INTRODUCING A DYNAMIC ALLOCATION METHOD TO A RENTAL SYSTEM WITH UNCERTAIN RENTAL DURATIONS AND DIFFERENT CUSTOMER CLASSES

MASTER THESIS

E.C. TER HORST

16/11/2022

UNIVERSITY OF TWENTE.

Master Thesis

University of Twente

Faculty of behavioural, management and social sciences

Industrial Engineering and Management

Introducing a dynamic allocation method to a rental system with uncertain rental durations and different customer classes

Erik ter Horst

Supervisors University of Twente:

Dr. D.R.J. Prak

Fabian Akkerman, MSc.

Supervisor SST Software:

Anniek Braham, MSc.

Date of publishing: 16-11-2022

MANAGEMENT SUMMARY

INTRODUCTION

This research is focused on a company that develops and exploits locker walls for the purpose of parcel distribution. The locker walls that are exploited by the company are located throughout the Netherlands and are used by multiple delivery companies. The delivery companies can deliver parcels in the locker wall that are picked up by customers. This is called the last-mile parcel stream. Furthermore, customers can use the wall to deliver a parcel that must be picked up by delivery companies, this is called the first-mile parcel stream. The lockers in the locker wall are allocated using a First-Come First-Served (FCFS) policy. This means that if a locker is free at the moment a locker is requested, the request is always accepted. When parcels that should be put in the wall do not fit anymore and are brought to another service point, delivery companies must make an extra stop and customers cannot pick up their parcel at their desired location. These two reasons combined raise the question if the FCFS policy is the most suitable or if there are other models that would work better.

CHARACTERISTICS OF PARCEL DISTRIBUTION USING LOCKER WALLS

The actor we focus on is the locker wall company. The most important relationships we focus on are with the customers and the delivery companies. As stated earlier, there are two different parcels streams: first-mile and last-mile. The last-mile parcels are delivered by the delivery company and the first-mile parcels are picked up by the delivery company. Furthermore, the last-mile parcels are picked up by the customer and the first-mile parcels are delivered by the customer. The delivery companies mainly arrive in the morning between 09:00 and 13:00. When they arrive, they first empty the first-mile parcels dedicated to them and then they deliver the last-mile parcels.

PROPOSED SOLUTION APPROACH

The problem is formed into a small problem instance to analyse possible scenarios. We assumed the supply and demand are balanced. Furthermore, two delivery companies arrive at a fixed time (in the morning). They always arrive at the same time and in the same order (10:00 and 12:00). Furthermore, the first-mile deliveries follow a non-homogeneous Poisson distribution with $\lambda_{fmd}(i)$ that varies per hour *i*. The last-mile parcels are also picked up with a non-homogeneous Poisson rate $\lambda_{lmp}(i)$. Based on literature, three policies were devised.

Threshold policy

We defined a simple threshold policy for our dynamic rental system problem. The threshold policy only considers the current number of parcels in the locker wall. The policy does not distinguish between different types of parcels in the wall.

Myopic allocation policy

The myopic allocation policy allow customers based on a probability that is calculated based on the $\lambda_{lmp}(i)$. The probability must be given as a parameter. Based on the time of day, the mean value function until the next delivery company arrival is calculated by the following formula:

$$m(t) = \int_0^t \lambda(u)du$$

Subsequently, the probability is calculated that enough lockers will be available tomorrow if the new customer is accepted using the following formula:

$$P(T_1 > t) = P(N(t) = 0) = e^{-m(t)} \quad t > 0$$

The customer is accepted if the probability is equal to or higher than the desired probability.

Markov Decision Process policy

Finally, a policy was created by solving a MDP with a value iteration algorithm. The state of the MDP looks as follows:

$$S(t) = [LM(t), FM_1(t), FM_2(t), time\_left]$$

The state keeps track of the number of last-mile parcels in the wall, the number of first-mile parcels in the wall per delivery company and the time that is left until the next delivery company arrives. The latter is kept track of to cope with the inhomogeneous Poisson arrival rates of the last-mile pickup and first-mile delivery streams. The transition probabilities are based on the $\lambda_{lmp}(i)$ and $\lambda_{fmd}(i)$ corresponding to the time of day. In the end, every hour a decision is made on how many first-mile parcels to accept of delivery company 2.

RESULTS

The myopic allocation policy only outperforms the current situation in crowded situations and if the income per parcel for delivery company 1 is higher than delivery company 2. In the figure below, the income per policy (five different settings for the myopic allocation policy) are shown. It can be seen that the First-Come First-Serve policy is outperformed by all other policies and the best setting would be a probability of 0.70.

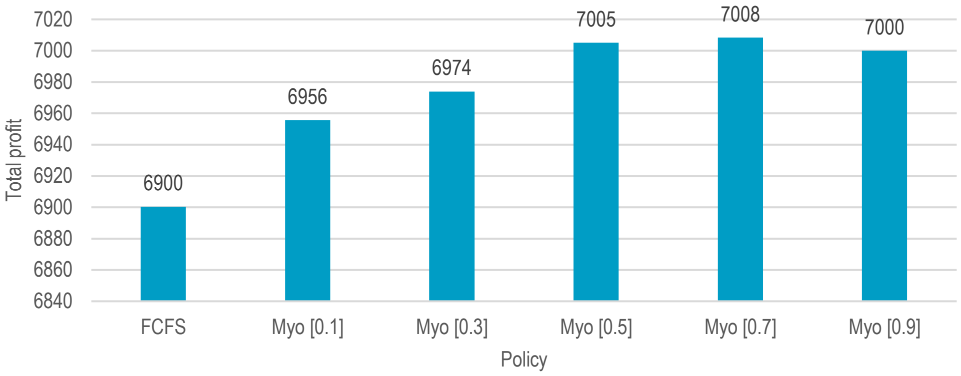

The MDP policy only outperformed the FCFS policy in extreme crowded situations. In the figure below these situations are shown. However, in the second case almost only no parcels of delivery company 2 were accepted.

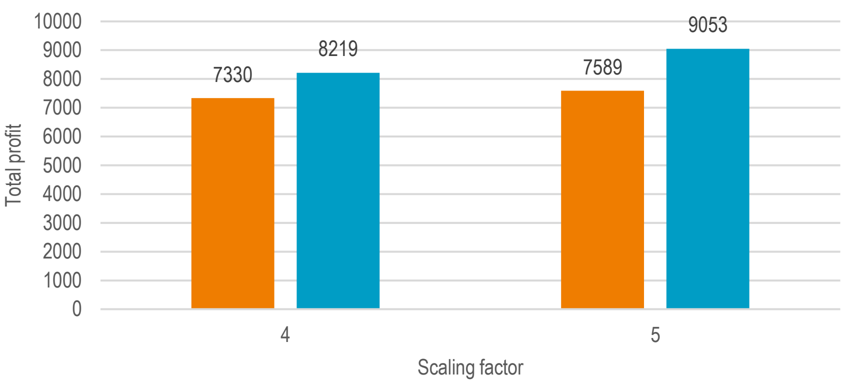

RECOMMENDATIONS

The myopic allocation policies can be implemented in crowded situations in which one delivery company should get priority compared to the other company. Based on the crowdedness and the importance of the first company, the strictness of the policy can be adjusted. In really crowded situations or with a large difference in income, a higher probability will lead to better results. The same holds if the difference in income is minimal and the penalty (in money or goodwill) of not having enough lockers available for the first company is high.

The Markov Decision Process policy that was developed did not outperform the FCFS policy in normal situations. However, when the situation got even more crowded, it started to perform better relative to the FCFS policy. In the most crowded situation, no more first-mile parcels of delivery company 2 were accepted. This turned out to work quite well for the first delivery company. Therefore, it could be interesting to close off the possibility to deliver first-mile parcels for other delivery companies in really crowded walls while still allowing last-mile parcel delivery.

PREFACE

With the finalisation of my master thesis, my time as an Industrial Engineering and Management student at the University of Twente comes to an end. I look back on six amazing years, in which I have met a lot of friends, got to challenge myself by going and living abroad and developed myself on a personal and professional level. I am very grateful to everyone who contributed to this. Furthermore, I would like to thank several individuals who contributed greatly to my master thesis.

First of all I would like to thank Anniek Braham of SST Software for her help and the many sparring conversations we have had. It was very interesting to get your view and advise on my thesis subject. Furthermore, I would like to thank everyone from SST Software for the great atmosphere and all the conversations during the lunch breaks and at the coffee machine. I would also like to thank Ralph, for the opportunity to work together, the time he made available for me and the many good and fun conversations.

I would also like to thank Dennis for being my first supervisor. From the moment I sent you the first mail with my research proposal you have been very enthusiastic and helpful. You always responded quickly and made time to give input when I was stuck. Furthermore, I would like to thank Fabian for being my second supervisor. Your knowledge of Markov Decision Processes and general comments have increased the level of my thesis a lot.

I hope you enjoy reading my thesis.

Erik ter Horst

Enschede, November 2022

TABLE OF CONTENT

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

In this chapter, an introduction into this thesis will be given. First, the company and context will be introduced in Section 1.1. Subsequently, we will describe the motivation behind the research in Section 1.2. Thereafter, we will describe the problem in Section 1.3. The research objective will be presented in Section 1.4 after which the research questions will be elaborated upon in Section 1.5. We will conclude this chapter by describing the research design in Section 1.6.

1.1 RESEARCH CONTEXT

This research is focused on a company that develops and exploits locker walls for the purpose of parcel distribution. The company owns walls throughout the Netherlands. The parcel distribution market has been growing rapidly over the past decade (eMarketer, 2017). This entails an increase in the demand for solutions like this company is offering. It is easier for the parcel distributors (they only have to drive to one destination instead of people's homes), it is easier for the customers because they do not have to stay home to wait for their parcel, it reduces carbon emissions and it increases the liveability of residential areas as the number of vans driving past homes will decrease. The demand for the locker spaces is increasing, but it is not possible nor cost efficient to keep placing as many locker walls as possible. From this, the wish of the company arises. Namely, using the existing lockers as efficiently as possible.

The locker wall company, the location partner or the delivery company can own locker walls. In the first two cases, the locker wall can be exploited openly or (partially) closed. In the former case, all lockers may be used by anyone. In the latter case, the locker wall is exclusively used by the delivery company and its customers. Rohmer & Gendron (2020) distinguish between exclusively owned locker walls and locker walls that are exploited openly. For the second category, they observe several challenges on strategic, tactical and operational level. The management on strategic level includes decisions about the location and configuration of the locker walls. Furthermore, the business model that is used is also a strategic choice. On a tactical level, the vehicle routing problems of the delivery companies are discussed. On operational level, the choice must be made which locker to assign to which customer and for what price.

1.2 RESEARCH MOTIVATION

The locker walls that are exploited by the company are located throughout the Netherlands and are used by multiple delivery companies. The delivery companies can deliver parcels in the locker wall that are picked up by customers. This is called the last-mile parcel stream. Furthermore, customers can use the wall to deliver a parcel that must be picked up by delivery companies, this is called the first-mile parcel stream. The lockers in the locker wall are allocated using a First-Come First-Served (FCFS) policy. This means that if a locker is free at the moment a locker is requested, the request is always accepted. A locker wall consists of lockers in different sizes.

Some delivery companies send information about the parcels that they will deliver to a certain parcel locker beforehand and some companies do not. This makes it difficult to estimate how busy a wall will be and if all requests in a day can be accepted. Delivery companies receive an overview in the morning with the number of first-mile parcels they have to pick up and the number of free lockers in total. The combination of these two gives a rough estimate on how much capacity will be available to them. However, in the meantime another delivery company may arrive or multiple first-mile parcels that are meant for another delivery company may be put in the wall. This may lead to the situation in which a deliverer cannot put all the parcels that are meant for the wall in a locker. The parcels that do not fit in the wall have to be taken with the deliverer again and are delivered to the nearest service point.

When parcels that should be put in the wall do not fit anymore and are brought to another service point, delivery companies must make an extra stop and customers cannot pick up their parcel at their desired location. These two reasons combined raise the question if the FCFS policy is the most suitable or if there are other models that would work better.

1.3 PROBLEM DESCRIPTION

A problem cluster that describes the situation is presented in Figure 1. The combination of missing information from the parcel companies and the unknown occupation duration of the lockers results in the fact that it is impossible to allocate the lockers before drivers arrive. Furthermore, there is no real-time availability overview. Therefore, a driver that wishes to use a locker wall always has to go there based on information that he or she received in the morning. Furthermore, the duration before the customer picks up its last-mile parcel cannot be influenced (or marginally). The complete registration of the parcels is done by some delivery companies but is often incorrect and for other delivery companies it is not possible yet. A real-time availability overview may not lead to improvement, as the routes of the delivery company are already made in the morning and cannot be adjusted when they are driving. Therefore, we will focus on the allocation model of the lockers, as this is the only core problem in Figure 1 that can be influenced and may lead to improvements.

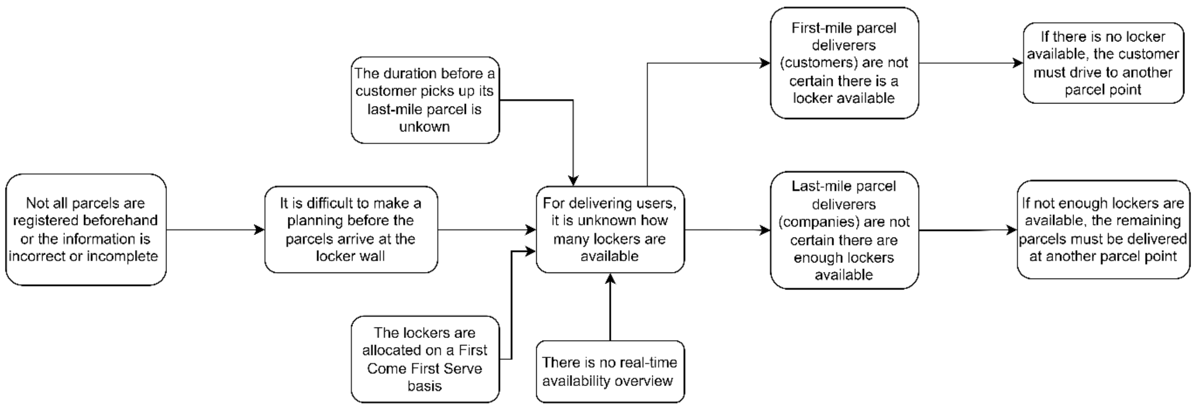

Figure 1 - Problem cluster of the locker wall problem

We formulate the core problem of this research as follows:

The allocation of the capacity and pricing of the locker walls is not done in a state-dependent way. There is no method based on historical data or demand information that dynamically allocates the capacity.

1.4 RESEARCH OBJECTIVE

As the popularity of the locker walls is growing and the capacity is limited, the company aims to optimize the utilization of the lockers. However, they are not certain how this can be realized. The most important performance indicators are the total profit and the satisfaction of the customers. The total profit will be defined to obtain the optimal utilization of the wall in different settings. The total profit will be calculated by multiplying the number of parcels with the income per parcel per unit of time. Furthermore, assumptions will be made to model the customer satisfaction as a service level. A certain service level must be obtained for a certain class of customers to be satisfied. There is not much literature available concerning allocation policies in locker walls. However, a parallel can be drawn with other rental systems with limited capacity and uncertain rental durations. In this branch of literature, admission policies and dynamic pricing policies are often proposed. However, these have never been analysed in a locker wall context. Therefore, the contribution of this research will be twofold.

First, we will investigate if there are policies that will optimize the usage of the lockers. We will look for various policies in literature about comparable rental systems and see if the policies can be applied to this context. We will analyse the different policies on total profit and will consider service levels for customers with whom explicit agreements are made. The total profit will be calculated using a prespecified income per parcel we will analyse the effect of different service level agreements on the admission and dynamic pricing policies. Therefore, the theoretical contribution will be the extension of existing literature to a new environment, an environment that has rapidly expanded over the last few years and will probably grow even further in the next decade. Investigating how an admission and dynamic pricing policy will behave in these situations will be valuable for future research in this direction.

Second, we will analyse how such a model would work in the context of the company where the research is executed. We will analyse their current situation and investigate how the policies that are found will behave in their situation. Based on this, we will give recommendations on how they can adjust the admission policy of their locker walls. This will enable them to optimally allocate the capacity of their locker walls to their different customers. As the e-commerce market will only grow over the next decade, it will only become more important to allocate the lockers in an intelligent way.

1.5 RESEARCH QUESTIONS

First, we investigated whether we could develop a dynamic method to operate the locker walls. We distinguished between contract and walk-in customers. The contract customers are the delivery companies. They send their demand information beforehand and are guaranteed a certain service level. When they wish to use more lockers, they can make a request and the system will provide an offer based on the state of the locker wall on that moment. The state of the locker wall depends on the number of lockers that are still free. Combining this with the core problem that we found in Section 1.3 gives us our main research question:

What would be the effect on the utilization of the lockers, the service level of the different classes of customers and the total profit if a state-dependent allocation policy would be implemented?

To be able to answer our main research question, we have divided our research and will answer multiple smaller research questions. The research questions will be formulated in the following paragraph.

Research question 1: What does the locker wall environment look like?

a. *Where are the locker walls located?*
b. *What are the characteristics of a locker wall?*
c. *Which actors play important roles and how do they interact with each other?*
d. *When are the first-mile parcels picked up and delivered?*
e. *When are the last-mile parcels picked up and delivered?*
f. *How does the utilization of a locker wall behave during the day (and week)?*

Research question 2: What literature exists concerning locker walls for parcel distribution and which relevant policies are described?

a. *What is written in the literature about the strategic, tactical and operational challenges of locker walls?*
b. *What relevant admission policies are described in literature concerning rental systems?*

Research question 3: Which methods will be used to model the locker wall environments and evaluate the policies that are found?

a. *How can the performance of the policies that are found be evaluated in a locker wall environment?*
b. *What should be the input of the models and what should be the experimental factors?*
c. *What should be the range of the experimental factors?*

Research question 4: What does the outcome of the performance analysis of the models mean for the locker wall company?

a. *What effect would implementing the new-found policy have on the utilization of the locker walls?*
b. *What effect would implementing the new-found policy have on the service levels of the customers?*

Research question 5: What effect would relaxing certain assumptions have on the locker wall utilization and service levels?

a. *What would be the effect of not accepting first-mile parcels anymore?*
b. *What would be the effect of not accepting last-mile parcels anymore?*
c. *What would be the effect of having more different delivery companies?*
d. *What would be the effect of spreading the delivery companies more throughout the day?*
e. *From which utilization should expansion be considered?*

1.6 RESEARCH DESIGN

In this research, we have aimed to answer all the research questions that are discussed in Section 1.5. The action problem in this thesis concerns the lack of an intelligent allocation model. However, before a solution for this was found, multiple knowledge problems had to be solved. To solve the knowledge problems, different methods have been used dependent on where the information could be found (Heerkens & Van Winden, 2017, pp. 21-25).

The first research questions and all its sub-questions are answered in Chapter 2. We will answer them by analysing historical data of the company. The first questions (locations, configuration) can be answered by looking at the data and talking with the company's director. Based on the data, an overview is made of the locker walls' characteristics and the most important actors and their interactions. Furthermore, the data is analysed to investigate the arrival patterns of first-mile and last-mile deliveries and pickups. The logs of the locker walls are used to get an overview which parts of the day are the busiest for all parcel streams. Furthermore, the logs of the locker walls are used to analyse the utilization of a locker wall (number of parcels in a locker wall per unit of time). For this analysis, five representative locker walls are chosen.

The second research question is answered in Chapter 3. We will answer the questions by investigating relevant literature. We started by looking for literature about locker walls for parcel distribution. We have looked for literature that elaborates on the strategical, tactical and operational challenges of locker walls for parcel distribution. As this area is still new, not much literature is available. Therefore, we extended our research by looking for literature about rental systems in general. Furthermore, we have searched for admission policies and dynamic pricing policies that we can use in our situation. By the end of the literature research, more is known about the challenges of the locker wall business and we have obtained policies that can be applied to our context (either directly or slightly adjusted).

The third research question is answered by looking at previous research. If the problem can be solved exactly, we may choose to use a mathematical model. If the problem environment becomes too complex, we may opt to use a simulation model. A combination is also possible. We may isolate a part of the problem to solve exactly and subsequently evaluate its performance by implementing it in a simulation model. This will be further discussed in Chapter 4 and 5.

The fourth and fifth research questions are answered by using the method that is found by answering research question 3. The method can be used to run experiments that will give answer to all the sub-questions of research questions 4 and 5. The setup for these experiments will be presented in Chapter 6. The outcome of this quantitative analysis will be the final result of this research. The results will be presented in Chapter 7. Subsequently, we will present some analyses for practical use in Chapter 8. Thereafter, we will draw conclusions and discuss the results in Chapter 9.

2. CURRENT SITUATION

In this section, the current situation will be discussed. In Section 2.1, the general characteristics of the locker walls are described. Furthermore, all actors are described in Section 2.2. In Section 2.3, an analysis on the basis of the delivery and pick-up times is presented. Finally, an analysis of the utilization of a locker wall is presented in Section 2.4.

2.1 LOCKER WALLS AND LOCATIONS

A regular locker wall consists of lockers in different sizes. There are various versions of the locker walls, depending on the physical restrictions and expected demand. The walls can be open for everyone. This means that all parcel delivery companies may use them but also normal customers. There are some exceptions to this. A (part of a) wall may be reserved for the location partner. A location partner is the partner where the locker wall is located. Furthermore, a locker wall can be reserved for one delivery company. In this situation, other delivery companies cannot use the wall for last-mile deliveries and last-mile pick-ups. The wall can be used by customers that want to send a first-mile parcel that must be transported by the specific delivery company.

Some delivery companies that use the locker wall already know in the morning how many parcels they will bring to a locker wall that day. They also have information about the size and weight of the parcels because this is measured during their sorting process. This information is sent to the locker wall company and the parcels are then registered in the system. When the size of the parcel is registered, the correct locker size is opened when the deliverer scans the parcel. He can also manually choose to open a bigger or smaller locker when the sizes are different or the parcel does not fit in the locker that was assigned. The parcels of the delivery company are registered in the system, but no reservations are made. This way, it is known which parcels will be delivered to speed up the process at the locker wall. However, no lockers are reserved and therefore the possibility exists that not all parcels will fit. Furthermore, the costs of all lockers are equal. However, when a private customer uses the locker wall, there is a difference in costs between the different sizes.

