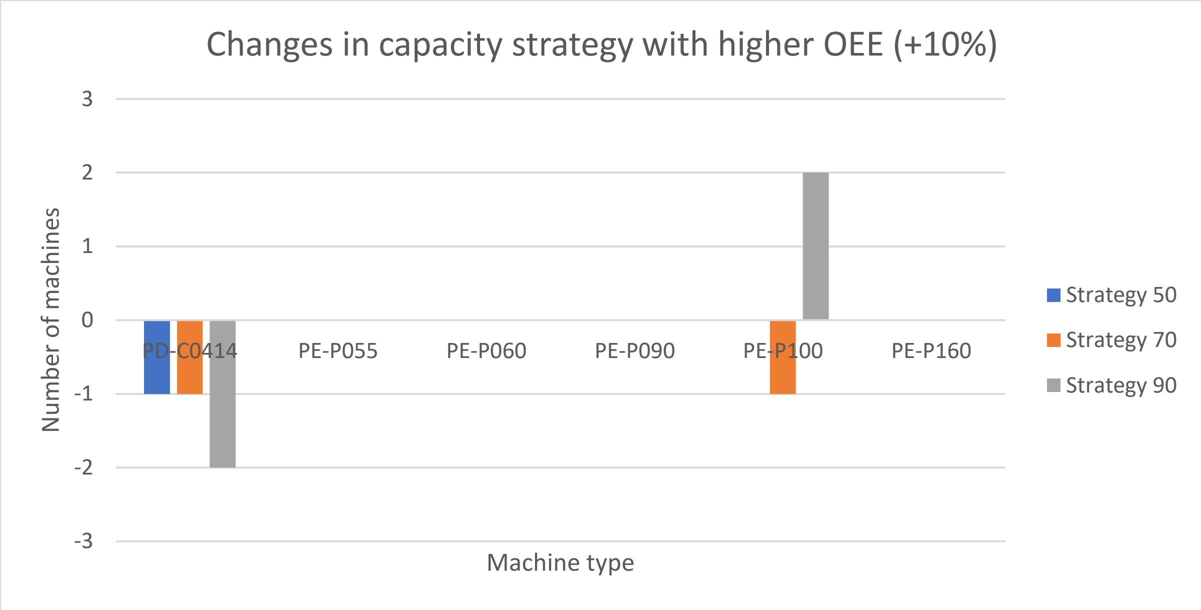

Figure 5-21. Changes in capacity strategies with higher OEE.

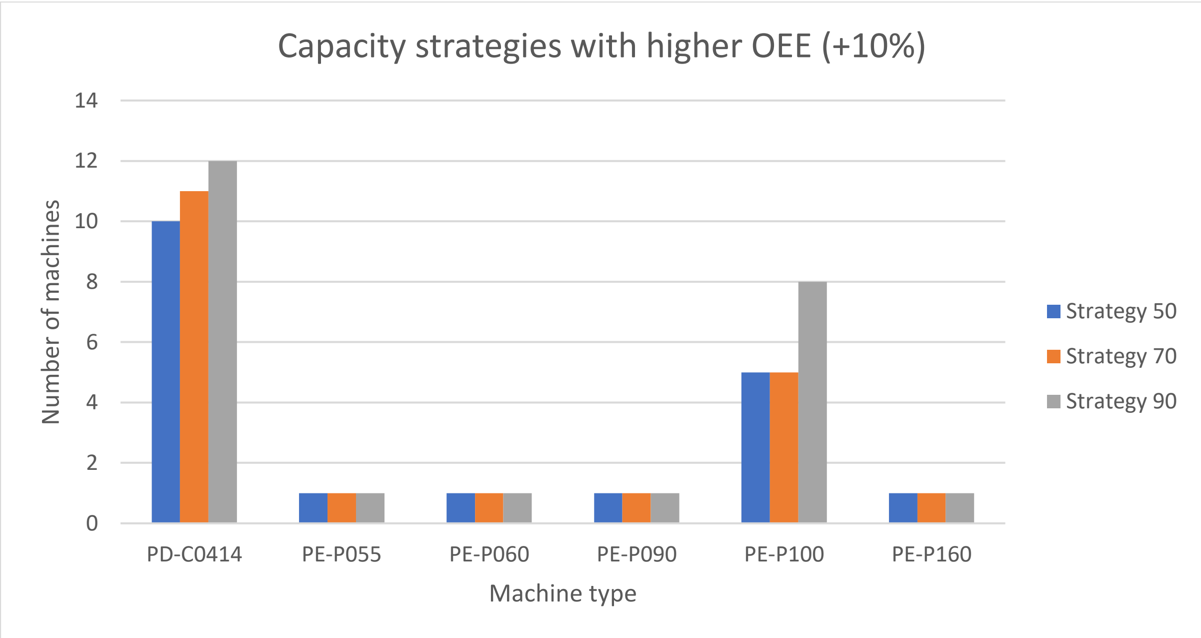

Figure 5-22. Capacity strategies with higher OEE.

As OEE decreases, we expect an increase in machine procurement, because more machines are required to do the same work. Figure 5-24 shows the capacity strategies with higher OEE and Figure 5-23 shows the changes of these compared to the default capacity strategies. We indeed observe an increase for dryer C0414, P060 and P100. Some machine types do not require an increase, because they had sufficient overcapacity in the default strategies. This happens because machine capacity increases in integer steps. Most notably, additional coating pan P100 does not have to be purchased in strategy 70, while this is the case for strategy 50 and 90. Six P100's is sufficient to fulfill demand in scenario 70 even with reduced OEE. Strategy 50 has sufficient P100's for scenario 70 as well (i.e. 6), so the reason that it was infeasible is the number of dryers.

Takeaway 11

A decrease in OEE leads to an increase in the required number of machines, except when there is sufficient overcapacity, because machine capacity increases stepwise.

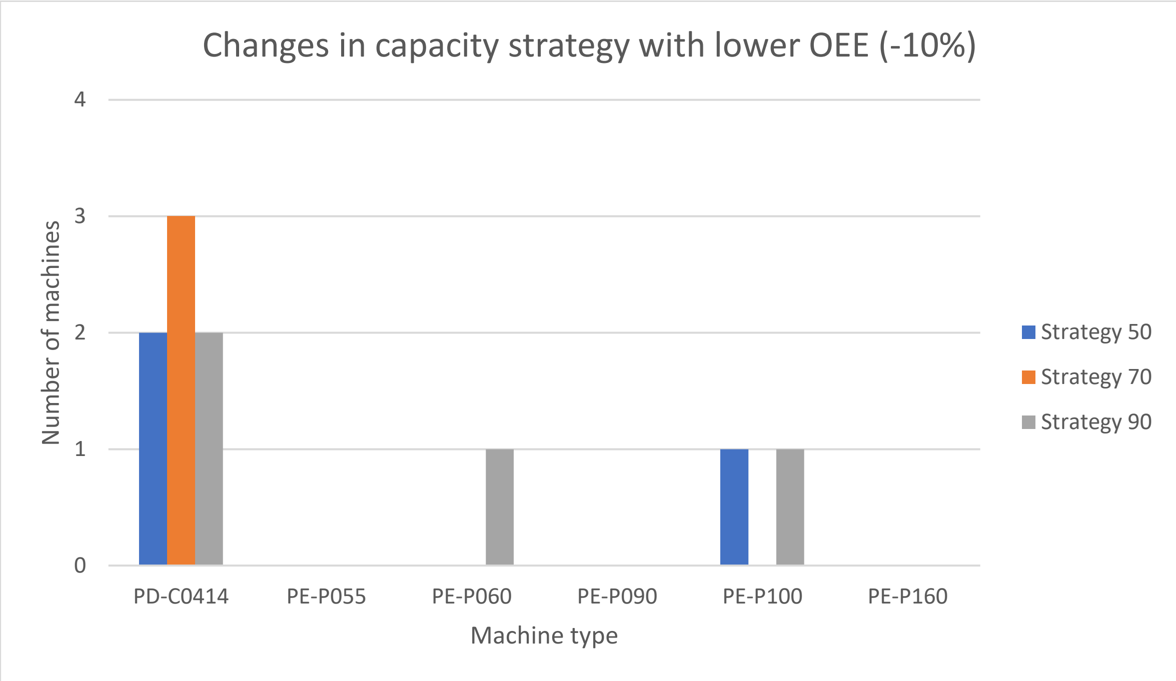

Figure 5-23. Changes in capacity strategies with lower OEE.