Especially in the western part of the Netherlands, the number of walls is relatively high. Therefore, there are situations in which customers live at walking distance from multiple walls. However, it is not possible to use a wall as back-up for another wall. If a wall has reached its maximum capacity, the parcel is delivered at a service point. This leads to the situation in which the deliverers have to drive to more locations than necessary. Furthermore, customers have to pick up their parcel at a location they did not pick, which decreases customer satisfaction. Additionally, an extra destination means the delivery van has to drive a bigger distance. This will lead to more emission and increases traffic in residential areas, which is evidently undesirable. Finally, not having enough space for all parcels will lead to a loss of income for the locker wall company, as the parcels will be delivered elsewhere and will never return.

2.2 ACTOR ANALYSIS

In this section, the different actors that are involved will be described. First, the company that exploits the locker walls will be discussed in Section 2.2.1. Subsequently, the customers that use the locker wall are described in Section 2.2.2. The delivery companies are described in Section 2.2.3. Finally, the location partners will be discussed in Section 2.2.4.

2.2.1 The company operating the locker walls

The company that exploits the locker walls, operates a network of locker walls throughout the Netherlands. Delivery companies and customers can use the lockers. The company that exploits the locker walls used to only work with one delivery company. For this reason, it was easy to guarantee capacity. Furthermore, the distributor had the same data as the company of the lockers. However, the market is growing and more distributors recognize the potential of the locker walls. This leads to more distributors wanting to collaborate and thus more parcels to distribute. The throughput time of parcels varies. However, a locker can be used multiple times on a day. If lockers are filled with a parcel in the morning and this parcel is picked up at noon, the locker becomes available again. It is also possible someone picks up a parcel and switches it for a parcel he or she wishes to return. Once the distributor picks up this parcel, the locker can be used again. Therefore, it is clear that a locker can be used multiple times per day. Allocation is currently first-come first-served and no real-time information is kept or shared.

Every morning, the company sends a report to the biggest delivery company that uses the locker walls. In this report, there is an overview of the parcels that are signed up by the parcel deliverer. Furthermore, the capacity that is currently available is given. This consists of the number of lockers that are available, lockers that are filled with a first-mile parcel destined for them and parcels that must be picked up again because the time limit has passed. Based on this, the expected over- or undercapacity is given. This way, the possibility is given to the delivery company to estimate whether it is desirable to visit the locker wall or not. However, in the meantime a customer can return a parcel or another deliverer can bring parcels. The overview only considers the delivery company in question and does not include the advance demand information that may have been sent by other companies. Therefore, the information is mostly correct when the delivery company arrives first, but this is mostly uncertain.

2.2.2 Customers

In this context, the customers are defined as the customers that use the wall for last-mile pickups or first-mile deliveries: people who order a parcel online and choose a locker wall as delivery option from the last-mile delivery pickup customers. They receive an email when their parcel is delivered. The email contains a QR-code with which the right locker can be opened at the wall. Depending on the deliverer, the customers have a number of days to pick up the parcel. Normally, this period consists of five or seven days. Customers who use the wall to send a parcel are seen as the first-mile delivery customers.

2.2.3 Delivery companies

The company works with different parcel deliverers. The agreements differ between the deliverers. All parties pay for a certain occupation duration. This means the parcel can stay in the locker for at most the number of days that is agreed upon before the parcel deliverer has to pick it up again.

For distributors, we believe that guarantees can be important because they otherwise take a risk by taking a detour with a full delivery van to a locker wall. When all the lockers are full, driving to the wall is useless. Too many of these useless trips may make collaborating with the locker wall exploiter less attractive. Therefore, it is important they are certain they can deliver (a big part of) their parcels in the lockers. All the parcels that do not fit in the wall are taken to the closest service point. As the parcels are delivered one by one, it can only be observed whether the wall is full after the first parcels have been delivered. Therefore, the part of the parcels that fit will always be delivered. However, this also depends on the size of the lockers that are still available. A deliverer cannot see how many lockers are still available. When he scans multiple large parcels (because those are on top of the pile) and none of them fit, he may draw the conclusion that the locker wall is full.

When the same customer orders multiple parcels, they have to be delivered in different lockers. Currently, the software is not developed to support multiple parcels per locker. There are some developments aimed to make this possible in the future.

The parcel flow depends on the season. For example, the busiest months are December and January. In December, a lot of customers order gifts online. In January, a lot of customers send gifts back. Furthermore, some locations are used more than others. The possibility of expanding locker walls depends on a number of factors. First of all, it is sometimes physically impossible. When the existing locker wall is stationed between two objects, it may not be possible to add another module to it. Furthermore, the location partner must agree to the expansion.

2.2.4 Location partners

Not all location partner agreements are the same. Locker walls can be owned by the location partner, by the delivery company or by the company that exploits the locker walls. In the first case, the location partner purchases the wall and pays a fee for the software license and maintenance. The location partner can choose to use the lockers (partly) for themselves. This way, other parties can only use these lockers if it is certain the location partner will not. In the second case, the company that exploits the locker walls is still owner of the locker wall. Depending on the agreement with the location partner, the wall can be used for open exploitation or for (partly) closed exploitation. The latter entails that only the location partner may use these lockers. Finally, a delivery company can own the locker wall. In this case, only that specific delivery company and its customers are allowed to use the wall. The costs for exploiting the wall consist of a software license fee, a maintenance fee and internet costs.

Because the locker walls that are exploited openly are more interesting from an operations management perspective, the choice is made to focus on locker walls that can be used by everyone. Additionally, the exploitation of the exclusively owned walls is more straightforward, as there are fewer different customers and thus there is less uncertainty. Furthermore, we will initially assume all lockers in a wall are open for everyone. So there are no lockers reserved for location partners.

2.2.5 Actor relations

The four important actors that were described in the previous sections all have a certain relation with each other. In Figure 2, the different relations between the actors are shown. The numbers in the arrows correspond with the relations that are described on the right-hand side.

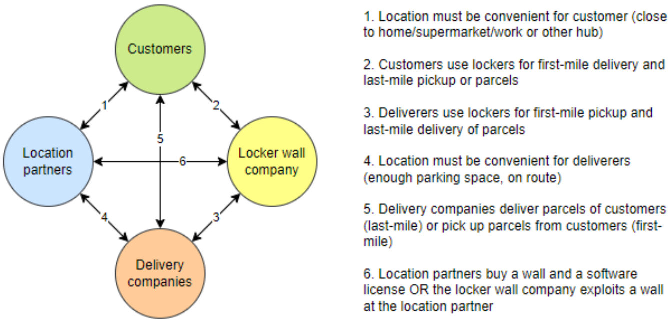

Figure 2 - Relations between the four actors

2.3 ANALYSIS OF DELIVERY AND PICKUP TIMES

The activities at the wall can be divided in two main different categories: first-mile and last-mile delivery. Other usages of the wall also exist but are left outside the scope of this research. In the first category, a customer brings his parcel to the locker wall to be picked up. In the latter category, a delivery company brings a parcel to the wall and the customer picks it up later. To get more insight in the behaviour of customers, data from five different locker walls throughout the Netherlands is analysed. The five walls are all entirely openly exploited and may be used by all delivery companies. Data from 1/1/2021 until 31/5/2022 is used. Three different delivery companies deliver the last-mile parcels and pick up the first-mile parcels. In Section 2.3.1, the first-mile delivery stream will be analysed. In Section 2.3.2, the last-mile delivery will be analysed. Finally, in Section 2.3.3, the differences and similarities between the first-mile and last-mile delivery will be discussed.

2.3.1 First-mile delivery

In first-mile delivery, customers bring a parcel to the locker wall. The parcel is usually picked up by a delivery company that ensures the parcel reaches its destination. In Figure 3, the distribution of the delivery time of parcels on four different locations in the Netherlands is shown per hour. The first parcels are returned in the morning between 06:00 and 09:00. The biggest part of the parcels are delivered between 13:00 and 17:00.

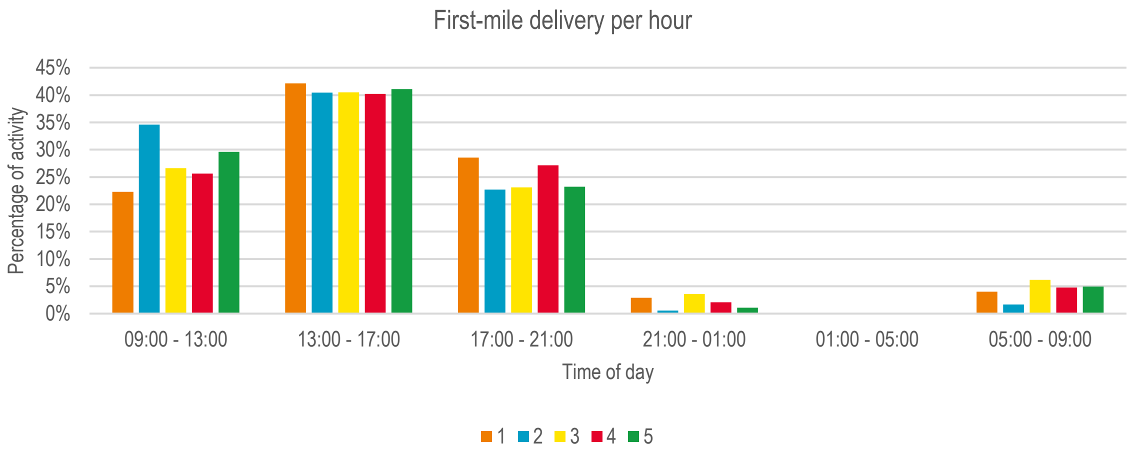

Figure 3 - Distribution per hour of the first-mile delivery of parcels to five different locker walls in the Netherlands

In Figure 4, the distribution of the pickup time of the first-mile parcels are shown. The delivery companies take the parcels. The big difference with Figure 3 is the grouping of the data. It can be seen that the biggest part of the parcels are picked up between 9:00 and 13:00 and a smaller part between 13:00 and 17:00. This may be explained by the fact that some delivery companies guarantee they will deliver a parcel the next day if it is delivered to the locker wall before a certain time. Therefore, only picking up the parcels that are delivered in the morning is not sufficient.

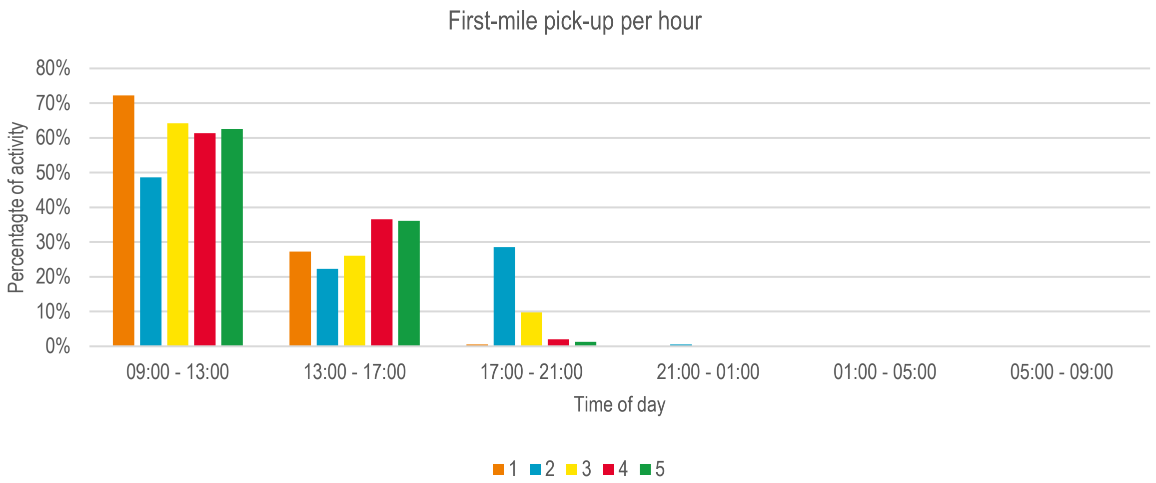

Figure 4 - Distribution per hour of the first-mile pickup of parcels from four different locations in the Netherlands

In Figure 5, the distribution of the pickup time is shown. This is defined as the time between the moment a customer delivers his or her parcel to a locker and the moment it is picked up by the delivery company. As can be seen, the biggest part of the parcels is picked up between half a day and a day (between 12 and 24 hours).

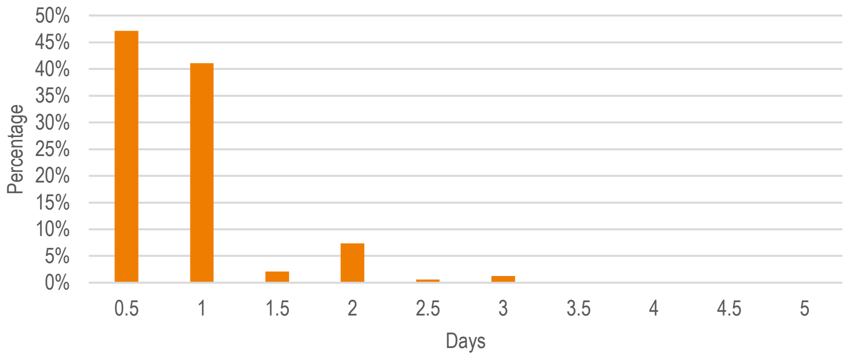

Figure 5 - The distribution of the first-mile pickup time of all delivery companies combined

2.3.2 Last-mile delivery

In last-mile delivery, the delivery companies bring the parcels to the locker wall where they are being picked up by the customers. In Figure 6, the distribution per hour of the last-mile delivery parcels is shown. It can be seen that almost all parcels are delivered between 9:00 and 13:00. Before delivering the parcels, the deliverer must empty all lockers in which a first-mile parcel lays that must be picked up by him.

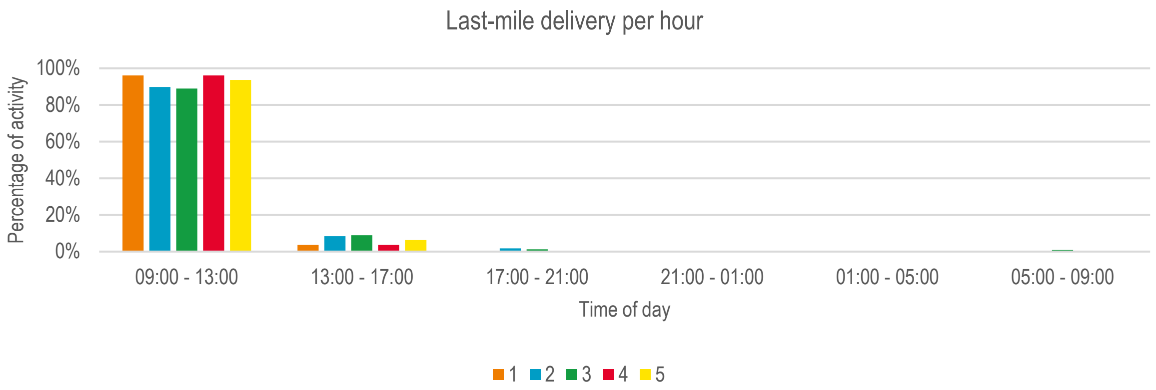

Figure 6 - Distribution per hour of the last-mile delivery of parcels to five different locker walls in the Netherlands

In Figure 7, the distribution per hour of the pickup of the last-mile parcels is shown. Customers arrive more spread out during the day. Most customers arrive between 13:00 and 17:00. Almost all customers arrive between 09:00 and 21:00.

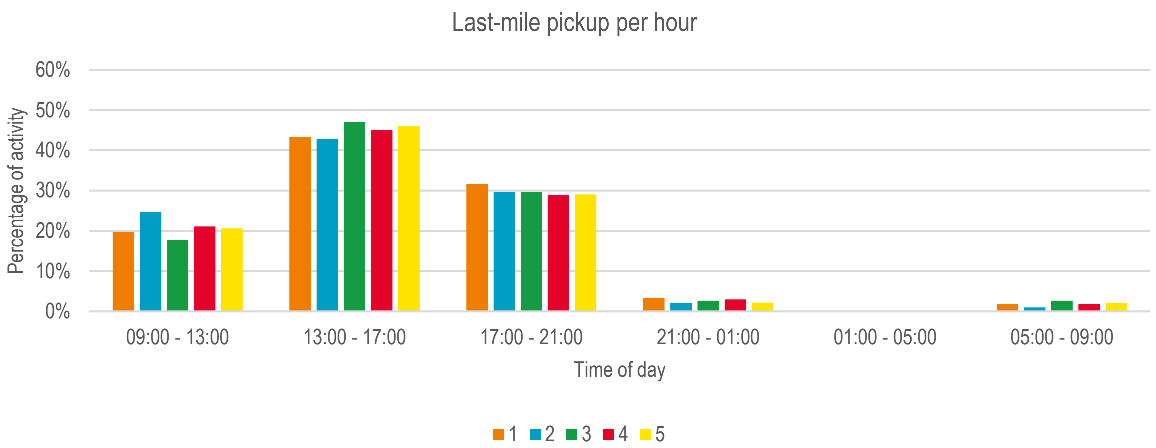

Figure 7 - Distribution per hour of the last-mile pickup of parcels from four different locations in the Netherlands

In Figure 8, the distributions of the parcel pick-up time in the last-mile are shown graphically. Almost all data points fall within the first day (around 70 percent). Theoretically, the tail on the right-hand side can infinitely long. However, the delivery company is not allowed to leave the parcel in a locker for more than a week. If the customer did not pick up its parcel by then, the locker must be emptied again by the deliverer.

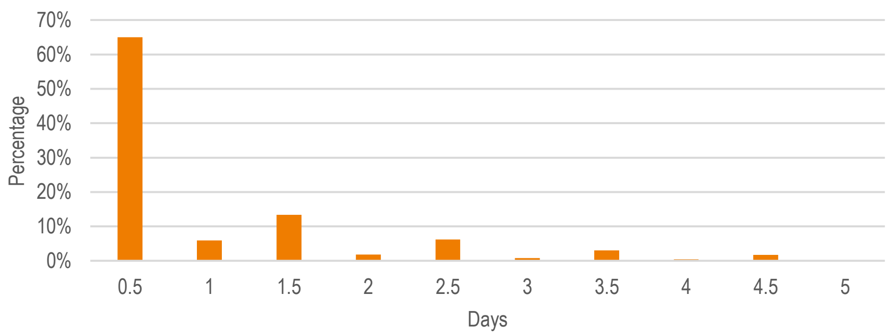

Figure 8 - The distributions of the last-mile customer pickup time of all delivery companies combined

2.3.3 Differences and similarities in first-mile and last-mile delivery

As we saw in the previous sections, there are four main different parcel streams. First-mile delivery, first-mile pickup, last-mile delivery and last-mile pickup. The customer does first-mile delivery and last-mile pickup. The first concerns individuals sending parcels which are taken by delivery companies and the second stream concerns customers picking up parcels that are delivered by delivery companies. Furthermore, the delivery companies do first-mile pickup and last-mile delivery. They either pick up the parcels that are left in the locker wall by customers or leave parcels in the locker wall that are ordered by customers. Because these streams are linked to the same stakeholder, the distribution of pickup and delivery times show similarities. This can also be seen in Figure 9 and Figure 10. Namely, in Figure 9, the last-mile pickup and the first-mile delivery stream activity are plotted next to each other. It is clear that there is almost no difference between the two of them. This is fairly easy to justify, as the behavioural pattern of the customer does not depend on whether he or she delivers or picks up a parcel.

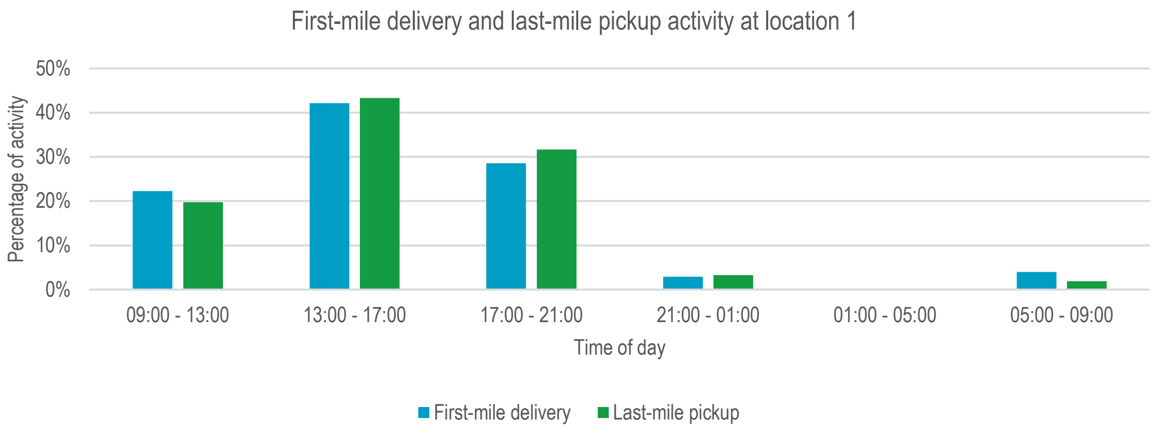

Figure 9 - An overview of the activity in the last-mile pickup and first-mile delivery streams during the day on one location

Furthermore, in Figure 10, the last-mile delivery and first-mile pickup stream activity are plotted. Between these two there is a small difference. In the morning, they are similar. Almost all activity is registered between 9:00 and 13:00. However, in the first-mile pickup stream there is also some activity in the afternoon. As stated earlier, this is explained by the fact that some delivery companies guarantee that a parcel that is delivered before a certain time (first-mile delivery) will be picked up on the same day. However, as can be seen in the graph, almost no parcels are delivered in the afternoon.