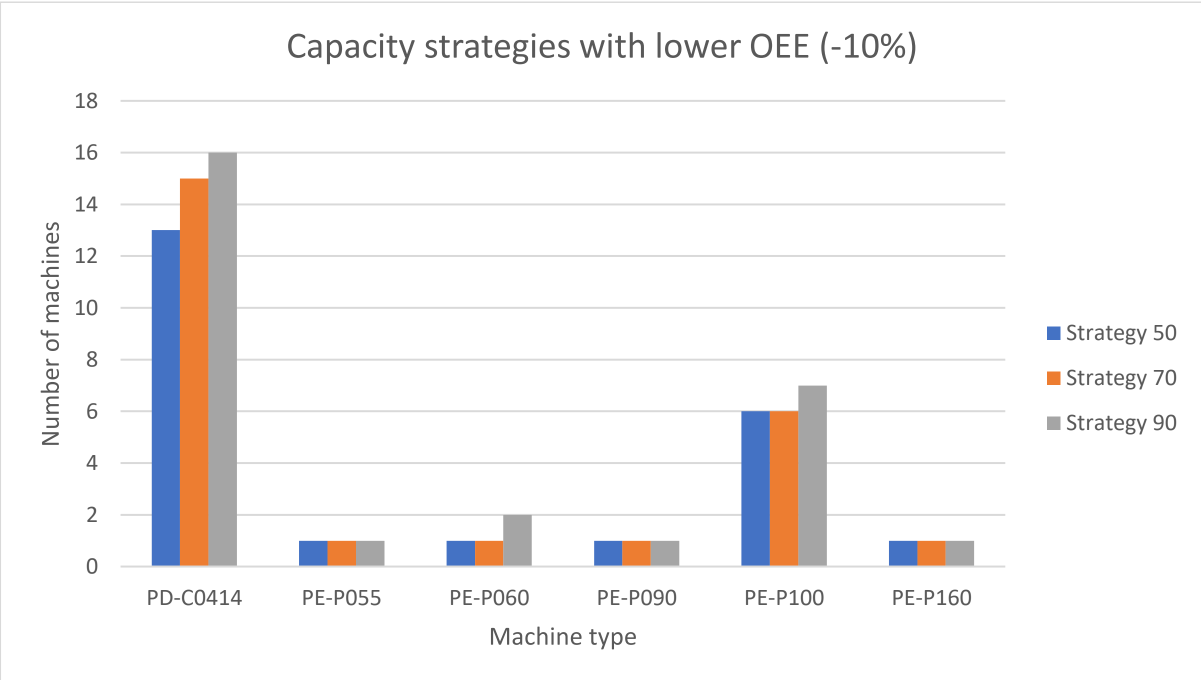

Figure 5-24. Capacity strategies with lower OEE.

5.6 DISAGGREGATING CAPACITY DEMAND

Section 4.1 describes a method to verify our assumption that capacity demand is sufficiently smooth within each month. If this assumption does not hold, the capacity plan that results from the model results in more capacity shortages than indicated by the coverage probability. We apply the method to the two most critical machines in terms of variability and quantity: dryers C0414 and coating pan P100. Figure 5-25 and Figure 5-26 show the average percentage of months and weeks with sufficient capacity, which are obtained by replicating the method 100 times. We choose 100 replications, because the average number of weeks does not change significantly (< 0.1 number of weeks) after 100 replications.

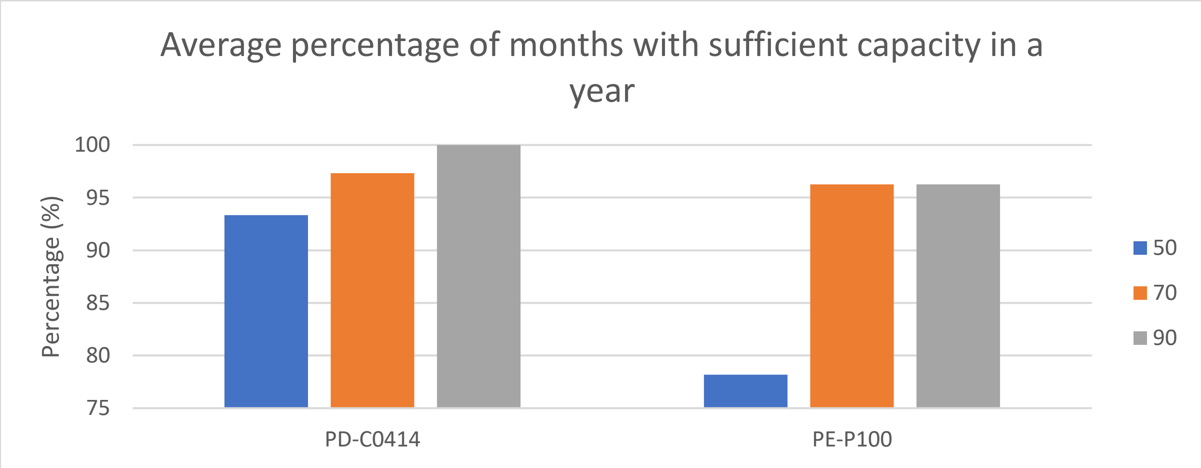

Figure 5-25. Average percentage of months with sufficient capacity.

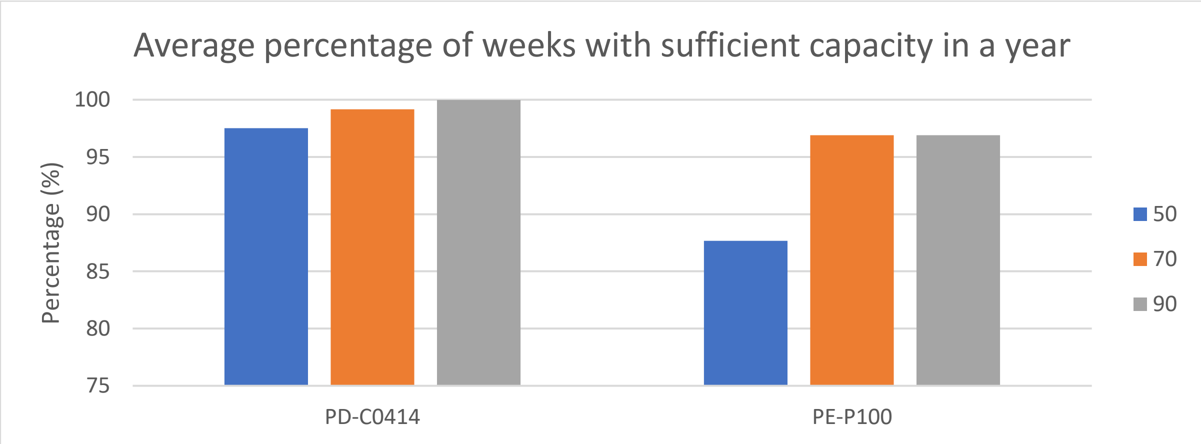

Figure 5-26. Average percentage of weeks with sufficient capacity.

Note that strategy 70 and 90 and equal for coating pan P100, therefore the percentage does not change. As expected, we observe an increase in the average number of weeks with sufficient capacity with an increasing coverage probability of each strategy. The percentage of weeks with sufficient capacity is much higher than the coverage probability of each capacity strategy. This is expected, because the capacity demand (based on the coverage probability) is a constraint for the available capacity. If the available capacity is exactly equal to the capacity demand, we expect that the percentage of weeks is close to the coverage probability. However, because the available capacity must satisfy peak demand, there is overcapacity in non-peak months. This overcapacity is a buffer against variability beyond the coverage probability, resulting in a higher percentage of weeks with sufficient capacity.

We observe that the average percentage is generally higher for weeks than for months, which means that disaggregation actually improves capacity coverage. Therefore, our assumption that demand is sufficiently smooth within each month holds.

We observe that the impact of a lower coverage probability is much higher for P100 than for C0414. The reason is that there is only one peak month for C0414, while there are four peak months for P100. Therefore, there is less overcapacity of P100 to absorb weekly variability. This highlights the issue of strategic capacity planning in case of highly seasonal demand. Many machines are required to deal with peak demand, while these are not needed during the rest of the year.

6 CONCLUSION

6.1 CONCLUSION

The main research question of this research is as follows.

> How can machine- and operator capacity planning deal with seasonal and variable demand to improve on-time delivery in a cost-efficient way?

To answer this research question, we have addressed three underlying core problems. The first core problem is that the current method to calculate capacity demand is inaccurate, it only considers the number of orders. We have designed a calculation model that is able to determine the monthly capacity demand for each machine type. Calculating per machine type is especially important, because seasonal patterns are different for each machine, due to crop seasonality. Many machines are only used for specific crops. We have shown that our calculation model is more accurate than the current method. Furthermore, our calculation model uses sales orders as input, so that we can calculate future capacity demand from judgmental sales forecasts.

The second core problem is that demand forecasts are unreliable, because Company A uses a seasonal naïve forecasting method, which does not reflect demand uncertainty. We have designed a forecasting method that is able to reflect demand uncertainty through prediction intervals, based on historical data from the past five years. To further improve the forecast reliability, we have designed a method to include judgmental forecasts. This methods involves adjustment factors and future sales orders, which are translated to capacity demand using our calculation model.

The third core problem is that Company A's capacity plan is misaligned; capacity decisions are taken individually, resulting in cost inefficiencies. We have designed an optimization model that is able to determine the optimal capacity plan for both machine- and operator decisions, given a set of input parameters. We have shown that the resulting capacity plan can achieve the same level of capacity feasibility as the current capacity plan, while saving costs.

The question was how Company A can improve on-time delivery. We have shown that Company A has sufficient capacity at a tactical (i.e. monthly) and strategic level to satisfy capacity demand with a 90% certainty. Therefore, we do not need to increase capacity levels compared to the current situation to improve on-time delivery. Instead, we suggest Company A to improve their scheduling methods to make more efficient use of the available capacity, which we discuss in Section 6.3. Using our model and its results, Company A can achieve the same delivery performance while being more cost-efficient. The model results show that Company A can reduce their total costs by reducing machine procurement and increasing the use of workforce flexibility.

The machine- and operator capacity planning can deal with seasonal demand by using double shifts during mid-season and peak-season, while using single shifts during low-season (August and September) to save costs. Regarding the strategic decisions, the number of machines procured are sufficient to deal with demand peaks, while using double shifts, at the desired coverage probability. If Company A is able to smooth these demand peaks, the number of machines can be further reduced. The machine- and operator capacity planning can deal with variable demand by using a higher coverage probability for the capacity strategy. We recommend a coverage probability of at least 70%, while 90% would be even better to deal with variable demand. Of course, that comes at a cost.

Possibly even more important is that Company A is able to use this model on a yearly basis to reevaluate their strategic and tactical capacity plan, as new information comes available. The model is able to use not only the most recent sales and production data, but also the latest judgmental sales forecasts to determine the optimal strategic and tactical capacity plan. The model enables Company A to make a trade-off between costs and delivery performance, by varying the coverage probability of the capacity demand scenarios.

6.2 CONTRIBUTIONS TO THEORY

Our forecasting method is primarily based on Hyndman's state space approach (2008) and our optimization model is primarily based on the strategic capacity optimization model by Bihlmaier et al. (2009) We made three contributions to theory.