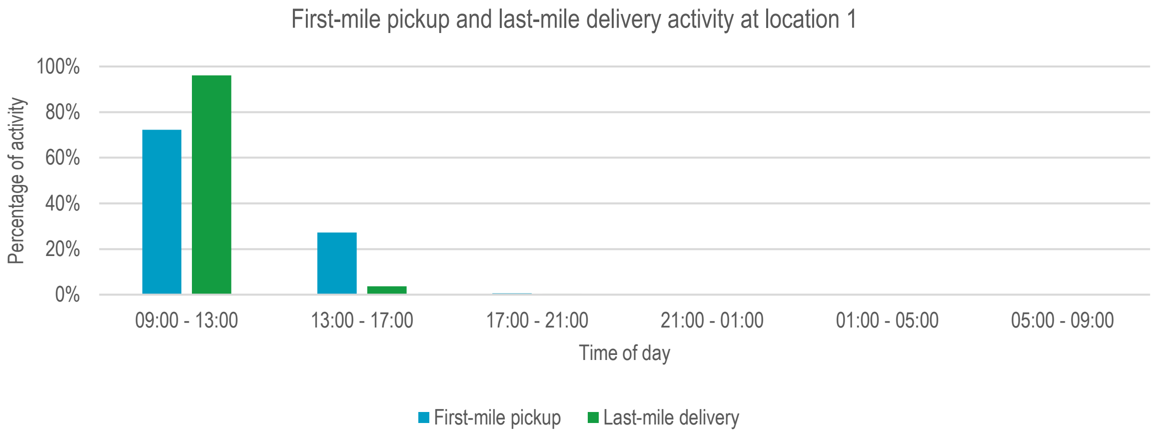

Figure 10 - An overview of the activity in the last-mile delivery and first-mile pickup streams during the day on one location

2.4 ANALYSIS OF THE UTILIZATION OF A LOCKER WALL

In Figure 11, an overview is given of the utilization of a locker wall at one of the five locations that are used. The number of occupied lockers is calculated by taking the starting volume and adding 1 if a locker is filled and subtracting 1 if a locker is emptied. By plotting the full number of lockers against the time, the graph is obtained. Three normal weekdays (Thursday, Friday, Wednesday) are plotted on top of each other. Each day starts at 08:30 am and ends at 10:00 pm. The capacity of the locker wall is 58 and the first day is 21/04/2022, which is a Thursday. The pattern that is observed is the same every day. The delivery company arrives in the morning to collect the first-mile parcels that are destined for him from the locker wall and fill it again with last-mile parcels. Generally, some delivery companies arrive again in the afternoon to collect first-mile parcels.

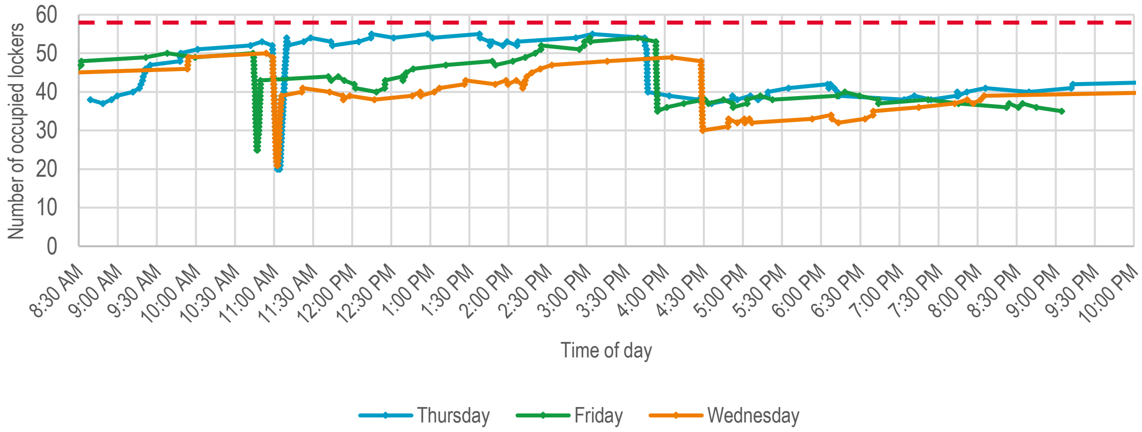

Figure 11 - An overview of the utilization of a locker wall during three weekdays in 2022

In Figure 12, the number of occupied lockers can be seen during the weekend. The first day – 23/04/2022 – is a Saturday (blue line). It can be seen that the delivery company arrives in the morning and collects the first-mile parcels. Subsequently, the last-mile parcels are put into the locker wall. After that, the line increases steadily during the day, which means more first-mile parcels are delivered by customers than last-mile parcels picked up. On Sunday (green line), no delivery companies visit the locker wall. On this day, the line increases until the full capacity is reached at 09:30 pm. Next day on Monday (orange line), the first activity is registered when the delivery company arrives and collects all the first-mile parcels. Subsequently, he delivers a number of last-mile parcels and then the pattern repeats itself.

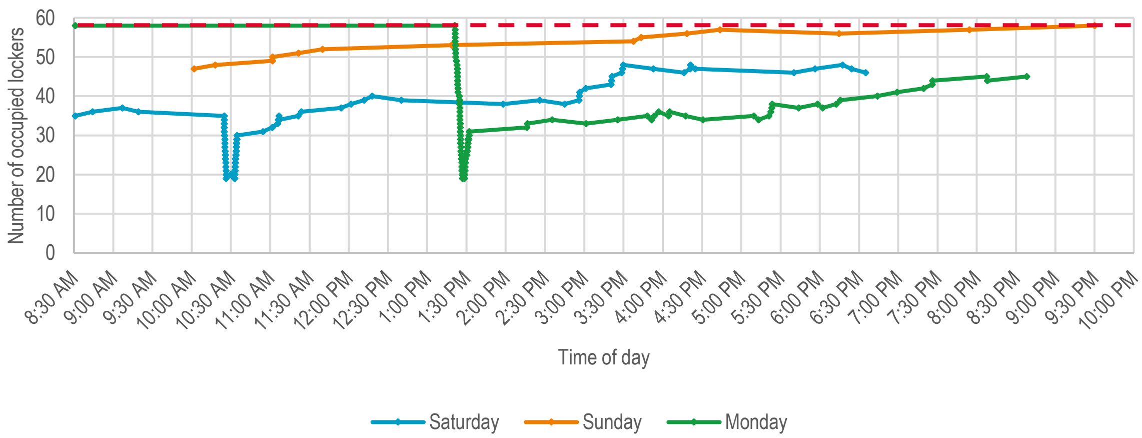

Figure 12 - An overview of the utilization of a locker wall during a weekend in 2022

3. LITERATURE

Not much research is done in the field of capacity planning for locker walls. The literature study therefore focuses on locker walls in general and pricing and capacity planning in rental systems with uncertain demand and uncertain rental durations. In Section 3.1, the difficulties in last-mile delivery and the potential effect of locker walls is discussed. In Section 3.2, literature on rental systems and accessory strategical decisions is presented. Subsequently, literature on capacity planning in rental systems with stochastic demand and rental durations with multiple customer classes are discussed in Section 3.3. Thereafter, we will discuss literature concerning dynamic pricing in comparable environments in Section 3.4.

3.1 LOCKER WALLS

In this section, we will introduce the characteristics of locker walls used for parcel distribution. In Section 3.1.1, we will discuss literature that gives a general introduction of the locker wall environment. We will then discuss literature concerning the customer's perspective in Section 3.1.2.

3.1.1 General introduction

Nowadays, the internet is available for almost everyone around the world. Consequentially, the number of digital buyers increases every year (eMarketer, 2017). In the beginning, mostly small parcels were ordered digitally. However, today everything can be ordered on the internet. Delivering all these goods to the customer requires a large and well-operating logistics network. A lot of parcels that must be delivered in the same area can be transported in bulk for a large part. However, in the end everyone desires to receive their parcel at a prespecified location. This last bit of the delivery chain is very complex. Furthermore, e-commerce shops are very concerned with the experience of their customers. Because they want to be seen as a convenient shop, they want to offer many delivery locations with time windows that are as small as possible (Macioszek, 2018). Furthermore, delivery at home can also be difficult because the delivery guy may have difficulties finding the right address or the receiving party may not be at home (Deutsch & Golany, 2018).

Rohmer & Gendron (2020) distinguish between two different types of locker networks. They can be owned by the carrier company or not. In the first case, the carrier often exclusively uses the walls. In the latter case, the lockers are exploited by a company that owns a locker wall network. The lockers are available for everyone. The introduction of locker walls poses numerous challenges. On a strategic level, decisions must be made concerning the number and locations of the locker walls, the configuration of the locker walls and the business model. On a tactical level, vehicle routing problems must be adjusted to include the hubs. The possibility that a wall is full and the deliverer still has to visit another location is one of the situations that makes the planning of the routes more difficult. Finally, on operational level the lockers must be assigned to the parcels. When a request is made, the decision must be made if and which locker is assigned and for which price. This depends on the size of the locker and the parcel. Pricing may be used to incentivize users to use the lockers in off-peak hours, so the demand and capacity are more in line. (Rohmer & Gendron, 2020) Furthermore, other specific restrictions may exist. An example are the locker walls stationed at pharmacies in which the products often must be stored in refrigerated lockers.

Rohmer & Gendron (2020) continue by discussing the pickup process. Namely, after a locker is assigned to a parcel, the customer must pick it up. The customer usually receives a pickup window. If the customer does not pick up the parcel within this time, the carrier takes the parcel with him again. They observe that more research should be done to see how customer behaviour can be influenced. For example, if a customer should pay a penalty if he fails to pick up his parcel in the pickup window. A short time window may incentivize customers to pick up a parcel faster and this increases the capacity of the locker walls. However, it may also result in a higher rate of failed deliveries. This increases costs and decreases customer satisfaction. Liu, Wang, & Susilo (2019) analyse the travel patterns of customers to and from collection-delivery points. They observed that not much research has been done in this field despite the fact that it is one of the most important parts of the pickup point delivery chain. Knowing when and how the customers will arrive to pick up or deliver their parcel is very valuable information that can be used in allocation policies. To optimally utilize the capacity of the locker walls, the planning and scheduling of the lockers should be investigated on a tactical level.

3.1.2 Customer's perspective

Conventionally, a parcel deliverer and customer are constrained by opening times at traditional pickup points. Locker walls offer the possibility to arrive at any time and deliver or pick up your parcel. Furthermore, compared to delivery to the door, customers do not have to stay home anymore, the number of failed deliveries decreases dramatically and the delivery costs are much lower. The latter can be explained by the fact that the deliverer can drive to one location instead of multiple houses. This dramatically decreases carbon emissions and improves liveability in residential areas (Davydenko & Hopman, 2020). Furthermore, Yuen, Wang, Ma, & Wong (2019) observed that customers of last-mile deliveries in China use the wall because they think the walls add value (according to the perceived value theory this means that customers will always choose the option that maximizes their utility). Furthermore, they state that they use the locker walls for parcel deliveries because they think it is dependable and convenient. Two other factors that were included in the research, privacy security and transactions costs, did not play a big part in the customers' intention to use the locker walls. Molin, Kosicki & Van Duin (2022) investigated the incentives that work best for Dutch customers to use locker walls for parcel delivery. They compared three different types of delivery: home delivery, service point pickup locations and locker walls. Their comparison entailed three attributes: price, delivery moment and the distance to the pickup location. One of the conclusions is that if home delivery prices increase and parcel lockers are available for free on a small distance, the majority of the customers would go for the last option. However, they argue that this can only be achieved if there exists a dense parcel locker network and it is important that the locker walls have a white label. This means that any delivery company can use it. Furthermore, they indicate that (local) governments should assist in facilitating the placement of white label locker walls. Iwan, Kijewska & Lemke (2016) also underline the key role local governments play in the development of last-mile deliveries in cities. They argue that municipalities too often focus on access restrictions instead of efficiency measures. When some half empty vans may enter residential areas and others not, some customers may not receive their parcels. Therefore, it would be better to focus on the efficiency.

Vakulenko, Hellström & Hjort (2018) executed a similar research in Sweden. The location was also one of the most important points according to the customers. This means the locker wall should either be located close to home or work or at the route between home and work. Furthermore, long opening hours (or even 24/7) are expected and seen as a positive aspect. Some of the negative points were the lack of track and trace information, unclear locations or service failures in combination with a lack of support. Another interesting point is the proposition that from the customer's perspective, value is not created instantly but evolves as he uses the locker wall more often. For example, when one uses the locker wall and has a positive experience, he will most likely return the next time he orders something. However, when a customer goes to the locker wall and no locker is available, he may never return as he does not want to take the risk the next time.

3.2 STRATEGIC CHOICES IN RENTAL SYSTEMS

Savin, Cohen, Gans & Katalan (2005) describe rental systems as a system in which the supply party (the rental company) invests in equipment with potential demand. The equipment can be used over a longer period of time and the rental periods are relatively short. The rental company wishes to utilize their rental capacity as efficiently as possible to earn the best return on their investment.

Lazov (2017) investigates the simplest version of a rental system. He investigates a location of a car rental company with one type of cars and one type of customers. He models the system as a Birth-Death process with homogeneous arrival and service rates. The company has M cars available, which can be seen as the number of available servers in a queueing system.

Biesinger, Hu, Stubenschrott, Ritzinger & Prandtstetter (2017) describe an instance of an electric car-sharing system in an urban area. Important strategical decisions are which number of hubs to place, where they should be located, the number of charging slots per hub and how many cars to purchase. Angelopoulos, Gavalas, Konstantopoulos, Kypriadis, & Pantziou (2016) research a comparable rental system: a bicycle sharing system in the centre of Athens. They develop a mathematical model for the determination of the locations of the stations. Furthermore, they show what the optimal number of bikes is and how they should initially be distributed. They aim to maximize the customers' utility. This is done by preventing the situation in which an individual wishes to rent a bike but the station is empty as much as possible. On the other hand, they aim to minimize the investment costs. They implement a mixed linear programming model to solve the capacitated facility location problem. In this problem, there are a finite number of possible facility locations and the aim is to find the optimal locations to cover all demand in the region. Yang, Lin, & Chang (2010) investigate a comparable situation. They analyse the situation in which an individual takes a bike from a vehicle station and drives it to another vehicle station, from which he departs to his destination. The included service level is twofold. First, customers that want to take a bike want to not be disappointed when they arrive at a vehicle station. However, when the customers arrive at their destination they wish to return the bike to a station. When all spots are full, the customer must drive to the next station, which is presumably further from his destination.

A rental system may consist of one type of rental units. However, most of the time there are different classes of rental units that are defined by some characteristic (size, luxury, age etc.). Customers that rent a certain rental unit class are often also satisfied with an upgrade if that specific class is not available. This is called an upgrade policy and is used to keep customers satisfied. When this phenomenon applies, it makes it easier for rental companies, as the demand of the group of customers of which the desired rental class is not available can be satisfied with an upgrade (Wu, Hartman, & Wilson, 2005). This can be translated to our context by assuming parcels may always be delivered in bigger lockers but it is physically impossible to deliver big parcels in smaller lockers.

From all different rental systems we obtain a number of strategical decisions that belong to the exploitation of rental systems. Firstly, the locations of the stations are important. Furthermore, the choice should be made how many rental units to deploy and how to distribute them over the rental stations.

3.3 ALLOCATION AND ADMISSION POLICIES OF RENTAL SYSTEMS WITH STOCHASTIC DEMAND AND DIFFERENT CUSTOMER CLASSES

The stream of literature that is described by Savin et al. (2005) relates to the more traditional revenue-management problems like plane seat reservation. However, there are two main differences. First, the plane seat allocation problem has a finite horizon: from the moment it is possible to book until the plane departs. As soon as the booking is open, everyone can book seats (including travelling organisations). But as soon as the plane leaves, the planning problem ends. Second, all seats can be used only once. Therefore, the allocation of the rental units is seen as a continuous-time infinite-horizon problem. The arrival pattern of customers is stochastic as is the rental period.

The growing pressure on rental systems incentivizes the development of new business models. Customers can be divided into different classes based on revenue, service level requirements or sales volume. Subsequently, each customer class can be treated differently to optimize capacity utilization and thereby revenue. For each class, a service level and a fixed fee can be agreed upon. Based on this agreement, the choice must be made to either accept or reject a rental request (Papier & Thonemann, 2010b). The trade-off focuses on choosing between a direct income and going to a state with one less rental unit or rejecting the direct income because a penalty may be incurred if no rental units are available in the near future (Jain, Moinzadeh, & Dumrongsiri, 2015).

Savin et al. (2005) show that a rental system with different classes of customers with class-dependent earnings and penalties can be reduced to a form of the stochastic knapsack problem. This method uses historical data to assume the arrival distributions of the different classes of customers. Furthermore, they assume the system behaves as a continuous-time Markov chain, only depending on the state (number of free rental units). They show that the complete sharing method can be optimal in two situations. The complete sharing method assigns its resources without considering types of customers (Wang & Pinsky, 1989). Complete sharing works well if the capacity is high relative to the demand. When this is not the case, it can still work well if the differences between the different classes is negligible. In these situations, the physical and economic utilization of the rental units are optimal. This means the revenue is as high as possible when the wall is as full as possible. However, in other scenarios, the physical and economic utilization of the wall may not behave the same.

Besides the complete sharing method, Wang & Pinsky (1989) describe four other resource allocation methods. They analyse a system with *N* identical resources that are used for different activities. The different classes of customers arrive with independent Poisson processes at the multi-server station. They wish to use a number of servers for a random period of time. The paper focuses on calculating the blocking probabilities, which is the chance that a request must be denied. The other methods they analyse are the complete partitioning method, the sharing with minimum allocation method, the complete sharing policy with an order constraint and the sharing with maximum allocation method. The complete partitioning method allocates a fixed number of resources to each customer and allocates all resources. The sharing with minimum allocation method also assigns a fixed number of dedicated resources to each customer. However, the resources that are not assigned are shared. The complete sharing policy with an order constraint does not allow a customer of type i to use more lockers than a customer of type $i+1$. Which means that some customers may not use more than other customers. The last method, sharing with maximum allocation, is the same as complete sharing but every customer may only use a maximum number of resources. Kaufman (1981) also analyses the blocking phenomenon. He analyses a system with a finite capacity of c servers. customers arrivals are modelled as a Poisson process. Customers are defined by two requirements, a spatial requirement and a temporal requirement. This means they request a number of servers for an amount of time. It is assumed all servers are the same. Customers whose demand cannot be met are blocked. This means they are not served and disappear. Blocking of customers does not further affect the system.

Altman, Jiménez & Koole (2001) study a similar situation in which they focus on a resource-sharing system. In the system, there are multiple classes of customers that all yield a certain class-dependent reward. Subsequently, they devise a dynamic program in which it is decided whether to accept or reject a certain customer. When all resources are used, the only possibility is to reject an arriving customer. Because the rewards differ per customer class, dynamic programming is used to estimate the future rewards of a certain choice. They develop an optimal admission policy and extend it by analysing a situation with three classes of customers.

Gans & Savin (2007) describe a rental system with two types of customers. The first type of customers pays a fixed fee which give them a guarantee they can have a certain amount of the rentals at any time for a fixed price. When there are not enough rentals available, the rental company pays a penalty. These are contract customers. The second group of customers are the walk-in customers. They are not guaranteed a certain service level and they receive an offer at request. The rental company does not have any obligation towards them. They develop a Markov decision process model that includes different classes of two types of customers. The value function is shown in Equation 1.

$$v(k) \;=\; \sum_{i=1}^{N} \lambda_i{}^{\varphi} H_i{}^{\varphi}[v(k)] + \sum_{j=1}^{M} \lambda_j{}^{\varpi} H_j{}^{\varpi}[v(k)] + \mu k v(k-1) + \mu(c-k)v(k) \qquad (1)$$

The *c* in Equation 1 represents the capacity of the rental system and *k* the number of lockers in use. The first summation part represents the *N* different classes of the premium customers. The arrival rate is multiplied with the expected value. Likewise, the second summation part represents the *M* walk-in customers. The third term represents the situation in which a locker becomes available again. The last term represents the possibility that nothing happens. The model is analysed at constant discrete moments. Therefore, the last term is needed to make sure the possibility nothing happens exists.

When a premium customer enters and there are lockers available ($k<c$) he can either be accepted or rejected. In the first case, the next stage is entered in the same state and a penalty is paid. In the second case, a direct reward is yielded and the next stage is entered. In this case, the state is increased with one. When there are no more lockers available ($k=c$), a penalty is paid and the next stage is entered in the same state. This is shown in Equation 2 and Equation 3, respectively.

$$H_i{}^{\varphi}[v(k)] \;=\; \max[f(k) \;-\; \pi_i, f(k+1) + r_i \text{ if } k \;<\; c \qquad (2)$$

$$H_i{}^{\varphi}[v(k)] \;=\; f(k) \;-\; \pi_i \text{ if } k \;=\; c \qquad (3)$$

When a walk-in customer enters and there are lockers available, a price ω_l is offered and customer *j* accepts this price with possibility ρ_{jl}. The price is chosen such that the value function is the highest. This is shown in Equation 4. When no lockers are available, nothing happens. This is shown in Equation 5.

$$H_j{}^{\varpi}[v(k)] \;=\; \max_l[\rho_{jl} * f(k+1) + (1-\rho_{jl}) * \; f(k)] \text{ if } k \;<\; c \qquad (4)$$

$$H_j{}^{\varpi}[v(k)] \;=\; f(k) \text{ if } k \;=\; c \qquad (5)$$

Furthermore, the aggregate event rate is given in Equation 6.

$$\Gamma = \sum_{i=1}^{N} \lambda_i{}^{\varphi} + \sum_{j=1}^{M} \lambda_j{}^{\varpi} + \mu c + \gamma \qquad (6)$$

The authors define the time scale in such a way so that $\Gamma = 1$. This means that every time step, one of the terms in 6 happens. The last term represents a discount rate, which is modelled as the probability that the next event is terminating. When this happens, the system stops and no more profits are earned. Adding a termination probability is said to be equivalent to discounting.

It is shown that threshold-based policies are optimal for the admission of contract customers. The fee that should be charged to walk-in customers should be calculated based on the congestion in the system. Furthermore, it is shown that the optimal threshold for a specific class decreases if the arrival intensity and the expected rental duration increase. Additionally, it depends on the fixed fee of other classes. The fee for walk-in customers increases when the arrival intensity, expected rental duration or fixed fee of contract customers increase. Finally, they show that higher contract fees guarantee a higher level of access.

One of the conclusions of the paper is that myopic management policies do not perform well if the demand significantly exceeds the capacity. The myopic policy always allows contract customers and a walk-in customer is offered a fee that maximizes the discounted revenue from his rental. When the demand is significantly larger than the available capacity, it becomes more important to allocate the capacity carefully. However, for systems in which the capacity and the demand are balanced, myopic allocation policies can be (near) optimal. This means decision are only based on the reward that can be earned now instead of also looking at the future.

Papier & Thonemann (2010) add to this by describing a system with two different types of customers: premium and classic. The premium customers pay a fixed fee and receive a service guarantee. When a service request is rejected, the rental company must pay a penalty. The classic customers can be seen as walk-in customers. They are not given any guarantee and can be denied without paying a penalty. The rental period is described by a random variable. In further research, they extend the system by assuming the premium customers provide their demand information beforehand (Papier & Thonemann, 2010a).

A classic customer is accepted if the revenue that can be earned is greater than the expected loss in profit of having one rental unit less in the future. In this situation, the number of premium customers in the coming time interval are known. However, the rental period is random. This means that if a classic customer is accepted, the capacity will be one less for a random period of time. This could either lead to not being able to accept a premium customer in the near future, not being able to accept a classic customer in the near future or neither of those cases. The probabilities of these states are used to calculate the expected profit of both choices. To be able to analyse the models, all demand is assumed to be homogeneous. To deal with inhomogeneous demand, they propose to implement the Average Stationary Approximation (ASA) by Whitt (1991). This method initially calculates the stationary performance of each period separately and then takes the weighted average over all periods. Papier & Thonemann (2008) use this method in their research to fleet planning. They develop and solve analytical models that can determine the size and the structure of a rental fleet. The latter is focused on the types of rental units that should be used (and in what proportions). They consider one type of rental unit but include seasonality. As the demand varies heavily per season, they show that fleet leasing can increase profit. In this scenario, the company owns a base number of rental units and lease extra units during peak demand. This outcome, however, is not applicable in our situation, as our rental units cannot be leased and the setup costs of rental units are higher. Furthermore, they distinguish between two commonly used objectives in their models. First, the profit can be optimized. However, it can be more realistic to minimize the costs using a service level constraint. This way, the penalty cost of not meeting demand (lost sales) does not have to be estimated. This may be better, as the penalty costs have a high impact on the solution and estimating it can be difficult.

Pazour & Roy (2015) investigate a vehicle rental company with two types of customers. At every arrival of a non-priority customer they must decide to rent them a car or reject them because they expect the arrival of a priority customer. They use a Markov-chain based solution approach to determine the optimal threshold. When the number of available cars is below that threshold, non-priority customers cannot rent a car. Additionally, they use a discrete-event simulation model to analyse the effect of the threshold policy on the waiting times of the different classes of customers.

Örmeci & Burnetas (2004) include another phenomenon in their description of admission control systems. Namely, the fact that rental companies often largely depend on orders from corporate clients rather than individual customers. They include this in their research by using batch arrivals to model the demand. Such systems are defined by a queueing system with a number of parallel identical servers, no waiting room and different job classes. The latter is defined by an arrival pattern and a job-specific service time. Adding to this, they consider a system in which batches can also be accepted partially (Örmeci & Burnetas, 2005). Furthermore, they consider dynamic admission control policies instead of static admission control policies. Dynamic admission control policies can increase profit compared to static admission control policies because the decisions are based on the current state of the system.

Satır, Erenay & Bookbinder (2018) describe a vehicle capacity allocation problem. For every vehicle, they have to decide which goods to accept and which to reject. They distinguish between two types of orders: expedited and standard. They model their problem as a continuous-time Markov decision process. The arrival rates of the different types of orders follow a compound Poisson distribution. They each have an independent arrival rate and a discrete random variable which determines the size of the arrival order. They develop an optimal policy and define the state space of the problem as the amount of normal and expedited orders that have already been accepted.

In this section we have seen that a lot of research has been done into rental systems with different customer classes. Most research is rooted in the wish of rental companies to use their rental units more efficiently. By differentiating in customer classes, one can justify treating groups of customers differently. For example, some pay a higher fee for a higher service level and some do not have expectations and try to get the best deal at a certain moment in time. Most literature focuses on determining optimal (threshold) admission policies. For this, dynamic programming and Markov decision chains are used. They are used to determine what action would be the smartest in a certain state. These models are extended by adding batch arrivals, varying service times per customer classes and customers who send their demand information in advance.

3.4 DYNAMIC PRICING

In this section we will discuss the topic of dynamic pricing. First we will introduce dynamic pricing in a general context in Section 3.4.1. Subsequently, we will discuss literature that investigates dynamic pricing strategies in problem instances with a finite horizon in Section 3.4.2. Finally, we will discuss dynamic pricing in environments with an infinite horizon in Section 3.4.3.

3.4.1 Introduction to dynamic pricing

The topic of dynamic pricing concerns environments in which prices for products or services can be adjusted if circumstances change. The prices can be adjusted monthly, weekly, daily or even hourly. Dynamic pricing strategies are mostly seen in companies that operate digitally and have a lot of available data. From the data, important customer patterns and behaviour can be deduced (Den Boer, 2015). Exploiting this knowledge can help companies to maximize their revenue and/or balance the workload. The process in which all the sales data is used to determine the best price for a product on a certain moment is called dynamic pricing and learning. This topic concerns dynamic pricing algorithms that learn based on the data what the price should be (Den Boer, 2015). Elmaghraby & Keskinocak (2003) observe an expansion in types of companies that use dynamic pricing strategies. Initially, it was adopted by companies with a fixed capacity and a small horizon, e.g., airlines, cruises, hotels, energy providers and healthcare. However, they saw that companies with bigger capacity and longer horizons also started to introduce dynamic pricing strategies. Three main causes are discussed. First, the amount of available data rapidly increases. Second, it has become very easy to change your prices due to technological developments. Finally, more methods are available to easily process demand data and to even suggest decisions. On the other side, customers now have access to the internet from everywhere, which gives them the chance to compare multiple providers of the same product in seconds. Therefore, it is important to stay competitive as a seller and dynamic pricing strategies may help (Bitran & Caldentey, 2003).

Historically, determining a price was based on the operating costs and availability of certain products. These variables were known and therefore it was relatively easy to determine the price of a certain product. Nowadays, two other factors are included: how much the customer wants the product and the expected demand in the (near) future. The last two factors are not known and must be estimated and forecasted. There are a lot of different techniques for this and a lot of data is needed. Therefore, dynamic pricing strategies are more difficult to implement than traditional pricing strategies. Roughly, a distinction between two different dynamic pricing strategies can be made. Namely, posted-price policies, which concerns a non-negotiable price determined by the sellers and price-discovery policies in which the buyers determine the price. The latter can for example be implemented as an online auction. Furthermore, customers can also be divided in two different classes: myopic and strategic. Myopic customers will accept a price if it is beneath a certain boundary. This boundary can very per (class of) customer and does not change. The strategic customer follows the development of a price and includes this in his decision process. It is evident that dynamic pricing is harder when a customer falls in the second category. Assuming a customer behaves myopically can be correct in a number of cases. First, if a customer really needs an item and cannot longer wait. Second, when the price changes are so small a strategic customer would also not wait. Third, when there are so many customers that one rejection does not affect the system much. Finally, if it concerns an impulse buy (Elmaghraby & Keskinocak, 2003). Based on a combination of the first, second and third reason, we would argue customers of a locker wall can be seen as myopic. Within certain boundaries, customers who already decided to use a locker wall will not return if a small price change is observed. Furthermore, when this happens occasionally, the lockers will be filled with other parcels.

3.4.2 Dynamic pricing with a finite horizon

Farias & Roy (2010) analyse a dynamic pricing problem with uncertain demand, limited inventory and with the goal to maximise their discounted earnings over an infinite time horizon. Customers arrive according to a Poisson process of which the rate is uncertain. The customers make a myopic decision when they are offered a price. The customers only purchase a product if the price they are offered is equal to or below their reservation price. A reservation price is the boundary myopic customers have. They do not consider historical data or expected future prices. They develop a simple heuristic as dynamic pricing strategy which is called decay balancing. Furthermore, they extend their model by distinguishing between different products which can be offered to different classes of customers for different prices.

Gallego & Van Ryzin (1994) also focus on a problem in which the earning must be maximized. However, they consider a finite horizon problem. This type of problem is applicable to the types of business we mentioned earlier: airlines, cruises and hotel rooms. These problems are defined by posted-price strategies from the seller side and myopic customers on the other side. In the paper, a function is developed that calculates the price a certain product should cost at a given moment depending on the number of products in stock and the length of the horizon. They still saw a lot of difficulties in introducing dynamic pricing strategies as it was difficult to implement and constant price adjustments were seen as undesirable. However, as stated earlier, Elmaghraby & Keskinocak (2003) and Bitran & Caldentey (2003) explained that due to the increased amount of data, easiness in adjusting prices online and the development of decision tools this has become more popular. Therefore, introducing a state-dependent function to determine a price at any moment is not unthinkable anymore in any context.

3.4.3 Dynamic pricing with an infinite horizon

Gosavi, Bandla & Das (2002) use reinforcement learning to solve a semi-MDP over an infinite time horizon. The problem concerns a revenue management problem in the airline industry. For a single flight, they create a model that considers different classes, possible overbooking, different arrival patterns of different customer classes and possible cancellations. They observe that prices are heavily influenced by the prices that are set by competitors. However, there are two factors that can be affected. The first is seat allocation and the second is overbooking. The former concerns the concept of selling seats in the same cabin for different prices to different classes of customers. A number of seats may be saved to prevent higher class customers from not having a seat. However, it is difficult to determine the number of seats that should be saved. It is evident it depends on the class-dependent fares. Customers can be allocated to classes based on how many weeks they book in advance, their itinerary Furthermore, the overbooking problem focuses on the problem that customers sometimes do not show without cancelling their ticket. Without overbooking, the plane would leave with some empty seats and this is obviously a missed chance on revenue. However, the question arises how many seats should be overbooked to compensate for the no-shows.

The overbooking problem is less interesting for our context. However, the seat allocation problem shows resemblances with the locker wall problem. Normally, different classes of customers are made based on different characteristics and customers within the same class are treated equally. In the paper of Gosavi et al. (2002), customers arrive according to independent Poisson processes. Furthermore, the probability that a customer cancels his ticket is fixed but differs per customer class. The problem is modelled as a semi-MDP and the system state is defined by four state variables. Namely, the class of the most recent customer, a vector with the number of seats that are sold per class, a vector with the arrival times of all customers per class and the time that is left until the plane leaves. The state can change if any of these three events occur: a new customer arrives at the system requesting a ticket, a cancellation occurs or the flight departs (the time horizon is reached). The time between the customer request arrivals are seen as the time between the decisions, which is also known as the time epochs. Furthermore, the arrival rates of the customer do not change over time. This could, however, easily be implemented using reinforcement learning (Gosavi et al., 2002).

Carroll & Grimes (1995) describe a dynamic pricing strategy in the car rental business. In the paper, a yield management system is described that was used at the car rental company Hertz. In the situation that is described, customers do not always return the car after the specified period. Therefore, Hertz developed a model to estimate the probability that the car would be returned on a certain day. This would nowadays maybe not apply to the car rental business, but the system is comparable to the one we are investigating. Like airline bookings, the products are perishable. This means a product disappears if it is not sold before a specific day or time. This is often represented by a finite horizon in operation research problems. Normally, products perish after the day or time. However, in the car rental industry, cars that are not rented out on a certain day, can still be rented out the next day. Therefore, the product is perishable on a given day, but the next day it starts again. This can be seen as small sequential finite horizon problems. However, one could also choose to not rent out a car on a given day because he expects that he will generate more revenue by renting it out the next day. Extending a problem with these kind of decisions would make it an infinite horizon problem and a lot more difficult. Furthermore, it is heavily affected by the duration of the rentals as we have also seen in Section 3.3.

Another area in which dynamic pricing is considered in an infinite horizon setting is cloud computing capacity providers. Customers pay a time-varying rate for the time they use the capacity of for their calculations. Xu & Li (2012) implement a revenue management policy to solve the problem. They ultimately investigate optimal pricing strategies for the infinite horizon revenue management problem. They also observe that capacity (in this specific case computing resources) is perishable if they are not used they do not add value. However, in environments with stochastic demand pricing is difficult. As accepting customers now for a low rate could lead to not having capacity for a customer tomorrow who is willing to pay more. On the other hand, rejecting a customer now could lead to not having a customer the coming days. This phenomenon is meant by products being perishable. Furthermore, setting the price high will give you a direct high reward but may decrease the future demand because customers will not return. Moving from a finite horizon problem to an infinite horizon problem, the problem state does not include any information about the time anymore. The state is now solely defined by the utilization of the system. However, the earnings are calculated based on the time someone uses the cloud capacity. Therefore, the time a customer is in the system is modelled explicitly. The departure process is modelled as a price-dependent Poisson process. The rate increases with the price.

3.5 CONCLUSION

In this chapter, we have presented and discussed literature from four different streams. Namely, literature concerning locker walls in general, literature concerning the strategic choice in rental systems, literature about admission policies rental systems with stochasticity and dynamic pricing. The first part explained the challenges and the current situation concerning locker walls. It explained that customers are willing to use the locker walls but it highly depends on usability. Furthermore, an important role is seen for local governments to focus on solutions like this. It will reduce emissions and improve liveability of residential areas. Based on this, we can say with confidence that this is a promising area and therefore research into the exploitation of locker walls can be valuable.

The strategic choices regarding locker walls for parcel distribution boil down to a few important ones. First, the location of the locker wall should be decided. As already said in the previous section, customers value usability and the distance to a locker wall is an important factor. Furthermore, the configuration of the locker wall must be decided upon. The number of lockers that are used and in what ratio the different sizes should be chosen. In this research, we will not focus on the strategic choices regarding the locker walls. However, we will focus on finding the right moment for expansion based on the utilization in the analyses for practical use.

Given a number of locker walls, a few delivery companies and customers that wish to use the walls, another problem arises. Namely, who may use the wall given a certain state and time. The literature we presented in Section 3.3 concerns rental systems with different types of customers. The common divisor is that all rental unit owners wish to optimize the utilization and earnings of their business. By dividing the customers in different categories and estimating the arrival patterns, admission policies can be devised that optimize the profit of the companies. This can be used in the context of the locker walls. By distinguishing between different users, the lockers can be divided based on category-dependent earnings and penalties. Therefore, we will use the MDP model that was presented by Gans & Savin (2007) and apply it to our situation. Furthermore, we will see if myopic allocation policies work in our situation. We will investigate in which settings they perform best.

Finally, we have discussed literature about dynamic pricing. The admission policy that will be developed and evaluated can be extended by a dynamic pricing method. Instead of accepting or rejecting a customer when the locker wall is in a certain state, a state-dependent price can be offered (which can also depend on the type of customer that wishes to use a locker). In this research, only the admission policies will be studied. However, we will describe how they can be extended to also include dynamic pricing strategies.

4. PROBLEM FORMULATION

In this chapter, the research problem will be elaborated upon. Furthermore, it will be linked to the literature that we discussed in Chapter 3. After this chapter, the research problem will be clear to the reader and it will be known which parts of the literature can be used in this context and why. In Section 4.1, a more in-depth problem description will be provided and it will be linked to the literature. In Section 4.2, the problem settings will be discussed. In Section 4.3, the problem will be formulated as a Markov decision process.

4.1 PROBLEM DESCRIPTION

Rohmer & Gendron (2020) describe choices on the operational level of locker walls. This includes when to assign a locker to which customer and for what price. The literature that was discussed in Section 3.3 all use different customer classes to describe their different customers. The classes can be used to agree on different service levels and prices. The consequence of not having a locker available at arrival of a customer may also be class dependent . This means that the penalty may not be the same for all classes of customers. For example, rejecting a contract customer could have a bigger negative effect than rejecting a walk-in customer.

This segmentation enables a company to treat different classes of customers differently. The classes can be divided based on revenue, sales volume and service requirements. Contract or premium customers pay a fixed fee and receive a guaranteed service. Other customers see if anything is available at request and then decide whether they think it is worth it. This gives the rental company the possibility to accept or reject specific classes of customers in a certain state or to adjust the prices based on the congestion of the system.

To be able to investigate certain areas, we will have to make some assumptions to be able to create a reasonable framework in which we can do our research. For this rental system, we will assume delivery companies (at least the ones that will be considered 'premium') send all information about their parcels beforehand (e.g. in the morning). Additionally, they will pay a lump sum premium for the service. The advance demand information will be used to determine if a walk-in customer may use a locker and for what price. For this purpose, the number of occupied and expected premium parcels will be added to analyse the expected congestion in the system. Delivery companies that do not send their information beforehand will be allowed to use the wall but will be considered walk-in customers.

Additionally, the service levels that are communicated to the contract customers will be measured over the long term. So, when a service level of 0.90 is promised, a request from them should be accepted nine out of ten times on average. The walk-in customers are not guaranteed any service and can either be accepted or rejected, depending on the crowdedness in the locker wall at their arrival.

In Figure 13, a flowchart is shown that illustrates three different arrival processes . The first customer stream is called the contract customer stream. It depicts the stream of customers that is marked as premium in the literature. A customer is marked as premium if they send their demand information beforehand. On arrival, the system checks whether the customer is accepted or not. This may depend on the state of the system, the current service level of the partner and the service level that is promised to him. Furthermore, it may depend on the expectation of other customers in the near future and what their current and promised service levels are. Penalties are only paid if the service level is not met over a longer period of time. Therefore, they are not captured in this flowchart.

If non-premium customers arrive at the system, the process is different. The flowchart in the middle of Figure 13 shows the process if an admission policy is implemented at the locker wall. When a customer arrives, the system determines based on the current state whether the customer can be accepted or not. The flowchart in the bottom of Figure 13 shows the process if a dynamic pricing policy is implemented. At arrival, the system checks the current state of the locker wall. In combination with the type of customer and time of day, a price is calculated which is offered to the customer who wishes to use the wall. Subsequently, the customer can either accept or reject the offer and the state of the locker wall is updated accordingly.

Like Pazour & Roy (2015), we will use discrete-event simulation to test the policies and compare it to the current situation in which a FCFS policy is used. This way, we will evaluate what the effects of implementing the admission policies and dynamic pricing policies would be that are found with our mathematical models. We will be able to use empirical data and analyse the potential effect of the policies in a realistic environment. It also makes it easier to analyse the impact of changes that are made in certain parameters. For example, the distribution of the arrival times of the delivery companies and the number of lockers. The simulation model will be further elaborated upon in Chapter 6.

Figure 13 - A flowchart of the arrival process of a contract customer (above), the arrival process of a walk-in customer to a locker wall with an admission policy (middle) and the arrival process of a walk-in customer to a locker wall with a dynamic pricing policy (below)

4.2 PROBLEM SETTINGS

We will analyse a small problem instance to analyse possible scenarios. We will assume the supply and demand are balanced. Furthermore, we will assume the delivery companies arrive at a fixed time (in the morning). They always arrive at the same time and in the same order. The arrival rate varies per hour *i* and is based on historical data. Furthermore, the first-mile deliveries follow a non-homogeneous Poisson distribution with $\lambda_{fmd}(i)$ that varies per hour *i*. The arrival rates are based on the five locations that were also analysed in Chapter 2. The last-mile parcels are also picked up with a non-homogeneous Poisson rate $\lambda_{lmp}(i)$. The arrival rates are calculated per locker. When the arrival rates are used in a problem setting, they are multiplied with the number of lockers that are used. The pickup time of first-mile deliveries is a constant time T_{fmp}. The non-homogeneous Poisson arrival rates per hour are shown in Table 1.

Table 1 - Non-homogeneous arrival rates per hour of last-mile and first-mile deliveries scaled to one locker

From	Until	$\boldsymbol{\lambda_{lmp}(i)}$	$\boldsymbol{\lambda_{fmd}(i)}$
00:00	*01:00*	0.0002	0.0005
01:00	*02:00*	0.0001	0.0004
02:00	*03:00*	0.0001	0.0000
03:00	*04:00*	0.0000	0.0000
04:00	*05:00*	0.0000	0.0004
05:00	*06:00*	0.0001	0.0002
06:00	*07:00*	0.0002	0.0009
07:00	*08:00*	0.0009	0.0034
08:00	*09:00*	0.0043	0.0227
09:00	*10:00*	0.0058	0.0292
10:00	*11:00*	0.0070	0.0364
11:00	*12:00*	0.0170	0.0425
12:00	*13:00*	0.0278	0.0455
13:00	*14:00*	0.0312	0.0504
14:00	*15:00*	0.0283	0.0533
15:00	*16:00*	0.0294	0.0587
16:00	*17:00*	0.0358	0.0621
17:00	*18:00*	0.0345	0.0514
18:00	*19:00*	0.0216	0.0389
19:00	*20:00*	0.0161	0.0331
20:00	*21:00*	0.0084	0.0170
21:00	*22:00*	0.0044	0.0060
22:00	*23:00*	0.0018	0.0021
23:00	*00:00*	0.0012	0.0007

The decision to accept or reject a request to use a locker that must be made when customers or delivery companies arrive at a locker wall can be difficult. We will start by analysing a myopic allocation policy that only considers the direct probability that accepting a customer now will lead to not being able to accept another customer in the future. When we are going to distinguish between different classes of customers, these probabilities can be used to determine whether a customer should be accepted or rejected. Furthermore, the policy can be extended by using the probabilities to determine the dynamic price certain customers should pay based on the expected arrivals of other classes of customers. Because the service level agreements and penalties are class-dependent, the probabilities can be used to determine if accepting a customer now will lead to having to pay a certain penalty in the future. Multiplying the probabilities that are found with those penalties will give an insight into the consequences a certain choice might have. We will further define the problem instance and myopic allocation policy in Section 5.2.

4.3 MARKOV DECISION PROCESS

In this section, we will introduce an MDP formulation. We will discuss the situation and context in Section 4.3.1. Subsequently, we will introduce the state and action variable in Section 4.3.2 and 4.3.3, respectively. Furthermore, we will analyse the transition probabilities in Section 4.3.4. We will then discuss the reward function in Section 4.3.5 after which the value function will be discussed in Section 4.3.6. Finally, we will introduce a possible solving method in Section 4.3.7.

4.3.1 Case description

The Markov decision process model developed by Gans & Savin (2007), that was described in Section 3.3, will be used to analyse the problem mathematically. In the model, different classes of premium customers can be analysed with different Poisson arrival rates. Furthermore, the rental service time is exponentially distributed for all customers. The model will be adapted to the problem setting we described in Section 4.2. The assumption was made that the delivery companies arrive in the same order. In between, last-mile customers can pick up a parcel from the locker wall and first-mile customers can bring parcels to the locker wall, which are dedicated to a delivery company. The arrival patterns of the last-mile and first-mile customers are represented by inhomogeneous Poisson distributions. To be able to cope with this in a MDP formulation, a time variable will be added to the state variable.

Puterman (1990) describes the problem formulation of MDPs and the various adaptations. First, the choice must be made at which time steps the system is analysed. The time set can either be continuous or discrete and finite or infinite. We will implement an MDP with discrete time steps of 1 hour and a finite horizon. The system consists of two delivery companies, a locker wall with a fixed capacity *c* and customers that arrive to either pick up or deliver a parcel to the locker wall. Parcels are dedicated to one of the two delivery companies. A first-mile parcel dedicated to delivery company 2 cannot be picked up by delivery company 1 and the other way around. In this problem instance, it is assumed the delivery companies always arrive at the same time and in the same order. For this reason, it is always known when first-mile parcels will be picked up and when last-mile parcels will be delivered. As the time horizon goes from the moment delivery company 2 leaves until the moment delivery company 1 arrives, accepting first-mile parcels of delivery company 2 will lead to a smaller available capacity for delivery company 1. When a delivery company arrives, the first-mile parcels that are dedicated to them are emptied and the last-mile parcels that they carry with them are delivered (if accepted).