6.2.1 Transforming product demand to capacity demand

Piecewise linear transformations are commonly used in capacity optimization models to determine capacity demand from product demand forecasts. For Company A, this method results in inaccurate capacity demand forecasts, and thus an inaccurate capacity plan. We designed a method to address this problem. First, we calculate the historical capacity demand from historical product demand, using a highly customized calculation model. Second, we forecast the capacity demand, instead of the product demand. Finally, we take this capacity demand forecast as input for the capacity optimization model.

We suggest future research to formulate a more generic base calculation model, that can be used for other companies. Future research can identify the ways in which capacity demand has a nonlinear relation with product demand. Company A's case of seed quantity is only one example.

6.2.2 Adjusted capacity demand forecasts

In literature, demand scenarios are often defined as a demand quantity and probability, based on which the capacity strategy is optimized using a stochastic model. For Company A however, capacity demand does not increase linearly with demand quantity. A larger demand quantity in few large orders results in less capacity demand than a lower demand quantity in many small orders. Not only that, but the machines used can be totally different. We designed two methods to deal with this: adjustment factors and future sales orders.

The first method uses adjustment factors that can be applied to either specific customers, products or a combination thereof. We multiply the historical capacity demand for the respective customers and/or products by this adjustment factor, before generating forecasts. This method allows Company A to take external effects into account, such as the largest customers moving towards in-house production and the banning of agrochemicals used in seed treatments.

The second method uses sales order forecasts (i.e. number and size of orders for a product) to calculate the capacity demand. These sales order forecasts are especially useful for products where history is not representative for the future, such as new products. It allows Company A to determine in advance how much capacity they need to deal with future demand for their organic treatments.

6.2.3 Using prediction intervals to reflect demand uncertainty

To deal with demand uncertainty in capacity planning, there are two widely used methods in literature. One is simulation modeling and the other is stochastic optimization. We designed a new method that requires less input data than simulation modeling and remains accurate when the production process undergoes changes.

We use one-tailed upper prediction intervals as capacity demand scenarios. It tells us that the true capacity demand will be lower than the prediction interval with a certain probability. Compared to a confidence interval, the prediction interval must account for both the uncertainty in knowing the value of the mean and the data scatter, so it will always be wider than the confidence interval. These prediction intervals provide a simple and intuitive way to find alternative capacity strategies, such that management can decide on the trade-off between costs and certainty.

However, the interpretation of the resulting capacity strategy is more difficult than coverage probabilities from simulation models. Statistically speaking, the probability is for each machine in each month, so not for the entire capacity plan. Furthermore, we use the prediction interval as a capacity constraint, Often the available capacity will be higher than the capacity demand, therefore having a higher certainty.

We suggest future research to evaluate the statistical implications of using prediction intervals for capacity demand scenarios in a capacity optimization model.

6.3 RECOMMENDATIONS FOR PRACTICE

We structure our recommendations in three parts. First, the results of our model for the strategic and tactical capacity plan. Second, the implementation and use of our model. Third, further research that Company A can pursue.

6.3.1 Strategic and tactical capacity plan

When it comes to capacity strategy, we suggest that Company A stays at their current level of 90% certainty, while being more cost-efficient. That way on-time delivery will not decrease, because the probability of no capacity shortage remains the same. We suggest not to reduce the certainty to 70%, because the cost decrease is only 3.2% (i.e. 383,000 over 10 years) in the expected scenario. Furthermore, high delivery reliability is in line with Company A's business strategy. The cost increase is relatively low for a higher coverage probability, because increases in procurement costs are offset by decreases in workforce costs.

Company A can increase cost-OEE by working with more double shifts, and use overtime incidentally to deal with peaks. To be specific, our model results suggest to use double shifts in all months except August and September, while currently Company A uses double shifts during seven months. Using more of this workforce flexibility allows Company A to purchase less machines. More specifically, when purchasing new coating pans, Company A can reduce the number of coating pans P055 by two, P060 by one, P100 by four, and P160 by one. Our model suggests to keep the same number of dryers to satisfy demand with 90% certainty each month.

6.3.2 Model implementation and use

The implementation of our model has been designed for practical use. We have created a simple Excel dashboard from which all models can be run using the click of a button. Furthermore, the basic parameters can be changed on this dashboard. More advanced adjustments can be made in the optimization Excel file. Interactions with external tools, such as the calculation and forecasting model in R, have been automated. Both documentation and training has been provided to ensure Company A will be able to use the model in the future independently.

6.3.3 Further research opportunities

Recall from Chapter 5 that we have shown there is sufficient capacity on a strategic and tactical level. The cause for decreased on-time delivery performance during winter is probably a scheduling issue. Company A currently uses backward scheduling, working back from the requested delivery date. Each order is scheduled as soon as it is released and the schedule cannot be changed afterwards. It is no surprise that this inflexible scheduling method leads to delivery issues. There are methods available to improve scheduling, such as the shifting bottleneck heuristic. Recall from Chapter 1 that we did not choose to improve the scheduling method for two reasons. First, because Company A must use SAP to plan orders, so the implementation of a better scheduling method can be difficult. Second and most important, Company A first wanted to have insight in their strategic and tactical capacity plan before diving in on an operational level. Now that we have provided insight in their strategic and tactical capacity plan, we suggest Company A to take the next step.