The time horizon will be 22 hours. At t=0, the second delivery company will just have left the system (12:00 at noon) and at t=22 (10:00 in the morning), the other delivery company will arrive. Because the delivery company heavily affects the state of the system, we only capture the time between the delivery companies with our MDP. At the end of the horizon of the MDP, the delivery company will arrive. All first-mile parcels dedicated to him will be collected and new last-mile parcels will be delivered. However, this is not a stationary process and therefore we choose to focus on the hours between the arrival of the two delivery companies. The parcels that are delivered by the delivery companies are either last-mile parcels from delivery company 1, first-mile parcels for delivery company 1, last-mile parcels from delivery company 2 or first-mile parcels for delivery company 2.

In this model, the number of parcels delivery company 1 will wish to deliver to the wall is known. Therefore, during the day on arrivals, decisions must be made whether to accept or reject arriving customers. Not being able to provide lockers for all last-mile parcels will lead to a penalty whereas declining first-mile parcel requests leads to a direct loss of income.

4.3.2 State variable

The state of the system is defined by three variables. The number of first-mile parcels for delivery company 1, the number of first-mile parcels for delivery company 2, the number of last-mile parcels that are in the locker wall and the time that is left until the end of the time horizon is reached. No difference is made between last-mile parcels for delivery company 1 or delivery company 2, because they are picked up with the same distribution and it does not matter for the decision or action. We have to make the state time-dependent because the arrival patterns of first-mile deliveries and last-mile pickups are represented by a non-homogeneous Poisson distribution. The Markovian property dictates that the system should only depend on its currents state and not on history (Puterman, 1990). Therefore, the changing $\lambda_{lmp}(i)$ and $\lambda_{fmd}(i)$ can only be included if the time i is included in the state. This follows the same structure as the state vector that was presented by Gosavi et al. (2002) and discussed in Section 3.4.2. Therefore, our state vector looks as follows:

$$S(t) = [LM(t), FM_1(t), FM_2(t), time\_left]$$

$LM(t)$: the number of last-mile parcels in the locker wall at time t

$FM_1(t)$: the number of first-mile parcels of delivery company 1 in the locker wall at time t

$FM_2(t)$: the number of first-mile parcels of delivery company 2 in the locker wall at time t

$time\_left$: number of hours that are left until the time horizon

We will initially analyse a locker wall with 10 lockers. Every sub-state can take any value between 0 and 10. However, their sum can never exceed the total capacity, which is equal to the total number of lockers that are available. We can calculate the number of possible states by analysing it like a statistical combination problem. We have three different types of parcels and ten spots. We can calculate the number of possible combinations with repetition, which means the number of different combinations we can make with the three different parcels if we can use each of them as many times as we wish. The formula for this is provided in Equation 7 (Taboga, 2021).

$$C_{n,k} = \binom{n+k-1}{k} = \frac{(n+k-1)!}{(n+k-1-k)!k!} = \frac{(n+k-1)!}{(n-1)!k!} \quad (7)$$

The number of different objects n and the number of spots k, are - if all spots are filled - equal to 3 and 10, respectively. Therefore, the number of possible different states amounts to 66. This calculation must be done for all possible numbers of filled lockers.

Table 2 - Number of possible combinations for all possible numbers of full lockers

n	k	$C_{n,k}$
3	10	66
3	9	55
3	8	45
3	7	36
3	6	28
3	5	21
3	4	15
3	3	10
3	2	6
3	1	3
3	0	1
	Total:	**286**

Additionally, the $time\_left$ variable can take on 23 different values. Therefore, the total number of possible states is equal to $286 * 23 = 6578$.

As said earlier, the problem horizon starts when the second delivery company has just left. The $time\_left$ variable is then equal to 22, as we assume the second delivery company leaves at 12am and the next delivery company arrives at 10am in the morning. During the time horizon, customers will arrive to pick up last-mile parcels and first-mile customers will arrive to deliver first-mile parcels. Both arrival patterns follow a Poisson arrival process. The distributions of the last-mile customers will be combined. The arrivals of first-mile parcel customers are analysed separately, as they are dedicated to a delivery company.

4.3.3 Action variable

When the agent is in a certain state, he must pick an action from a prespecified list of possible actions, picking an option triggers the transition to a next stage. Given a certain action, a probability distribution is specified on the next system state. Furthermore, a direct reward is earned based on the state and the action that was taken. The action does not only affect the direct income but also affects the evolution of the process. Eventually, the decision maker wants to make decisions that optimize the total reward over the whole time horizon (Puterman, 1990). In this situation, the decision maker only focuses on the first-mile parcels, because no decision can be made regarding last-mile parcels that are already accepted (they cannot be emptied preemptively). The decision can be made to allow 0 parcels or to allow up to the number of parcels that is equal to the number of empty lockers in the current state. However, the sum of allowed first-mile parcels for delivery company 1 and delivery company 2 must not exceed the available capacity. The action variable vector looks as follows:

$$A(t) = [\,allowed\_FM_{1,t}, allowed\_FM_{2,t}]$$

Given that:

$$allowed\_FM_{1,t} \; + \; allowed\_FM_{2,t} \leq available\_lockers$$

With:

$$available\_lockers = total\_capacity - (LM(t) + \; FM_1(t) + \; FM_2(t))$$

4.3.4 Transition probabilities

In Table 3, the arrival rates are shown that will be used for this problem. The arrival rates are the same arrival rates as presented in Table 1, but they are multiplied with a factor 10 to correct for the number of lockers. Furthermore, it is assumed the first-mile parcels of the first and second delivery company arrive with the same rate. Therefore, the arrival rates that were presented in Table 1 are multiplied with a factor of ten and then divided by two. In the left-most column, the $time\_left$ variable is shown. Furthermore, the decision will be made at the start of the hour. So, the first decision is made at $time\_left = 22$, which is at 12:00 in the morning.

Table 3 - Non-homogeneous arrival rates per hour of last-mile and first-mile deliveries scaled to ten lockers

$time\_left$	*From*	*Until*	$\boldsymbol{\lambda}_{lmp}(\boldsymbol{i})$	$\boldsymbol{\lambda}_{fmd\_dc1}(\boldsymbol{i})$	$\boldsymbol{\lambda}_{fmd\_dc2}(\boldsymbol{i})$
22	*12:00*	*13:00*	0.2779	0.2274	0.2274
21	*13:00*	*14:00*	0.3115	0.2520	0.2520
20	*14:00*	*15:00*	0.2834	0.2664	0.2664
19	*15:00*	*16:00*	0.2940	0.2933	0.2933
18	*16:00*	*17:00*	0.3581	0.3106	0.3106
17	*17:00*	*18:00*	0.3449	0.2571	0.2571
16	*18:00*	*19:00*	0.2163	0.1943	0.1943
15	*19:00*	*20:00*	0.1612	0.1653	0.1653
14	*20:00*	*21:00*	0.0844	0.0849	0.0849
13	*21:00*	*22:00*	0.0440	0.0299	0.0299
12	*22:00*	*23:00*	0.0178	0.0103	0.0103
11	*23:00*	*00:00*	0.0120	0.0033	0.0033
10	*00:00*	*01:00*	0.0024	0.0023	0.0023
9	*01:00*	*02:00*	0.0011	0.0018	0.0018
8	*02:00*	*03:00*	0.0008	0.0000	0.0000
7	*03:00*	*04:00*	0.0002	0.0000	0.0000
6	*04:00*	*05:00*	0.0000	0.0018	0.0018
5	*05:00*	*06:00*	0.0007	0.0008	0.0008
4	*06:00*	*07:00*	0.0024	0.0047	0.0047
3	*07:00*	*08:00*	0.0090	0.0170	0.0170
2	*08:00*	*09:00*	0.0426	0.1134	0.1134
1	*09:00*	*10:00*	0.0577	0.1461	0.1461

To calculate the transition probabilities, the probability mass function of the Poisson distribution is used. The probability density function is shown in Equation 8 (Kissell & Poserina, 2017).

$$P_x(k) = \frac{\lambda^k e^{-\lambda}}{k!} \tag{8}$$

The probability density function will be used to calculate the chance that a certain number of events happen in the next epoch. For example, if the state would be: $S(0) = [5, 0, 0, 22]$. The chance that the state in the next stage would be equal to $S(1) = [5, 1, 0, 21]$, would be equal to the probability that one first-mile customer of delivery company 1 has arrived multiplied with the probability that no first-mile parcels of delivery company 2 have arrived multiplied with the probability that no last-mile customer has picked up a parcel, given that the action taken allows new first-mile parcels in the system. We assume that the probabilities are independently distributed. The chances would look as follows:

$$P(one\ FM\ arrival\ of\ DC1) = P_{fmd\_dc1}(1) = \frac{0.1073^1 e^{-0.1073}}{1!} = 0.0964$$

$$P(no\ FM\ arrival\ of\ DC2) = P_{fmd\_dc2}(0) = \frac{0.1073^0 e^{-0.1073}}{0!} = 0.8983$$

$$P(no\ LM\ pickups) = P_{lmp}(0) = \frac{1.2046^0 e^{-1.2046}}{0!} = 0.2998$$

Therefore, the total probability of moving from $S(0) = [5, 0, 0, 22]$ to $S(1) = [5, 1, 0, 21]$ would be equal to:

$$P((S(0) = [5, 0, 0, 22]) \rightarrow S(1) = [5, 1, 0, 21])) = 0.0964 * 0.8983 * 0.2998 = 0.0230$$

Furthermore, when the next stage is entered it is only possible to go to a state of which the $time\_left$ variable is one less than the state in the previous stage. The transition probabilities to all other states will be equal to zero. The states in which the $time\_left$ variable is equal to zero are absorbing states. This means that the only transition probability that is bigger than 0 is towards themselves and is equal to 1.

4.3.5 Reward function

The rewards that can be earned in the system are twofold. First, a direct reward can be earned by accepting first-mile parcels. However, if there are not enough lockers available for the delivery company that arrives at t=0, a penalty will be incurred. This is represented by terminal rewards. Terminal rewards are earned if the agent ends up in a certain state (Puterman, 1990). The terminal rewards are negative in this case. It is assumed, the number of parcels the next deliverer will bring is known. Per parcel that does not fit in the wall anymore, a penalty must be paid.

As stated earlier, the direct rewards are earned by accepting first-mile parcels. The reward function is shown in Equation 9.

$$r_t(s, s') = income\_parcel * ((FM_{1,s'}(t) - FM_{1,s}(t)) + (FM_{2,s'}(t) - FM_{2,s}(t))) \quad (9)$$

The terminal rewards are dependent on the number of parcels that will be delivered by the end of the time horizon, the number of parcels that fit in the wall and the number of lockers that are occupied at the end of the time horizon. The number of lockers that are short will be multiplied with a prespecified penalty to obtain the terminal reward for that state. If the number of lockers that are needed is smaller than the available lockers, no penalty will be incurred. Therefore, a maximum operator is included, which takes the value 0 in the latter case.

$$terminal\_reward = \max(0, (lockers\_needed - available\_lockers)) * penalty$$

With:

$$available\_lockers = total\_capacity - (LM(T+1) + FM_1(T+1) + FM_2(T+1))$$

4.3.6 Value function

The value function does not only consider the reward that is earned by moving from one state to another. It optimizes the rewards earned over the whole horizon, including the terminal penalties that may be incurred. The value function is shown in Equation 10.

$$V(s) = \max_a (r(t) + E[V(s'|s, a)]) \quad (10)$$

The value function forces the system to take the actions that maximize the rewards that are obtained and the terminal penalties that are incurred. Normally, future rewards are discounted with a factor to adjust for the fact that income now is more valuable than income later. However, because the complete time horizon of the problem is only a day and the income per parcel is marginal, we choose to not use a discount factor.

4.3.7 Value iteration algorithm

We analyse our problem using the value iteration algorithm. This algorithm first initializes the values of each state to the immediate rewards of that state. Thereafter, the values are updated until the change in values is lower than ε= 0.01 or the maximum number of iterations is reached. The value of a state is updated by the value of the next best state multiplied with the probability of getting there given a certain action. In the end, the best possible action is obtained from each state. The combination of all best actions in all states is called a policy. The algorithm is shown below (Pashenkova, Rish, & Dechter, 1996):

```
Set Values = {state : R(s) for each state s}
Until values don't change or max iterations is reached or difference is smaller than ε:
    copy_values = values
    for each state s:
        initialize best_Expected_Value
        for each action:
            Expected_Value = 0
            for each next state ns:
                Expected_Value += trans_probability * copy_values[ns]
            best_Expected_Value = max (Expected_Value, best_Expected_Value)
        values[s] = R(s) + discount * best_Expected_Value
```

5. SOLUTION APPROACH

In this chapter, we will discuss different methods that are used to form a solution. First, we will discuss a simple threshold policy in Section 5.1. Thereafter, we will elaborate on and work out the myopic policy and the Markov decision process policy that were discussed in Chapter 4. First, we will present a myopic allocation policy in Section 5.2. Thereafter, we will discuss the MDP we devised in Section 5.3.

5.1 THRESHOLD POLICY

We introduced the phenomenon of threshold policies in rental systems in Section 3.3. We will now define a simple threshold policy for our dynamic rental system problem. The threshold policy will only consider the total number of parcels in the locker wall. Based on this, the decision will be made whether an arriving first-mile customer may use a locker or not. The policy will not distinguish between different types of parcels in the wall. For example, a threshold of 0.90 means no one may use a new locker anymore if 9 out of 10 lockers are filled. That is, until a last-mile parcel customer arrives and collects a parcel. This will bring the utilization down to 0.8 and will make it possible for a new customer to use the wall. In the discrete-event simulation model, we will experiment with various threshold values ranging between 0.1 and 0.9.

5.2 MYOPIC ALLOCATION POLICY

In this section, we will discuss a myopic allocation policy. The myopic allocation policy will be more complex than the threshold policy as it will be based on information about historical arrival patterns. In Section 5.2.1, we will introduce the method and discuss the theoretical background. In Section 5.2.2, we will present calculations and probabilities that lay on the basis of this policy. In Section 5.2.3, we will conclude and define the policy.

5.2.1 Introduction

Gans & Savin (2007) showed that within certain conditions, myopic policy management may be near optimal. We propose to analyse a myopic policy and analyse this mathematically. For this, we can use a small problem and use the exponential interarrival times of the different classes to analyse the possibilities that assigning a locker now will lead to loss of income and/or a penalty in the near future. We can analyse this situation for different interarrival times and different rental durations. We can also analyse the difference between exponentially distributed rental durations and rental durations that are constant. The latter could be the case in first-mile delivery pick-ups, because you almost always know the parcels will be picked up between 10:00 and 12:00 or in the afternoon.

From Huseby (2021), we obtain some useful theorems. For example, the mean value function that is shown in Equation 11.

$$m(t) \ = \ \int_0^t \lambda(u)du \tag{11}$$

Additionally, to calculate the possibility that no arrivals occur until t is given in Equation 12.

$$P(T_1 > t) = P(N(t) = 0) = e^{-m(t)} \quad t > 0 \tag{12}$$

Moreover, the density of T_1 is given in Equation 13.

$$f_{T_1}(t) \ = \ \lambda(t)e^{-m(t)} \, t > 0 \tag{13}$$

We will use these properties to analyse certain scenarios in which one locker is empty or all lockers are full. Subsequently, we will have to decide on the spot whether to accept a locker request when a locker is free. Or we can calculate the chance that a customer will arrive before any locker is free. We can experiment with the duration of the first-mile pick up deliveries (constant time) and the decisions we make in certain scenarios. Furthermore, we can calculate the probability at any given time a last-mile parcel will be picked up before the next delivery company arrives again. This can be valuable information because the virtual capacity at the delivery company's arrival then increases and another first-mile parcel may be accepted.

To calculate the probability that a last-mile parcel that is currently in the wall will be emptied, the $\lambda_{lmp}(i)$ from Table 1 are used. The $\lambda_{lmp}(i)$ are scaled per locker. However, the utilization of the locker walls in reality is approximately equal to 0.50. Furthermore, half of the parcels in the wall are first-mile parcels. To be able to use the Poisson arrival rates per individual parcel, we therefore multiply the $\lambda_{lmp}(i)$ from Table 1 with a factor 4 (times two because of the 0.50 utilization and times two because half of the utilization are first-mile parcels). Suppose the time of day is 15:00 and the next delivery company arrives at 10:00. Using the mean value function shown in Equation 11 and the $\lambda_{lmp}(i)$ from Table 1, we obtain $m(t) = 0.6596$. Inserting this in Equation 12, we obtain $P(N(t) = 0) = e^{-m(t)} = e^{-0.6596} = 0.52$. Therefore, the probability that the parcel is picked up before the next delivery company arrives is equal to $1 - 0.52 = 0.48$. This will be further explained and presented in the next section. Furthermore, when two last-mile parcels are in the wall at 15:00, the probability that at least one parcel is picked up before the next delivery company arrives again is calculated using the probabilities above and the possible scenarios. The probability that both parcels will be picked up is equal to the multiplication of both separate chances, as we assume independence. Therefore, the probability that both lockers will be free when the next delivery company arrives is equal to $0.52^2 = 0.27$, this will be later presented in Figure 15. Furthermore, the probability that at least one of the two lockers will be available in the morning consists of the sum of the following probabilities:

- The probability that parcel 1 will be collected and parcel 2 not;
- The probability that parcel 1 will not be collected and parcel 2 is;
- The probability that both parcels are collected.

The first two probabilities are equal to $0.48 * 0.52 = 0.25$. Therefore, the sum of the three probabilities that is equal to $0.25 + 0.25 + 0.27 = 0.77$. When more lockers are considered, the probabilities of all possible combinations are calculated and added. We will now further present the probabilities in Section 5.2.2.

5.2.2 Myopic policy

In this section, we will present the probabilities that will form the basis of our myopic policy. The model will only consider the current state of the wall and - based on historical data and arrival patterns - we will determine how many lockers can be given away at what time. In this problem instance, we distinguish four different arrivals. Namely, the arrivals of the first-mile delivery, first-mile pickup, last-mile delivery and last-mile pickup. Delivery company 1 always arrives at 10:00 in the morning and delivery company 2 always arrives at 12:00 in the morning. The delivery companies deliver last-mile parcels. Furthermore, the last-mile pickup customers arrive with a non-homogeneous Poisson distribution. The first-mile parcels that are delivered are picked up by the delivery company (the parcels are dedicated to delivery company 1 or 2). For clarity, the arrival distributions are shown in Table 4.

Table 4 - Arrival patterns of the four different arrivals

Arrival	How
Last-mile delivery	Fixed moment during the day
Last-mile pickup	Non-homogeneous Poisson
First-mile delivery	Non-homogeneous Poisson
First-mile pickup	Constant

First-mile parcels that are delivered and must be taken by delivery company 2, cannot be taken by delivery company 1. Because delivery company 1 always arrives before delivery company 2, the lockers that are filled with first-mile parcels for delivery company 2 cannot be used by delivery company 1. When a delivery company arrives, it first empties all the first-mile parcels that are dedicated to them and subsequently fills the wall with the last-mile parcels they have. The last-mile parcels that are delivered to the wall are picked up by customers. The customers arrive following a non-homogeneous Poisson distribution. It is assumed the number of lockers that is needed by the delivery companies is known one day beforehand. Therefore, the allocation policy can be used with a horizon of one day.

In the next paragraphs, we will analyse the probabilities that last-mile parcels in the wall are picked up before the next delivery company arrives. In Section 5.2.3, we will extend this by explaining how the probabilities can be used in an admission policy.

Probabilities last-mile parcels are picked up before the next delivery company arrives

In scenarios where first-mile customers of delivery company 2 arrive and delivery company 1 is the next arriving deliverer, accepting the customers means that the locker will not be available for delivery company 1. The lockers that can be used by delivery company 1 are the lockers that are:

- empty on that moment;
- lockers that contain first-mile parcels of delivery company 1;
- lockers that contain last-mile parcels that will be picked up before delivery company 1 arrives.

Out of these three possibilities, only one contains stochasticity. First-mile customers of delivery company 1 may always be accepted after delivery company 2 has visited, because the first-mile parcels will be taken before delivery company 2 uses the wall again. It is evident empty lockers can always be used. But in a scenario where first-mile customers of delivery company 2 arrive in the afternoon and you know (approximately) how many lockers delivery company 1 will need tomorrow, you must decide to accept or reject them. In that moment, only the last-mile parcels that are in the wall can still be emptied before the next delivery company arrives. Because the first-mile parcel that is accepted will definitely be not empty when the deliverer arrives. In this problem instance, the number of last-mile parcels in the wall makes a difference. Therefore, we analysed the scenarios that 1, 2, 3 or 4 last-mile parcels are in the locker wall. For the last scenario, the probability is calculated that at least 1, at least 2, at least 3 and at least 4 parcels are picked up before the next delivery company arrives. For the other scenarios, the same is done depending on the maximum number of last-mile parcels in the wall. The probabilities are calculated for each hour after delivery company 2 has visited at 12:00 and before delivery company 1 arrives at 10:00 in the morning.

Scenario 1: One last-mile parcel in the wall

When there is only one last-mile parcel in the wall, the probability that it is emptied before the next deliverer arrives is between 0.6 and 0.7 after deliverer 2 has left. The probability decreases steadily during the day and stays approximately constant during the night, as can be seen in Figure 14.

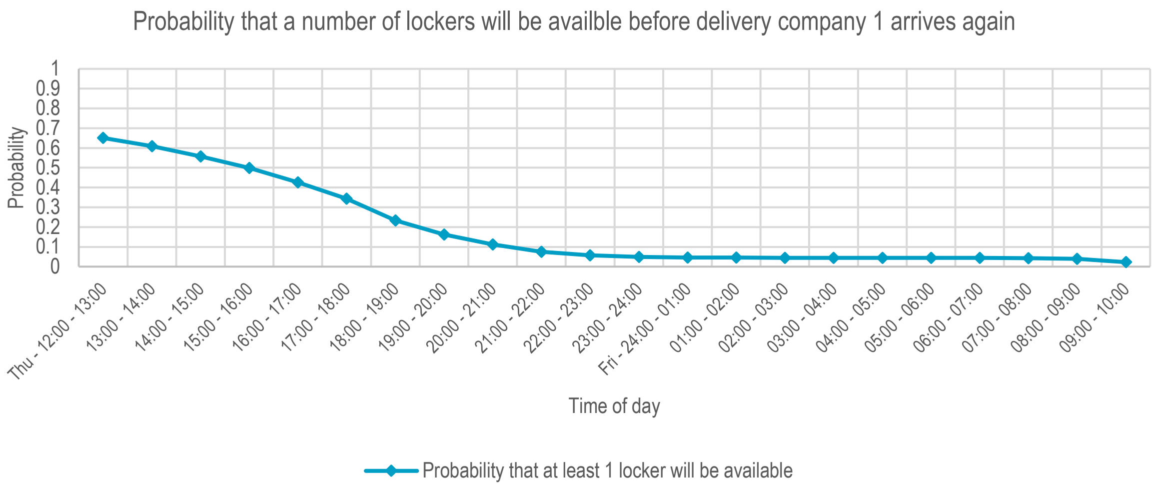

Figure 14 - The probability that one last-mile parcel will be collected before the next delivery moment

Scenario 2: Two last-mile parcels in the wall

When there are two last-mile parcels in the wall, there is a chance that least one of them will be emptied and the chance that they will both be emptied before the next deliverer will arrive. In Figure 15, these probabilities are projected. The chance that either of them will be collected during the day is fairly high. However, the chance that they will both be collected is approximately 0.4 at the highest.