A second research opportunity is demand smoothing. Recall from Chapter 1 that we did not consider this method, because there is a small time window for orders due to crop seasonality. However, the results in Chapter 5 show that smoothing over just a few weeks around March can be extremely beneficial to reduce the number of C0414 dryers required. Our model also shows the potential cost savings that can be realized by demand smoothing. These can be used to incentivize customers to shift their demand by a few weeks, which might make customers surprisingly flexible.

7 REFERENCES

Ahmed, S. & Sahinidis, N. V., 2003. An approximation scheme for stochastic integer programs arising in capacity expansion. *Operations Research,* 51(3), pp. 461-471.

Balzer, W. K., Doherty, M. E. & O'Connor, R., 1989. Effects of cognitive feedback on performance. *Pyschological Bulletin,* Volume 106, pp. 410-433.

Bergmeir, C. H. R. J. B. J. M., 2016. Bagging exponential smoothing methods using STL decomposition and Box-Cox transformation. *International Journal of Forecasting,* 32(2), pp. 303-312.

Bihlmaier, R., Koberstein, A. & Obst, R., 2009. Modeling and optimizing of strategic and tactical production planning in the automotive industry under uncertainty. *OR Spectrum : Quantitative Approaches in Management,* 31(2), pp. 311-336.

Box, G. E. P. & Jenkins, G. M. R. G. C. L. G. M., 2015. *Time series analysis: Forecasting and control.* 5th ed. New Jersey: John Wiley & Sons.

Bradley, J. R. & Glynn, P. W., 2002. Managing capacity and inventory jointly in manufacturing systems. *Management Science,* 34(2), pp. 273-288.

Buehler, R., Messervey, D. & Griffin, D., 2005. Collaborative planning and prediction: Does group discussion affect optimistic biases in time estimation?. *Organizational Behavior and Human Decision Processes,* 97(1), pp. 47-63.

Carbone, R. & Gorr, W., 1983. Accuracy of judgmental forecasting of time series. *Decision Sciences,* Volume 16, pp. 153-160.

Chand, S., McClurg, T. & Ward, J., 2000. A model for parallel machine replacement with capacity expansion. *European Journal of Operations Research,* 121(3), pp. 519-531.

Chen, Y. Y., Chen, T. L. & Liou, C. D., 2013. Medium-term multi-plant capacity planning problems considering auxiliary tools for the semiconductor foundry.. *International Journal of Advanced Manufacturing Technology,* 64(9-12), pp. 1213-1230.

Chopra, S. & Meindl, P., 2013. *Supply Chain Management.* Essex: Pearson Education Limited.

Eppen, G. D., Martin, R. K. & Schrage, L., 1989. OR Practice—A Scenario Approach to Capacity Planning. *Operations Research,* Volume 37(4), pp. 514-673.

Fleischmann, B., Ferber, S. & Henrich, P., 2006. Strategic Planning of BMW's Global Production Network. *INFORMS Journal on Applied Analytics,* 36(3), pp. 194-208.

Gardner, E. S., 1985. Exponential smoothing: The state of the art. *Jornal of Forecasting,* 4(1), pp. 1-28.

Geng, N. & Jiang, Z., 2009. A review on strategic capacity planning for the semiconductor manufacturing industry. *International Journal of Production Research,* 47(13), pp. 3639-3655.

Goodwin, P., 2000. Correct or combine? Mechanically integrating judgmental forecasts with statistical methods.. *International Journal of Forecasting,* Volume 16, pp. 261-275.

Goodwin, P. & Fildes, R., 1999. Judgmental forecasts of time series affected by special events: Does providing a statistical forecast improve accuracy?. *Journal of Behavioral Decision Making,* Volume 12, pp. 37-53.

Goodwin, P. & Wright, G., 1993. Improving judgmental time series forecasting: A review of the guidance provided by research.. *International Journal of Forecasting,* Volume 9, pp. 147-161.

Green, K. C. & Armstrong, J. S., 2007. Structured analogies for forecasting. *International Journal of Forecasting,* 23(3), pp. 365-376.

Ha, C., Seok, H. & Ok, C., 2018. Evaluation of forecasting methods in aggregate production planning: A Cumulative Absolute Forecast Error (CAFE). *Computers & Industrial Engineering,* Volume 118, pp. 329-339.

Harrell, F. E., 2015. *Regression Modeling Strategies.* Second Edition ed. New York: Springer.

Harvey, N. & Harries, C., 2004. Effects of judges' forecasts on their later combination of forecasts for the same outcomes.. *International Journal of Forecasting,* Volume 20, pp. 391-409.

Heerkens, J. M. G. & Van Winden, A., 2012. *Geen probleem, een aanpak voor alle bedrijfskundige vragen en mysteries.* Buren: Business School Nederland.

Holt, C. E., 1957. Forecasting seasonals and trends by exponentially weighted averages. *Office of Naval Research,* Issue 52.

Huh, W. T. & Roundy, R. O., 2006. A general strategic capacity planning model under demand uncertainty. *Naval Research Logistics,* 53(2), pp. 137-150.

Hyndman, R. & Athanasopoulos, G., 2018. *Forecasting: principles and practice.* Melbourne: OTexts.

Hyndman, R. J., 2013. *The difference between prediction intervals and confidence intervals.* [Online] Available at: https://robjhyndman.com/hyndsight/intervals/ [Accessed April 2020].

Hyndman, R. J. & Billah, B., 2003. Unmasking the Theta method. *International Journal of Forecasting,* 19(2), pp. 287-290.

Hyndman, R. J. & Koehler, A. B., 2006. Another look at measures of forecast accuracy. *International Journal of Forecasting,* 22(4), pp. 679-688.

Hyndman, R. J., Koehler, A., Ord, J. & Snyder, R., 2008. *Forecasting with Exponential Smoothing The State Space Approach.* s.l.:Springer.

Company A, 2020. *Seed technology.* [Online] Available at: https://www.Company A.com/en-gb/technologies/ [Accessed 01 04 2020].

Karabuk, S. & Wu, S. D., 2003. Coordinating strategic capacity planning in the semiconductor industry. *Operations Research,* 51(5), pp. 839-849.

Klayman, J., 1988. On the how and why (not) of learning from outcomes. *Human judgment: The SJT view,* pp. 115-156.

Lawrence, M., Goodwin, P., O'Connor, M. & Önkal, D., 2006. Judgmental forecasting: A review of progress over the last 25 years. *International Journal of Forecasting,* Volume 22, pp. 493-518.

Levis, A. A. & Papageorgiou, L. G., 2004. A hierarchical solution approach for multi-site capacity planning under uncertainty in the pharmaceutical industry. *Computers and Chemical Engineering,* 28(5), pp. 707-725.

Levis, A. A. P. L. L., 2004. A hierarchical solution approach for multi-site capacity planning under uncertainty in the pharmaceutical industry.. *Computers and Checmical Engineering,* 28(5), pp. 707-725.

Makridakis, S., 1993. Accuracy measures: theoretical and practical concerns. *International Journal of Forecasting,* 9(4), pp. 111-153.

Makridakis, S., Andersen, A., Carbone, R. & Fildes, R., 1982. The accuracy of extrapolation (time series) methods: Results of a forecasting competition. *Journal of Forecasting,* 1(2), pp. 111-153.

Makridakis, S., Spiliotis, E. & Assimakopoulos, V., 2020. The M4 Competition: 100,000 time series and 61 forecasting methods. *International Journal of Forecasting,* 36(1), pp. 54-74.

Martínez-Costa, C. et al., 2014. A review of mathematical programming models for strategic capacity planning in manufacturing. *International Journal of Production Economics,* Volume 153, pp. 66-85.

Martínez-Costa, C., Mas-Machuca, M., Benedito, E. & Corominas, A., 2014. A review of mathematical programming models for strategic capacity planning in manufacturing. *International Journal of Production Economics,* Volume 153, pp. 66-85.

Montero-Manso, P., Athanasopoulos, G. & Hyndman, R. J., 2020. FFORMA: feature-based forecast model averaging. *International Journal of Forecasting,* 36(1), pp. 86-92.

O'Connor, M., Remus, W. & Griggs, K., 1993. Judgmental forecasting in times of change. *International Journal of Forecasting,* Volume 9, pp. 163-172.

O'Connor, M., Remus, W. & Griggs, K., 1997. Going up-going down: How good are people at forecasting trends and changes in trends?. *Journal of Forecasting,* 16(3), pp. 165-176.

Peidro, D., Mula, J., Poler, R. & Verdegay, J., 2009. Fuzzy optimization for supply chain planning under suppply, demand and process uncertainties. *Fuzzy Sets and Systems,* 160(18), pp. 2640-2657.

Rajasekharan, M. & Peters, B. A., 2000. Strategic configuration of flexible electronics assembly facilities facing stochastic requirements. *International Journal of Production Research,* 38(3), pp. 639-656.

Randall, D. M. & Wolff, J. A., 1994. The time interval in the intention-behaviour relationship: Meta-analysis. *British Journal of Social Psychology,* 33(4), pp. 405-418.

Rastogi, A. P. et al., 2011. Supply network capacity planning for semiconductor manufacturing with uncertain demand and correlation in demand considerations. *International Journal of Production Economics,* 134(2), pp. 322-332.

Rowe, G. & Wright, G., 1999. The Delphi technique as a forecasting tool: issues and analysis. *International Journal of Forecasting,* Volume 15, pp. 353-375.

Sanders, N. R. & Ritzman, L. P., 1995. Bringing judgment into combination forecasts.. *Journal of Operations Management,* Volume 13, pp. 311-321.

Scherbakov, M., Brebels, A. & Shcherbakova, N., 2013. A survey of forecast error measures. *World Applied Sciences Journal,* Volume 24, pp. 171-176.

Slack, N. & Lewis, M., 2011. *Operations Strategy Third Edition.* Essex: Pearson Education Limited.