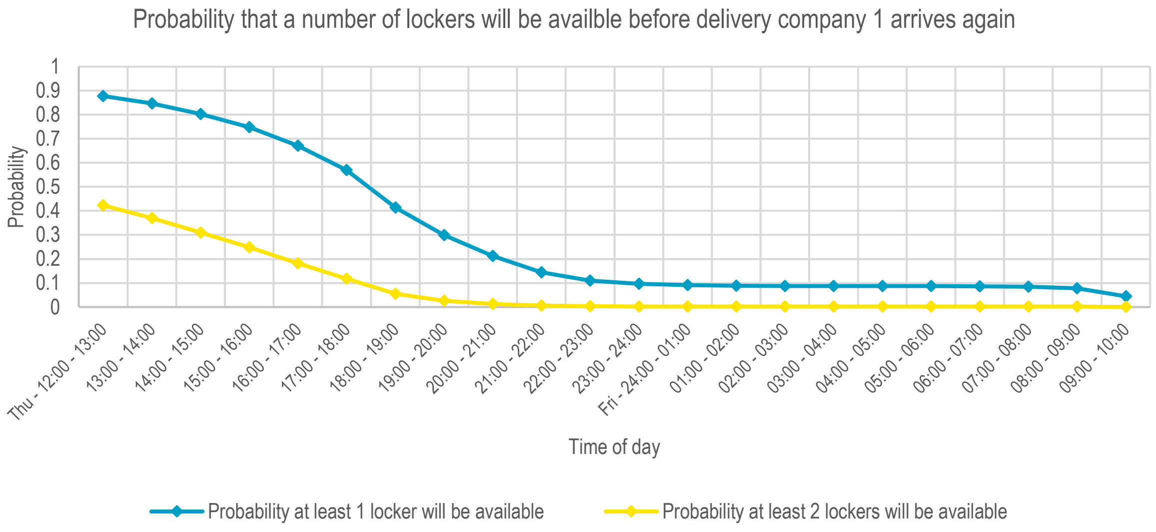

Figure 15 - The probability that at least one or two last-mile parcels will be collected before the next delivery moment

Scenario 3: Three last-mile parcels in the wall

In the scenario with three last-mile parcels in the wall, the chance that at least one parcel will be collected before the next deliverer arrives is almost 1 in the beginning. The chance that at least two parcels will be collected starts just above 0.7. The probabilities are shown in Figure 16.

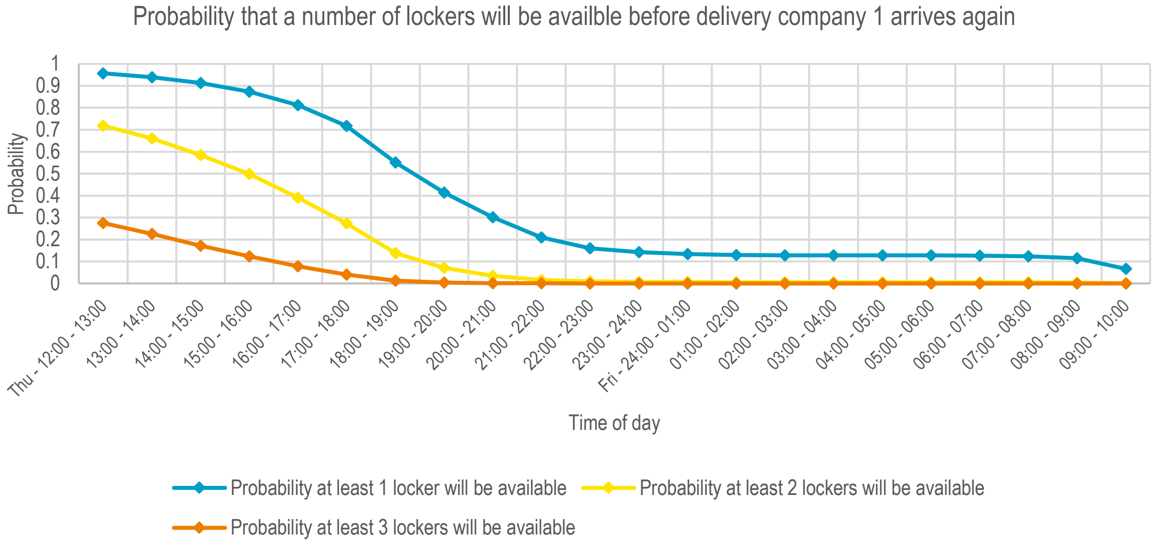

Figure 16 - The probability that at least one, two or three last-mile parcels will be collected before the next delivery moment

Scenario 4: Four last-mile parcels in the wall

Delivery companies arrive each day from Tuesday until Saturday. On Sundays and Mondays they do not visit the wall. To illustrate the effect, we analysed the problem instance from Thursday to Sunday. First, a normal weekday is shown in Figure 17. The graph starts just after delivery company 2 has left and ends at 10:00 in the next morning, when the next delivery company arrives. The different lines are chances that at least that number of lockers will become available before the next delivery company arrives. For example, the blue line is the chance that at least one locker becomes available. This probability consist of the chance that 1, 2, 3 or 4 lockers will become available before the next morning.

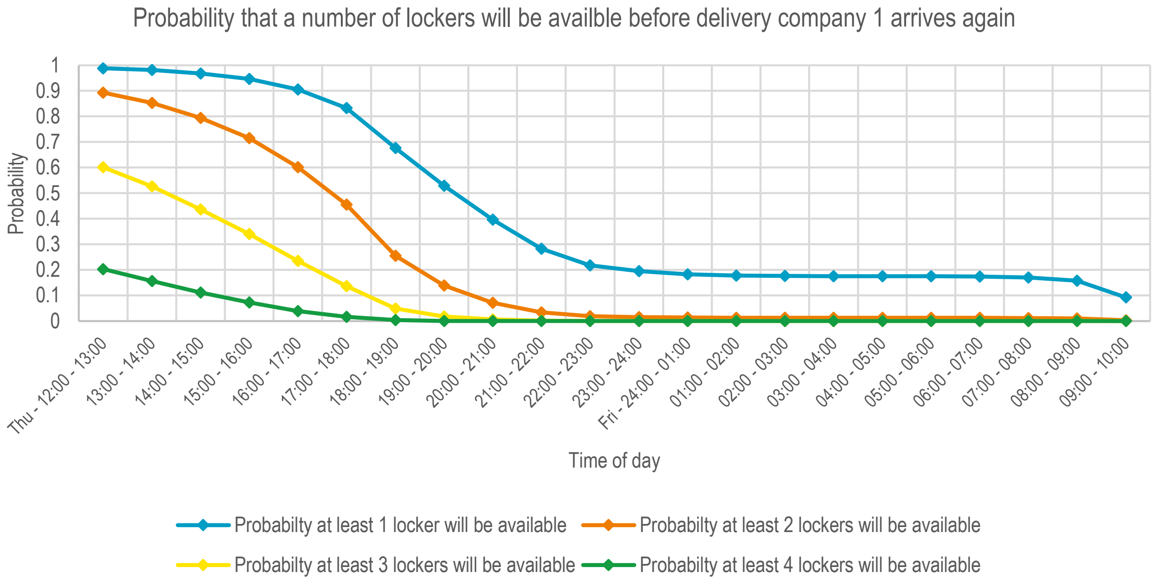

Figure 17 - Probability that a specific number of last-mile parcels will be picked up before the next delivery company arrives

In the graph it can be seen that the chance that one last-mile parcel will be picked up before the next delivery company arrives starts around 100%. Until approximately 19:30 the chance is above the 50%. This means that when a last-mile parcel is in the wall at that time, the chance is still fifty-fifty that it will be picked up before 10:00 in the morning. When a first-mile customer of delivery company 2 wishes to use the wall and there are lockers available, the choice must be made to either accept or reject the customer. When it is known that there are 5 lockers free at 19:30 and it is known delivery company 1 needs 5 lockers in the morning, accepting a first-mile customer of delivery company 2 means that there will be a 50% chance delivery company will have enough available lockers tomorrow.

As stated earlier, the delivery companies do not visit the wall on Sundays and Mondays. Therefore, the chances on Saturday afternoon and Sunday that parcels will be picked up before the next delivery company arrives are higher than on normal days. To illustrate this, we plotted the probabilities in Figure 18. As can be seen, the probabilities that one, two or three lockers will be emptied is very high during the weekend.

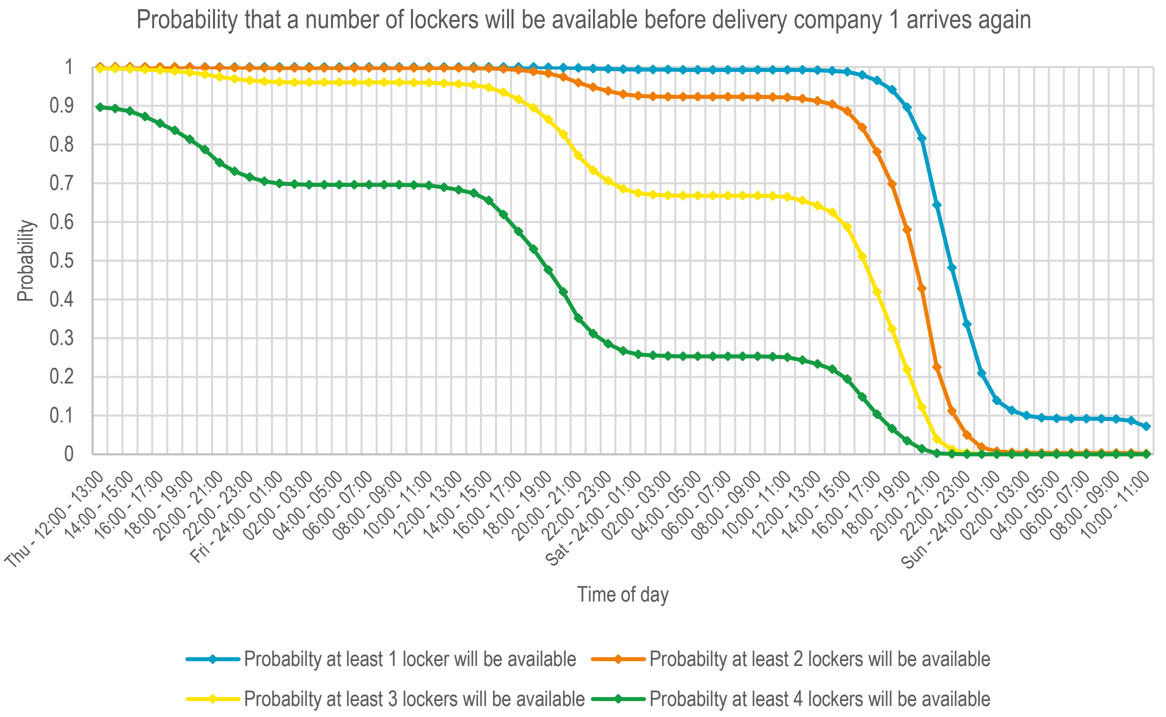

Figure 18 - Probabilities that a number of lockers will be available before delivery company 1 arrives again during the weekend

To get a good overview of the behaviour of the probabilities over the days and weekends the probabilities from Thursday to Tuesday are presented in Figure 19.

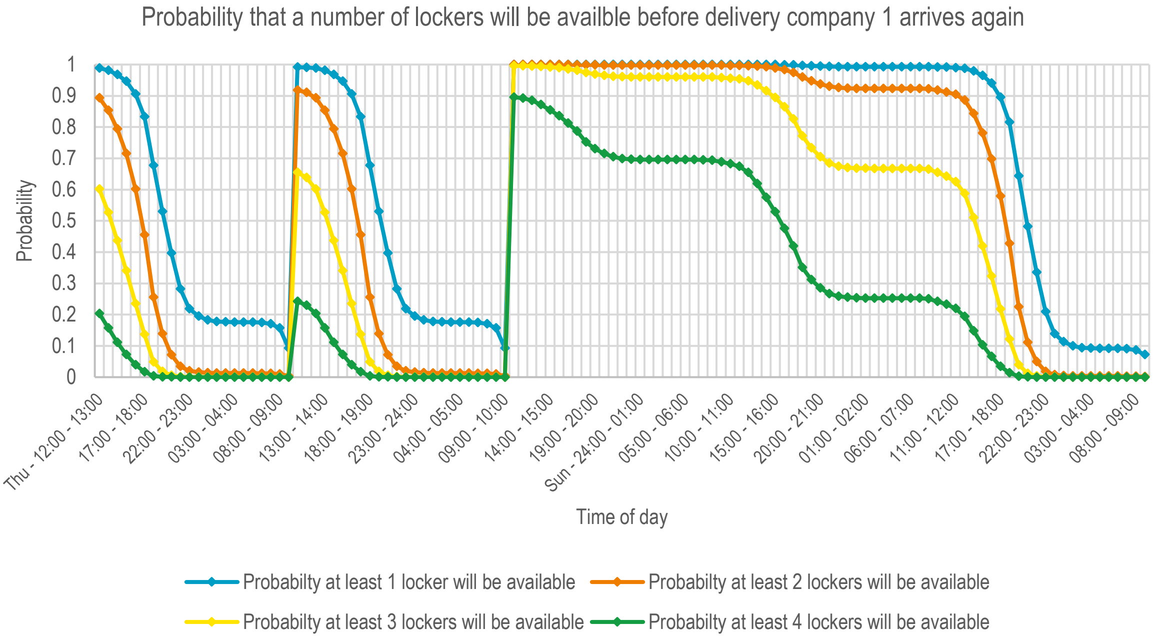

Figure 19 - Probabilties that a number of lockers will be emptied before the next delivery company arrives (from Thursday until Tuesday)

5.2.3 Admission policy

The probabilities that were calculated and presented in Section 5.2.2 can be the basis of an allocation policy to determine whether to accept or reject certain customers. When first-mile customers arrive at the wall and are not dedicated to the next arriving deliverer, accepting them will definitely lead to having a smaller capacity compared to when they are rejected. The capacity that is expected can be calculated by the number of available lockers, the number of lockers that are filled with first-mile parcels that are dedicated to the next arriving deliverer and the number of lockers that contain a last-mile parcel and are expected to be empty before the next deliverer arrives.

Using the policy, it must be possible to decide at every moment based on the state of the wall, type of customer and the time of day. In Figure 20, a flowchart is shown that shows the myopic allocation policy. First, it is checked whether the arriving customer belongs to the first arriving delivery company. If so, the customer is accepted because his parcel will first be emptied by the delivery company and will therefore not occupy a locker that may have to be used for last-mile parcels. If not, the capacity of the locker wall is checked. If the sum of first-mile parcels for delivery company 1 and empty lockers is bigger than the expected number of needed lockers, the customer is always accepted. If this is not the case, a calculation is triggered. This moment is represented by the green highlighted square. The calculation was explained in Section 5.2.1 and 5.2.2 and is based on the number of last-mile parcels that are currently present in the locker wall. Based on this probability, the decision is made to either accept or reject the customer. When accepting a customer now leads to the probability of not having enough lockers in the morning being smaller than the service level, the customer is not accepted. Otherwise, the customer is accepted.

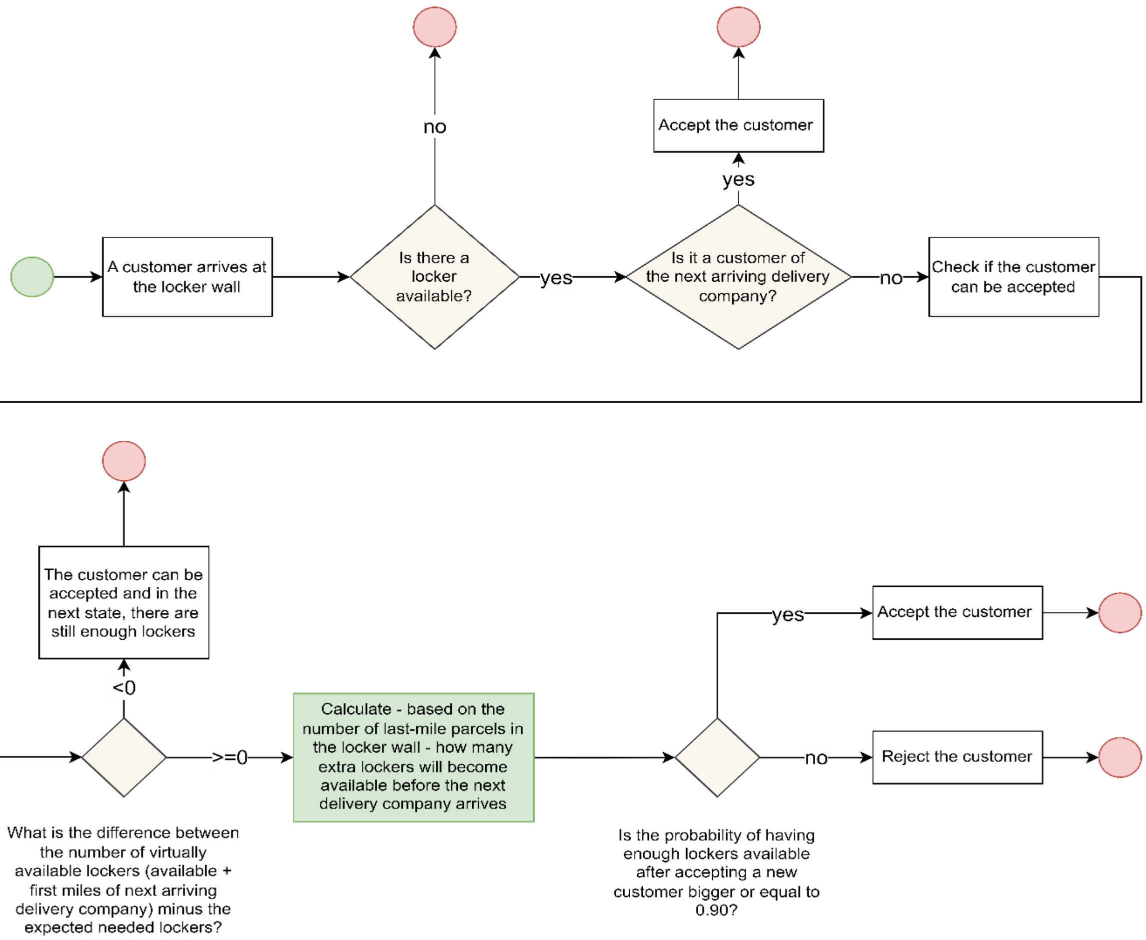

Figure 20 - A flowchart of the myopic allocation policy

As can be seen in the flowchart, the green square is not reached if there is already enough capacity available. In the event not enough capacity is available, the green square is entered. First, the number of last-mile parcels in the wall is determined. Based on this and the time of day, the calculation is made with which probability sufficient last-mile parcels will be picked up before the next delivery company arrives. In Table 5, the expected number of last-mile parcels that will be picked up before the next delivery company arrives is shown. The probability must be at least 90% and it depends on how many parcels are in the wall at that moment. For example, from 18:00 it cannot be said with a probability higher than 0.90 anymore that any of the last-mile parcels will be picked up, so no locker can be assigned anymore that is needed in the morning if the promised service level is 0.90.

We will illustrate the policy using an example. Let us assume that we know delivery company 1 that comes first tomorrow morning needs 4 lockers and it is currently 15:00. The state of the wall is as follows:

- Three lockers are empty;
- One locker contains a first-mile parcel for delivery company 1;
- Two lockers contain a first-mile parcel for delivery company 2;
- Four lockers contain last-mile parcels.

The question is whether we can now accept a customer of delivery company 2. When the delivery company arrives, he will first empty the first-mile parcel dedicated to him and he will be able to use the available lockers. That will give him four lockers in total. Looking at Table 5, we see that we expect that one of the four last-mile parcels will be collected before the next morning (4 last-mile parcels, 15:00 in the afternoon). Therefore, we can accept the new customer as we will still have enough lockers the next morning if one last-mile parcel is picked up.

Table 5 – Maximum number of last-mile parcels that are expected to be collected from the locker wall with a probability of at least 0.90.

	Number of last-mile parcels in the wall			
Tuesday - Saturday	*1*	*2*	*3*	*4*
12:00	0	1	1	2
13:00	0	0	1	1
14:00	0	0	1	1
15:00	0	0	1	1
16:00	0	0	0	1
17:00	0	0	0	1
18:00	0	0	0	0
19:00	0	0	0	0

5.3 MARKOV DECISION PROCESS POLICY

In this section, we will elaborate on the policy that is obtained by solving the MDP defined in Section 4.3. This policy will be more complex than the threshold policy and the myopic allocation policy because we have analysed the best decision to take in every feasible state. It includes both the inhomogeneous arrival pattern of last-mile parcel customers as the inhomogeneous arrival patterns of first-mile parcel customers. We will discuss the definition and input parameters in Section 5.3.1. Thereafter, we will present the resulting policy in Section 5.3.2.

5.3.1 Definition and parameters

The MDP formulation that was given in Section 4.3 was solved using the algorithm that was given in Section 4.3.7. The decision space consisted of nine possible actions. Each hour, the decision must be made how many first-mile parcels may be accepted from each delivery company. Due to computational limits, nine possible actions were considered. Each hour, the decision was made to accept 0, 1 or 2 parcels per delivery company in the next hour. Two delivery companies are involved, so nine different decision can be made in total. The parameters of the model and their values are shown in Table 6. The reward for a first-mile parcel of delivery company 1 is equal to 0.5, the reward for a first-mile parcel of delivery company 2 is equal to 0.7 and the penalty for not having enough space for a last-mile parcel in the morning is equal to 0.8. The latter is the sum of three different costs. First, the direct loss of 0.5 in income of a last-mile parcel, the loss of goodwill of the customer because his parcel is not delivered at the desired location is valued equal to 0.25 and the loss of goodwill of the delivery company is valued at 0.05.