Syntetos, A. A., Boylan, J. E. & Croston, J. D., 2005. On the categorization of demand patterns. *Journal of the Operational Research Society,* 56(5), pp. 495-503.

Theil, H., 1971. *Applied economic forecasting.* Amsterdam: North-Holland Publishing Company.

Wang, S. M., Chen, J. C. & Wang, K. J., 2007. Resource portfolio planning planning of make-to-stock products using a constraint programming-based genetic algorithm. *Omega,* 35(2), pp. 237-246.

Winters, P. R., 1960. Forecasting sales by exponentially weighted moving averages. *Management Science,* Volume 6, pp. 324-342.

Yaniv, I., 2004. Receiving other people's advice: Influence and benefit. *Organization Behavior and Human Decision Processes,* Volume 93, pp. 1-13.

8 APPENDIX: MODEL IMPLEMENTATION CHOICES

To implement the designed model, the right software tools must be chosen. This choice is critical for both the ability to implement the model as designed, as well as for the integration in Company A. In this section, the decision criteria for the choice and the decisions for each model part are discussed.

8.1.1 Decision criteria

To decide between the available tools for each model part, the following criteria are used, in order of descending importance.

1. **Completeness**: the ability to implement all components and complexities of the model.
2. **Ease of use**: the ease by which the components and complexities of the model can be implemented in both time to implement and difficulty to learn.
3. **Availability**: the availability of the tool, documentation, and people that can work with it. A higher availability helps with implementation as well as integration in Company A. For example, open source has a higher availability than most commercial software.
4. **Visibility**: the ability to view and interact with the model and its inputs and outputs. For example, a Graphical User Interface and debugging tools.
5. **Performance**: the time needed for the tool to generate outputs from inputs.

Each criteria can be scored as either poor, sufficient, or good. Poor means a tool does not meet minimum requirements for this model. Sufficient means a tool meets the minimum requirements. Good means a tool provides additional benefits above the minimum requirements.

8.1.2 Transformation from sales data to capacity requirements

The most widely used tool for transformation is Excel. However, due to the data size and complexity of data transformation, Excel is not practical. ETL tools are commonly used for these type of data transformation tasks. ETL is the process of taking data from a source system (Extract), converting into the desired format (Transform), and stored in a data warehouse (Load). Pentaho Kettle is a widely used and open source ETL tool. Disadvantage is that this is made for transformation of databases, thus is less flexible in available functions. A third option is to use a more generic language, such as R or Python. These high-level and open source programming languages have many packages that help with data transformation. Tidyverse in R is the most widely used package for data transformation, offering a wide range of high performance tools. The table below summarizes the scores, from which can be concluded that Tidyverse in R is the best tool to use.

Tool / Criteria	Completeness	Ease of use	Availability	Visibility	Performance
Kettle (Pentaho)	Sufficient	Good	Sufficient	Good	Good
Excel (VBA)	Poor	Poor	Good	Poor	Poor
Tidyverse (R)	Good	Good	Sufficient	Sufficient	Good

Figure 8-1. Decision matrix for transformation model.

8.1.3 Forecasting capacity requirements

Basic forecasts can be produced in Excel using VBA, however these need to be built from the ground up and the calculation times are quite long. In literature, forecasting models are most commonly implemented in R. Some exceptions are C++, Matlab, and Python. C++ is low-level, therefore the most difficult to implement. The main advantage of R is that it is purpose-built for data analysis and statistics, with many package available that enable an easier implementation of forecasting methods. The most well-known are forecast and fable, of which the first is more complete and the second is more modern, as it integrates better with Tidyverse. However, since completeness is crucial in this research, the forecast package is chosen.

Tool / Criteria	Completeness	Ease of use	Availability	Visibility	Performance
Forecast (R)	Good	Good	Sufficient	Sufficient	Good
Fable (R)	Sufficient	Good	Sufficient	Sufficient	Good
Excel (VBA)	Poor	Sufficient	Good	Poor	Sufficient

Figure 8-2. Decision matrix for forecasting model.

8.1.4 Optimization of capacity planning

Mathematical programming and optimization can be implemented in a variety of software and languages. Algebraic modeling languages have been created that describe and solve complex and large scale optimization problems. Most commonly used languages are AIMMS, AMPL, and GAMS. The disadvantage is that these are commercial software. Another option is to use a more generic language that has built-in packages for optimization problems, such as R and Python. COIN-OR, a scientific organization that provides open source computational tools, created Symphony, for which packages are available for R and Python. Disadvantage of these packages is a poor visibility of the model itself; the entire model is one large matrix. OpenSolver, also from COIN-OR, is an add-in for Excel, which is visible and easy to use. Disadvantage is performance. Since this MILP model is not very large, performance is sufficient. Thus, OpenSolver is chosen.

Tool / Criteria	Completeness	Ease of use	Availability	Visibility	Performance
RSymphony (R)	Sufficient	Sufficient	Sufficient	Poor	Good
OpenSolver (Excel)	Sufficient	Good	Good	Good	Sufficient
AIMMS (Commercial)	Good	Sufficient	Poor	Sufficient	Good

Figure 8-3. Decision matrix for optimization model.