Table 6 - Parameters and their values of the MDP

Parameter	Value
Reward first-mile parcel delivery company 1	0.5
Reward first-mile parcel delivery company 2	0.7
Penalty not enough space per last-mile parcel	0.8

5.3.2 Policy

The resulting policy gives an action for every possible state of the locker wall. In total, there are 6578 states (see Section 4.3.2). For every state, the policy prescribes to take any action in the list [0,1,2,3,4,5,6,7,8]. The number of allowed first-mile parcels of delivery company 1 and delivery company 2 that belong to each action. However, because in this situation delivery company 1 will be our premium customer, we will always accept first-mile parcels of delivery company 1. Therefore, we will implement the policy based on the actions for delivery company 2. Which means first-mile parcels of delivery company 2 will be accepted according to the policy and first-mile parcels of company 1 will always be accepted if there are lockers available.

Table 7 - The different actions that can be taken and the corresponding number of parcels that are allowed

Action	Number of allowed first-mile parcels of delivery company 1	Number of allowed first-mile parcels of delivery company 2
0	0	0
1	1	0
2	2	0
3	0	1
4	1	1
5	2	1
6	0	2
7	1	2
8	2	2

6. SIMULATION MODEL

Like Pazour & Roy (2015), we will use discrete-event simulation to analyse the performance of our policies. We have developed a simulation model using Plant Simulation. With this model, the performance of the policies found in Chapter 5 can be analysed, different configurations can be experimented with and alternative policies can be analysed (Robinson, 2014). Because the system we analyse is subject to a lot of variability and the processes are interconnected, analysing the policies in a simulation environment may give valuable insights into the behaviour of the different policies. First of all, a conceptual model is devised in which all the inputs and outputs of the model are described. Initially, a base model will be developed that will resemble a simple version of reality. In Section 6.1, an overview will be given of the conceptual model and the in- and outputs of the system. In Section 6.2, we will present the configuration settings of the simulation environment. Subsequently, we will define the experiments in Section 6.3.

6.1 CONCEPTUAL MODEL

The base model will be a simple representation of a locker wall. The system will be modelled using discrete event simulation. We will now describe the most important features on the basis of four concepts: the scope, the level of detail, the input and the output.

6.1.1 Scope

The model will represent one locker wall. The locker wall consists of 10 lockers. Initially, the lockers will all have the same size. Each delivery company will deliver once per day (at 10am and at 12am) and will bring a random number of parcels. This amount will be based on the historical data of the past year. The parcels will stay in the wall for a certain amount of time. This duration is taken from an exponential distribution which is based on the historical data. When parcels have not been picked up after seven days, the parcel is removed from the wall. This represents the event in which the parcel deliverer has to take the parcels with him again.

Furthermore, the arrival process of first-mile deliveries will be modelled as a non-homogeneous Poisson process. The arrival rate will vary by hour following $\lambda_{lmp}(i)$ and $\lambda_{fmd}(i)$ that were shown in the fourth column of Table 1. The arrival rates are calculated on average per locker. When the simulation initializes, the total number of lockers is multiplied with the lambdas and that results in the arrival rates that are used in de simulation. The assumption is made that the deliverer of the last-mile parcels will always pick up all first-mile deliveries that are dedicated to him. Therefore, the pickup time of first-mile deliveries is constant.

6.1.2 Level of detail

The model that will be developed will be used for statistical analysis. The model will not be used for visualization purposes. For this reason, the model will be abstract. It will present all the objects and processes but will not include unnecessary visual elements. The arrival and pickup process will be represented by distributions.

6.1.3 Input

The main inputs of the base model are as follows:

- the distribution of the number of last-mile parcels that are delivered each morning. On average, 0.25 last-mile parcels are delivered per locker. The distribution is uniform. The average is multiplied with the scaling factor. The lower bound is 0.75 times the scaled average. The upper bound is 1.25 times the scaled average.
- the distribution of the last-mile pickup process representing the customers;

- the distribution of the delivery process of first-mile parcels of delivery company 1;
- the distribution of the delivery process of first-mile parcels of delivery company 2;
- the income per last-mile parcel of delivery company 1. The default value is equal to 0.5;
- the income per last-mile parcel of delivery company 2. The default value is equal to 0.7;
- the income per first-mile parcel of delivery company 1. The default value is equal to 0.5;
- the income per first-mile parcel of delivery company 2. The default value is equal to 0.7;
- a penalty per last-mile parcel of delivery company 1 that cannot be delivered. The default value is equal to 0.8;
- the configuration of the locker wall (number and size of lockers) . The default configuration will be 10 lockers of 1 size;
- the proportion of delivery company 1 and delivery company 2. The default setting is that the number of parcels they bring follow the same distribution;
- a scaling factor that makes it easier to experiment with busier situations. It multiplies all arrival rates with the factor. The default value is equal to 1.

6.1.4 Output

When a parcel arrives, the arrival time is stored in a table. Once the parcel is picked up again, the departure time is saved and the duration the parcel has occupied the locker will be calculated and saved. Based on this, the average utilization of the locker wall will be calculated. Furthermore, the main output we focus on is the percentage of accepted last-mile parcels and accepted first-mile parcels. The percentages will be stored per delivery company and type of stream (first-mile or last-mile). This is necessary so we can see the effect if we are going to discriminate between different (classes of) deliverers.

6.2 CONFIGURATION

In this section, we will determine the configuration of the simulation model. First, we will calculate the warm-up period and the run length in Section 6.2.1. Thereafter, we will calculate the number of replications in Section 6.2.2.

6.2.1 Warm-up period and run length

Our simulation output is not cyclic and reaches a steady state. We want to delete the data points that lie within the period in which the output is not steady yet, because this would otherwise bias our results. We use the graphical method to determine our warm-up period (Robinson, 2014, p. 166). To determine the warm-up period and the run length we consider the service levels of the four parcel streams as key performance indicators. A locker wall with ten equally sized lockers is analysed and every time a parcel is accepted or rejected; the corresponding service level is updated. The FCFS policy is used with a scaling factor of three. The latter choice is made because the output of the system will need more time to stabilize if it is more crowded. Therefore, the choice is made to analyse a crowded situation and set the scaling factor equal to three. We took the average of 30 replications of 1000 days long and plotted the outcome in a graph. For visualisation purposes, the data was stored as 1 – the service level. For example, a service level of 0.9 was stored as 1 – 0.9 = 0.1. The resulting graph is presented in Figure 21.

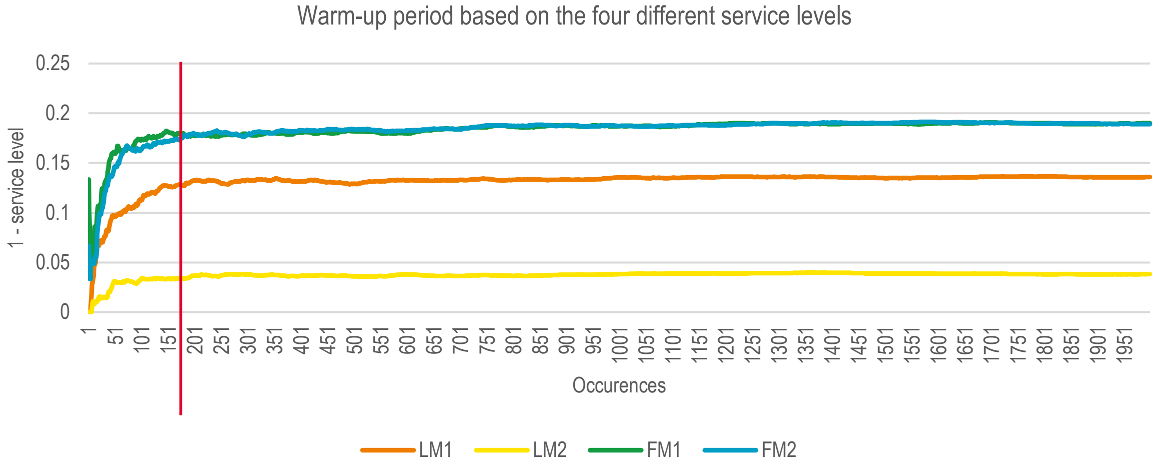

Figure 21 - Graph of the utilization over time for the warm-up period determination

As can be seen in Figure 21, the output becomes stable after approximately 250 occurrences. For the last-mile parcel streams, approximately 3300 data points were stored (only the first 2000 are shown in Figure 21). 250/3300 is approximately equal to 0.075. This means the warm-up period will be equal to 0.075 x 1000 (days) = 75 days.

Banks, Nelson, Carson, & Nicol (2009) state that the run length after the deletion point should be at least ten times as long as the period of which the data was deleted. Therefore, the run length is set equal to the warm-up length multiplied with 10. This results in a run length equal to 825 days (10 x 75 + warm-up period). Data will be stored from the 76^{th} day.

6.2.2 Number of replications

To determine the number of replications we need to get trustworthy results, we use the confidence interval method (Robinson, 2014, p. 184 - 186). We use a significance level of $\alpha = 0.05$. Initially, we execute 5 replications. All data is recorded after the warm-up period ends and the run length of all replications is equal to 750 days (excluding the warm-up period). For each run, the average per service level is taken and the cumulative average and standard deviation are calculated. These are used to calculate the confidence interval half width. Subsequently, the error relative to the mean is calculated by dividing the error by the mean. The first time the value is smaller (and stays smaller if the number of replications increase further) than the relative error, the number of replications is sufficient. Furthermore, at least 3 to 5 replications are recommended because the results of single replications should not be considered trustworthy (Robinson, 2014, p. 182). This was done for all four different service levels to ensure that enough replications are executed for all key performance indicators. In Table 8, the number of replications that are needed per performance indicator are presented. The full analysis per performance indicator can be found in Appendix A. The number of replications is 4.

Table 8 - The number of replications that are needed per performance indicator

Performance indicator	Number of replications needed to be able to say with 95% confidence the real mean is within the confidence interval
Service level of last-mile parcels of delivery company 1	2
Service level of last-mile parcels of delivery company 2	2
Service level of first-mile parcels of delivery company 1	3
Service level of first-mile parcels of delivery company 2	4

6.3 EXPERIMENTS

As stated earlier, the model revolves around one locker wall and two different delivery companies. The first delivery company arrives at 10:00 and the second delivery company arrives at 12:00. Initially, all lockers will have the same size and the locker wall will contain ten lockers. We will study four different situations:

- A benchmark situation in which a FCFS policy is used;
- A threshold policy as explained in Section 5.1;
- The myopic allocation policy as explained in Section 5.2;
- The MDP policy as explained in Section 5.3.

Furthermore, we will analyse the environment with a scaling factor of 1, a scaling factor of 2 and a scaling factor of 3. The former represents reality, the latter represents a situation in which there are more customers that wish to use the wall. Additionally, the number of parcels the delivery company bring follow the same uniform distribution. The threshold factor takes the values [0.1, 0.3, 0.5, 0.7, 0.9]. Furthermore, we will experiment with different goals for the service level in the myopic allocation policy. The service levels that will be aimed for take the values [0.1, 0.3, 0.5, 0.7, 0.9]. The policy that was obtained using the MDP has only one setting. The total number of experiments thus amounts to 12 (different settings) * 3 (scaling factor) = 36. The experimental settings are shown in Appendix B. Furthermore, an overview of the simulation model is shown in Appendix C.

7. NUMERICAL EXPERIMENTS AND RESULTS

In this section, we will discuss the results of the experiments that were presented in Section 6.3 and fully shown in Appendix C. We will analyse the outcomes in Section 7.1. Thereafter, we will execute additional experiments based on the most promising outcomes of the first thirty experiments in Section 7.2.

7.1 GENERAL OUTCOMES

In Table 9, the results of the thirty experiments are shown. Furthermore, the policy that was used per experiment is given in the second column. The full settings can be found in Appendix B. Furthermore, the service levels of the four different parcel streams are provided together with the total profit and the average utilization of the locker wall. The four different policies were evaluated with different settings and with three different scaling factors.

Table 9 - Results of the experiments

Scaling factor	Policy	LM1 service level	FM1 service level	LM2 service level	FM2 service level	Total profit	Utilization
1	FCFS	0.99	0.96	1.00	0.96	3295.28	0.52
1	Thr [0.1]	1.00	0.08	1.00	0.09	1114.87	0.19
1	Thr [0.3]	1.00	0.33	1.00	0.33	1736.14	0.29
1	Thr [0.5]	1.00	0.61	1.00	0.60	2417.16	0.40
1	Thr [0.7]	1.00	0.81	1.00	0.81	2930.25	0.47
1	Thr [0.9]	1.00	0.92	1.00	0.93	3217.07	0.51
1	Myo [0.1]	1.00	0.96	1.00	0.95	3292.43	0.52
1	Myo [0.3]	1.00	0.96	1.00	0.95	3287.54	0.52
1	Myo [0.5]	1.00	0.96	0.99	0.95	3282.48	0.52
1	Myo [0.7]	1.00	0.96	0.99	0.94	3273.19	0.52
1	Myo [0.9]	1.00	0.96	0.99	0.93	3264.38	0.52
1	MDP	1.00	0.97	0.99	0.77	3036.33	0.48
2	FCFS	0.96	0.74	0.97	0.74	5610.33	0.80
2	Thr [0.1]	0.99	0.04	0.96	0.04	2181.75	0.28
2	Thr [0.3]	0.99	0.16	0.96	0.16	2761.15	0.37
2	Thr [0.5]	0.99	0.31	0.96	0.31	3524.92	0.49
2	Thr [0.7]	0.99	0.48	0.96	0.49	4384.98	0.63
2	Thr [0.9]	0.97	0.66	0.97	0.66	5214.28	0.76
2	Myo [0.1]	0.98	0.75	0.96	0.68	5476.74	0.80
2	Myo [0.3]	0.98	0.76	0.96	0.65	5430.54	0.80
2	Myo [0.5]	0.99	0.78	0.95	0.62	5372.81	0.80
2	Myo [0.7]	0.99	0.79	0.95	0.59	5317.57	0.79
2	Myo [0.9]	1.00	0.81	0.95	0.56	5259.83	0.78
2	MDP	0.99	0.81	0.95	0.58	5314.30	0.75
3	FCFS	0.90	0.56	0.95	0.57	6932.95	0.90
3	Thr [0.1]	0.99	0.03	0.92	0.03	3225.58	0.29
3	Thr [0.3]	0.99	0.12	0.92	0.12	3862.87	0.38
3	Thr [0.5]	0.99	0.22	0.92	0.23	4667.92	0.51
3	Thr [0.7]	0.97	0.35	0.92	0.35	5557.64	0.66

3	Thr [0.9]	0.93	0.49	0.94	0.49	6496.39	0.82
3	Myo [0.1]	0.98	0.64	0.90	0.39	6597.62	0.87
3	Myo [0.3]	0.99	0.67	0.88	0.35	6520.53	0.86
3	Myo [0.5]	0.99	0.70	0.88	0.32	6479.44	0.86
3	Myo [0.7]	0.99	0.72	0.87	0.30	6420.50	0.85
3	Myo [0.9]	0.99	0.73	0.87	0.27	6361.96	0.84
3	MDP	0.96	0.64	0.91	0.46	6891.78	0.86

We will now discuss the different variables in the following paragraphs. To determine whether the differences in the results are significant, T-tests are executed for all results. A T-test is statistical test that can be conducted to see if, given a certain significance level, one may presume the means of two datasets are different. The default level of significance used in the T-tests is 95%. We will now first discuss the situation with a scaling factor of 1 (normal situation) in Section 7.1.1. Subsequently, we will discuss the results of the experiments with a scaling factor of 2 in Section 7.1.2. Finally, we will discuss the results of the experiments with a scaling factor of 3 in Section 7.1.3.

7.1.1 Scaling factor = 1

In the normal situation we expect to see less strict admission policies to perform well in terms of profit. Therefore, we expect the FCFS policy, the myopic allocation policies focused on lower probabilities to perform best. However, we do expect all other policies to outperform the FCFS policy in terms of the service levels of delivery company 1, while these policies focus on obtaining a high service level for delivery company 1. Furthermore, the utilization is expected to be highest using the FCFS policy as all parcels will be accepted.

The result in total profit for the normal situation (scaling factor = 1) is shown in Figure 22. It can be seen that the FCFS policy works best in this situation. The threshold policy performs best if the threshold is higher. When the threshold is equal to 1, the threshold policy behaves the same as the first-come first-serve policy. Therefore, it is logical that the result approaches the FCFS result. The myopic allocation policies perform approximately equally well and slightly underperform compared to the first-come first-serve policy. The performance of the myopic allocation policies increases as the service level it aims for decreases. The myopic allocation policy will behave as the FCFS policy when the service level it aims for is equal to zero. In this situation, the results of the T-tests showed that the results of the FCFS policy and all settings of the myopic allocation policies are not significantly different. Finally, the MDP policy only outperforms the lower threshold functions.

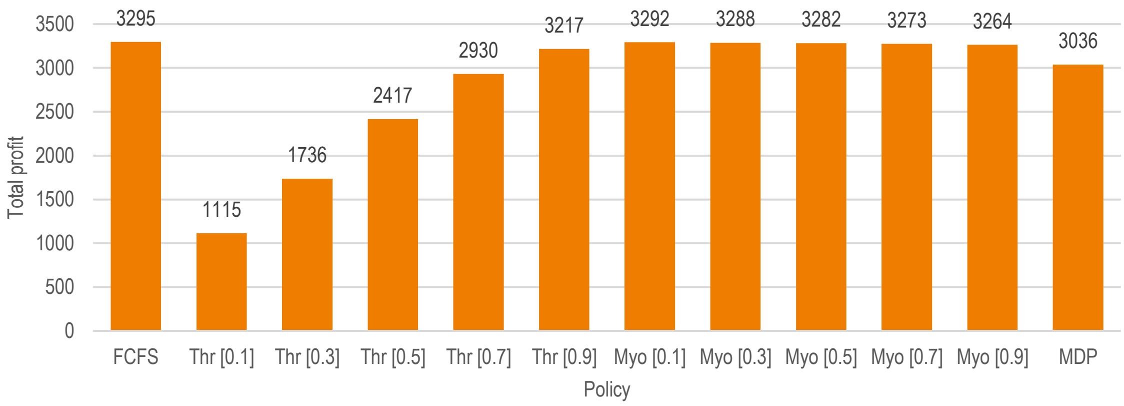

Figure 22 -The total profit per policy for the normal situation

In Figure 23, the service levels are shown per policy. Four different service levels are distinguished. Namely, the service level of last-mile parcels of delivery company 1 and delivery company 2 and the service level of first-mile parcels of delivery company 1 and delivery company 2. It can be seen that the FCFS policy approximately reaches service levels above 0.95 for all parcel streams. The threshold functions and the myopic allocation policies perform better on the last-mile streams but underperform on the first-mile stream (specifically the first-mile stream of delivery company 2). The MDP policy performs bad on the first-mile parcels of delivery company 2 and performs well on the last-mile parcel streams and the first-mile parcel stream of delivery company 1. The results in service level are as expected. The service levels of delivery company 1 are equal or higher using the admission policies compared to the FCFS policy. However, we saw in Figure 26 that the threshold and MDP policies are too strict in this situation as the slightly improved service level of delivery company 1 costs a big proportion of the total profit.

The myopic allocation policies and the FCFS policy all have a utilization of 0.52. As expected, the utilization is highest using the FCFS policy. Furthermore, it seems to be the maximum utilization obtainable in these circumstances as the service levels all approach 1.0. The service levels of the lower threshold policies are between 0.24 and 0.44. This can mainly be explained by the low service levels of the first-mile parcel streams.

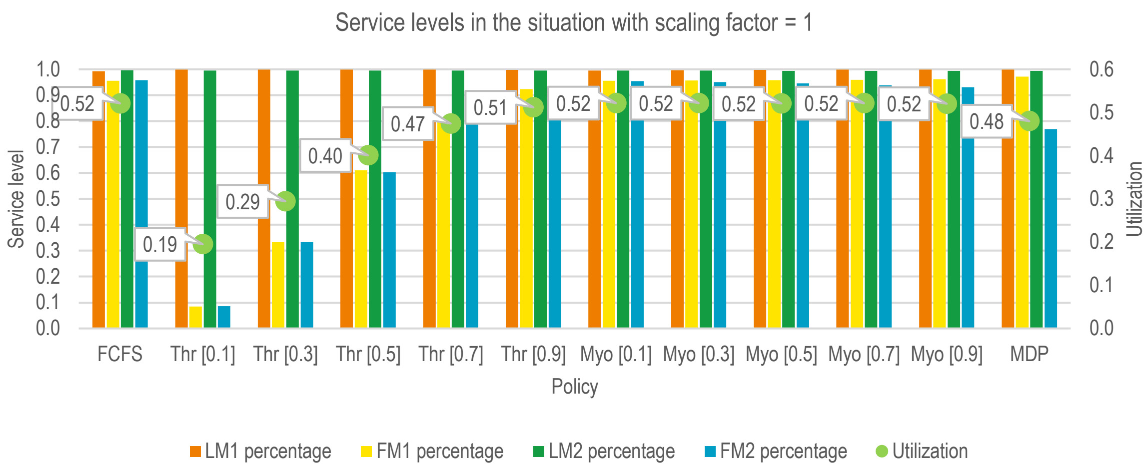

Figure 23 - Service levels obtained by the different policies in the normal situation

7.1.2 Scaling factor = 2

In a more crowded situation, we expect the myopic allocation policies and the MDP policies to perform better compared to the FCFS policy. The differences in the service levels of delivery company 1 will be bigger. The myopic allocation policies will deny more first-mile parcels of delivery company 2 to ensure a good performance on the service levels of delivery company 1. Therefore, we expect the first-mile service level of delivery company 2 to be lower than using the FCFS policy. The threshold policies will probably also perform better on the last-mile service level of delivery company 1. However, they may be too strict for the first-mile parcel streams .

The result in total profit for the situation in which it is more crowded (scaling factor = 2) is shown in Figure 24. It can be seen, that the FCFS policy still outperforms all other policies. The difference between the total profit using the FCFS policy and the other policies only became bigger. The lower threshold policies do not perform well. The myopic allocation policies perform approximately the same. The MDP policy only outperforms all threshold policies and the highest myopic policy.

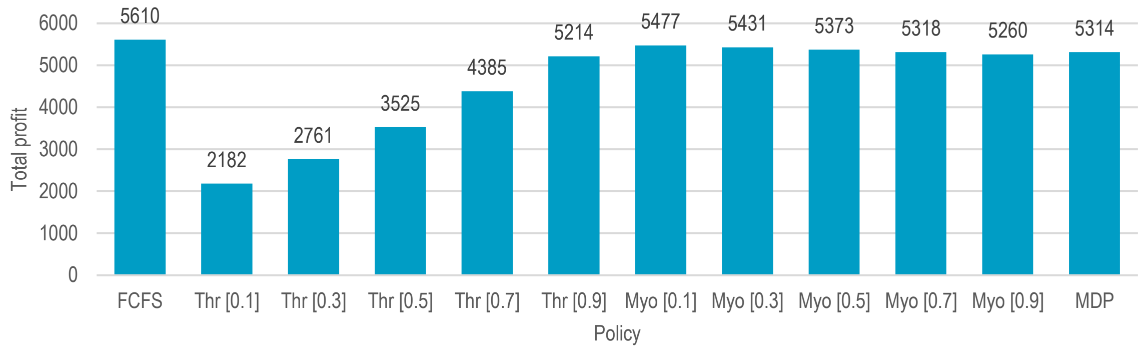

Figure 24 - The total profit per policy for the situation with a scaling factor of 2

In Figure 25, the results of the service levels are shown in the situation with a scaling factor equal to two. It can be seen that the last-mile service levels are still above 0.9 using the FCFS policy. The first-mile stream service levels do not exceed 0.75. With the first three threshold policies and the last myopic allocation policy a service level of 1.0 is obtained for the first last-mile stream. As we expected, the threshold admission policies are too strict for the first-mile parcel streams. In terms of total profit, the improved service level of delivery company 1 does not outweigh the missed income of the first-mile parcels. The first-mile parcel stream of delivery company 2 performs best using the FCFS policy. The myopic allocation policies perform better for the first-mile stream of delivery company 1 than the first-mile stream of delivery company 2. This is logical as myopic allocation policies only accept the first-mile parcels of delivery company 2 with a certain probability and the first-mile parcels of delivery company 1 always if there is space.

The utilizations of the FCFS and the myopic allocation policies are all around 0.80. The threshold utilizations vary between 0.32 and 0.69. The utilization of the MDP policy is equal to 0.75. This is lower than the utilization that was obtained using the myopic policies or the FCFS policy. This is mainly caused by the fact that the service level of the first-mile parcels of delivery company 2 is lower. Therefore, less parcels go through the wall as the other service levels seem comparable.

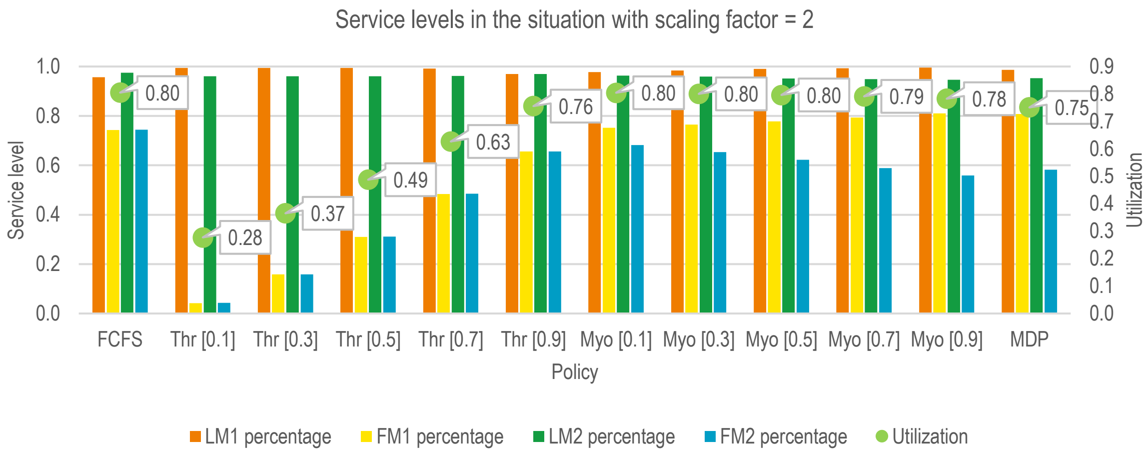

Figure 25 - Service levels obtained by the different policies in the situation with scaling factor = 2

7.1.3 Scaling factor = 3

In the most crowded situation, we expect the service levels to be lower. The last-mile service level of delivery company 1 should be highest not using the FCFS policy. The utilization will probably be the highest using the FCFS policy. In terms of total profit, we expect the myopic allocation and the MDP policies to perform well due to the extra parcels of delivery company 1 and the avoidance of penalties.

In Figure 26, the result in total profit of a situation that is three times as crowded as normally is shown. It can be seen that the FCFS policy still performs best, closely followed by the MDP policy. It is surprising the myopic allocation policies do still not outperform the FCFS policy. The threshold policy and myopic allocation policy still perform best when they approach the FCFS policy.

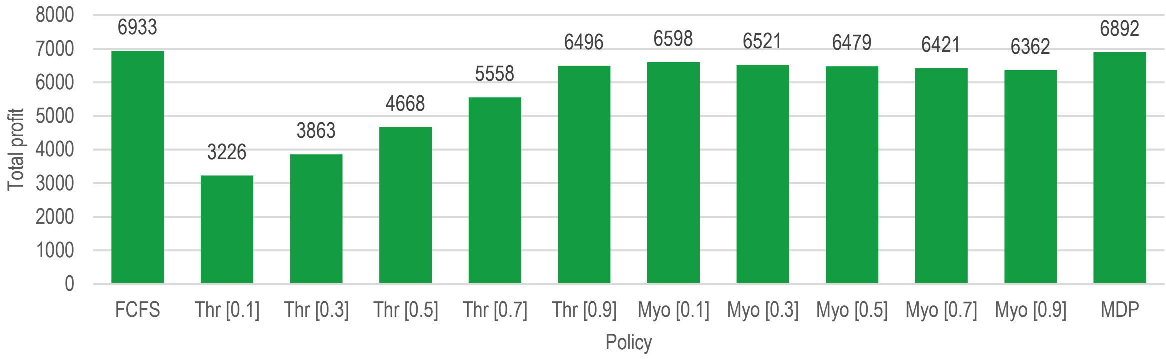

Figure 26 -The total profit per policy for the situation with a scaling factor of 3

In Figure 27, the results of the service levels in the situation with a scaling factor equal to three are shown. It can be seen that the service levels of the last-mile parcel streams still exceed 0.90 with all policies. However, the service levels of the first-mile parcel streams are low. Using the FCFS policy, they first-mile service levels do not exceed 0.60. The best service level for a first-mile parcel stream is obtained using the myopic allocation policies. The first-mile parcels stream of the first delivery company exceed 0.60 for all four settings. Furthermore, the last-mile service level of delivery company 1 is high using the myopic allocation policies.

The utilization is 0.90 using the FCFS policy and this is the highest utilization that was obtained during the experiments. The utilization using the myopic allocation policies is 0.87 at the highest. Using the MDP policy, a utilization of 0.86 is obtained. The difference between the best-performing myopic allocation policy and the FCFS policy is 0.03 in terms of utilization. This difference can be explained by the fact that the myopic allocation policy builds in a safety factor to ensure the service level of delivery company 1 can be obtained with a certain probability. The penalties that are avoided using this policy do not outweigh the income that is missed by the rejected parcels of delivery company 2.

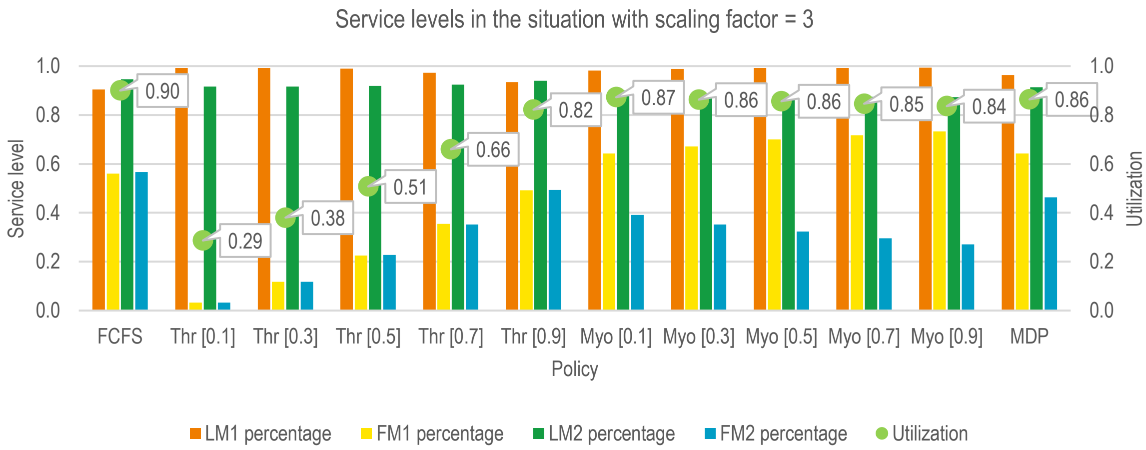

Figure 27 - Service levels obtained by the different policies in the situation with scaling factor = 3

7.2 ADDITIONAL EXPERIMENTS

In Section 7.1, we saw the results of 36 different experiments (Appendix B). The results of these experiments gave some interesting insights. To further deepen these insights, we will execute extra experiments in this section into two different directions. First, we will look at the myopic allocation policy and the effect of the cost setting in Section 7.2.1. Thereafter, we will look at the MDP policy in Section 7.2.2.

7.2.1 Myopic allocation policy

We have observed that the myopic allocation policy works better in terms of total profit if the service level that is aimed for is lower. When the probability with which first-mile parcels of delivery company 2 are accepted is lower, the total profit is higher. However, we only tried four different settings. Namely, the situations that will only accept first-mile parcels of delivery company 2 if the probability is 0.7, 0.8, 0.9 and 1.0 that enough lockers will be available in the morning. As the probability approaches 0, the myopic allocation policy will behave like the FCFS policy. Then the policy will accept first-mile parcels of delivery company 2 even when that will lead to the probability of having enough lockers in the morning being zero. However, we want to see if there is a setting with which the myopic allocation outperforms the FCFS policy. Therefore, we will try more different settings. Thereafter, we will see for the different myopic allocation settings in which environments they perform best. Therefore, we will experiment with the income per parcel and the penalty per locker to see what the effect is.

Different environment settings

In the previous paragraph, we concluded that the myopic allocation policies will not outperform the FCFS policy with the current settings. In this paragraph, we will investigate the effect of the other settings and see if there are situations in which the myopic allocation policy will outperform the FCFS policy. As we specified in 6.1.3, the income per parcel of delivery company 1 is equal to 0.5 and the income per parcel of delivery company 2 is equal to 0.7. Furthermore, if a last-mile parcel of delivery company 1 cannot be accepted a penalty of 0.8 will be incurred. We will analyse the situation in which the income per parcel of the different delivery companies is equal and the situation in which the parcels of delivery company 1 will yield 0.7 and the parcels of delivery company 2 0.5. We will execute the experiments in an environment with a scaling factor of 3, because the differences will be clearer if the situation is more crowded. The penalty of 0.8 will stay the same in all situations.

In Figure 28, the results in total profit are shown of the situation in which the income per parcel of both delivery companies is equal to 0.5. The FCFS policy performs best followed by the lower myopic allocation policies. The higher the probability of the myopic allocation policy, the worse it performs. So, in a situation in which the income per parcel is equal, the FCFS policy would still outperform the myopic allocation policies.

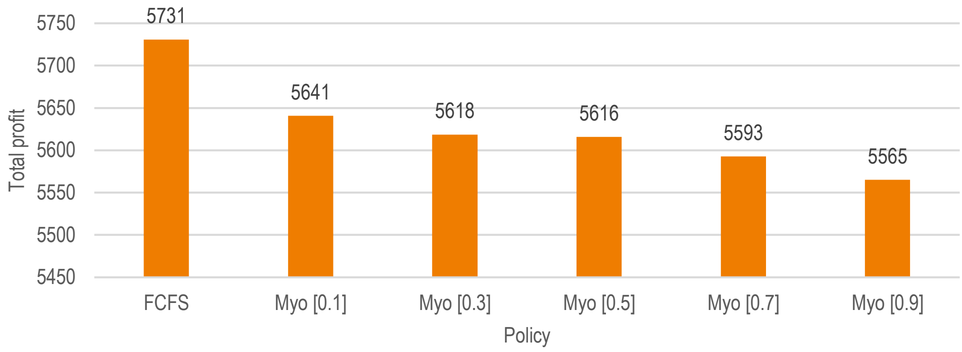

Figure 28 - The results in total profit of the situation in which the income per parcel is equal for both delivery companies

In Figure 29, the results are shown of the situation in which a parcel of delivery company 1 yields 0.7 and a parcel of delivery company 2 yields 0.5. The results show that the myopic allocation policies now outperform the FCFS policy. The performance of the myopic allocation policies increase with the probability that is configurated. The best-performing policy is the myopic [0.7] policy with a total profit of 7008.

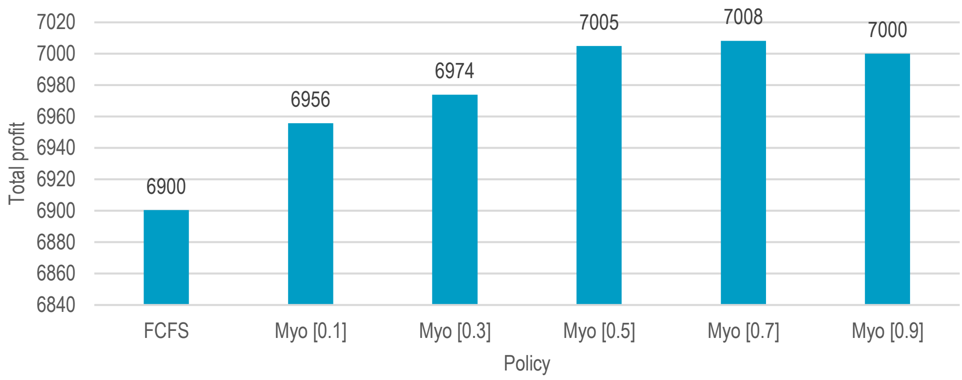

Figure 29 - The results in total profit of the situation in which the income per parcel of delivery company 1 is higher

Now, we will take the best performing myopic allocation policy and see what the influence is of changing the penalty that is incurred if a last-mile parcel of delivery company 1 cannot be delivered. We will use the last situation (delivery company 1 parcels yield 0.7 and delivery company 2 parcels yield 0.5). We will compare the FCFS policy with the myopic [0.7] policy on the result of total profit and we will analyse three different penalty values: 0.4, 0.8 (normal situation) and 1.2.

In Figure 30, we can see the results of three experiments. When the penalty is equal to 0.4, the FCFS policy and the myopic [0.7] policy are almost equal in terms of total profit. The myopic policy will have less income due to the higher number of rejections of first-mile parcels of delivery company 2. However, it compensates for that by having less costs because the service level of last-mile parcels of delivery company 1 is higher. If the penalty is lower, rejecting first-mile parcels of delivery company 2 will start to cost more (missed income) than the higher service level of last-mile parcels of delivery company 1 will compensate for. On the other hand, if the penalty is adjusted to the other direction and set equal to 1.2, the difference grows. It becomes more valuable to choose to reject first-mile parcels of delivery company 2 to ensure a high service level of last-mile parcels for delivery company 1.

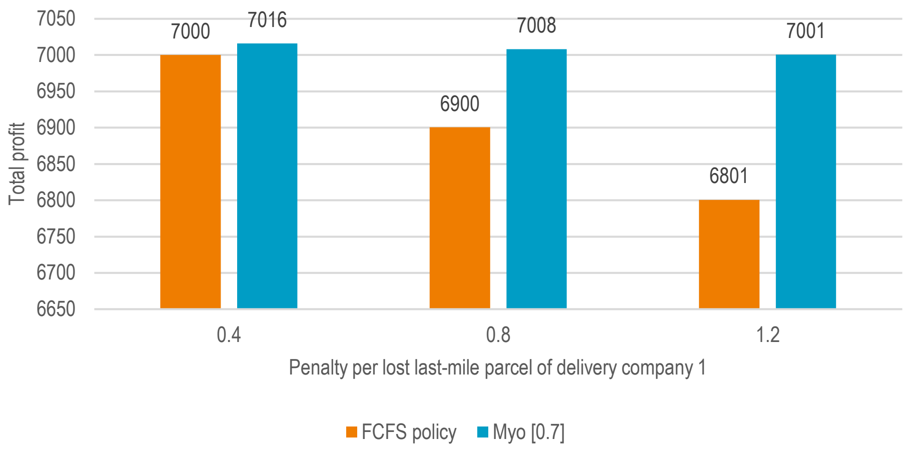

Figure 30 - The results of the three experiments with different penalties for lost last-mile parcels of delivery company 1

7.2.2 MDP policy

We observed that the MDP policy performed relatively better when the scaling factor was higher. So, when we analyse more crowded situations it performs better in terms of total profit. Therefore, we want to find out when the MDP policy outperforms the FCFS policy. In Section 7.1.3, we saw that the MDP policy started to perform better with an increasing scaling factor. Therefore, we will analyse the situations in which the scaling factor is equal to 4 and 5.

In Figure 31, the results of the experiments are shown. It can be seen the MDP policy now outperforms the FCFS policy and the difference only grows when it gets more crowded. The results in terms of service level are shown in Table 10. It can be seen, the service level of last-mile parcels of delivery company 1 is still 0.99. However, the service levels of the first-mile stream of delivery company 2 is really low in the first situation and in the second situation even equal to 0. The results show that it may be interesting to open really crowded walls for the most important party entirely and for another party only for last-mile delivery.

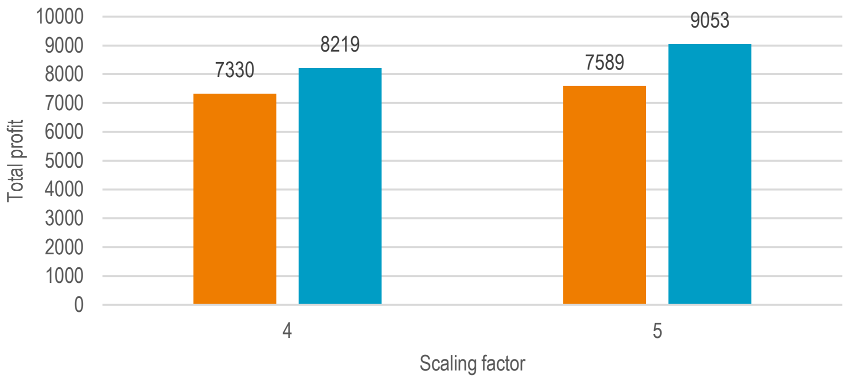

Figure 31 - The results of the experiments with the FCFS and MDP policy for scaling factors equal to 4 and 5

Table 10 - The results in terms of service levels of the experiments in crowded situations with the FCFS and MDP policy

Policy	Scaling factor	LM1 service level	FM1 service level	LM2 service level	FM2 service level
FCFS	4	0.82	0.44	0.85	0.44
MDP	4	0.99	0.68	0.68	0.12
FCFS	5	0.76	0.37	0.80	0.37
MDP	5	0.99	0.67	0.49	0.00

8. ANALYSES FOR PRACTICAL USE

In this chapter, we will analyse different scenarios that can provide helpful insights. Initially, we will analyse the different scenarios based on the policy that is currently being used (FCFS). First, we will analyse the scenario in which no first-mile parcels are accepted in Section 8.1. Thereafter, we will analyse a scenario in which only last-mile parcels are accepted in Section 8.2. Furthermore, we will analyse the scenario in which the second arrival of last-mile parcels is bigger than the first last-mile arrivals. Then, we will discuss the effect of spreading the delivery companies more over the day in Section 8.4. Finally, we will try to determine when a locker wall should be considered to be expanded in Section 8.5.

8.1 NO FIRST-MILE PARCELS

In Figure 32, the service levels and utilization are shown of the situation in which no first-mile parcels are accepted. The situation is shown for three different scaling factors. When the scaling factor is equal to 1 (the normal situation), all last-mile parcels can be accepted and the average utilization is equal to 0.16. Furthermore, when the scaling factor is equal to two, almost all parcels of delivery company 1 can be accepted and approximately 95% of the parcels of delivery company 2. The average utilization increases to 0.23. In the last situation (scaling factor = 3), the service level of delivery company 1 is still almost equal to 1.0. The service level of the second delivery company is approximately equal to 0.9. The average utilization increases to 0.25.

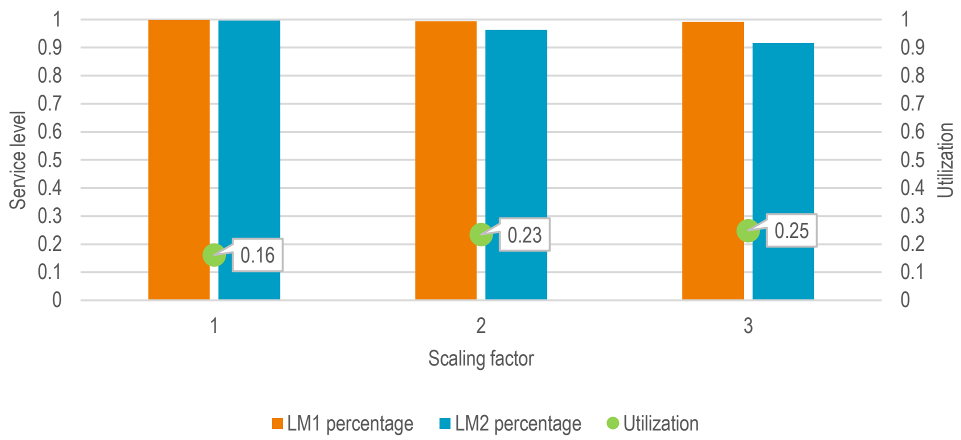

Figure 32 - Service levels and utilization of three situations in which no first-mile parcels are accepted

In Table 11, the results of the scenarios without first-mile parcels are shown next to the situation in which there are first-mile parcels. When the scaling factor is equal to 1, the service levels are all equal to 0.99 or 1.00. However, the utilization without first-mile parcels is really low. It is only equal to 0.16. This may be explained by the fact that the wall is never filled entirely and most of the parcels will be picked up before the evening. What is interesting to see, is that in the situation with a scaling factor equal to 3, the last-mile service level of delivery company 2 decreases from 0.95 to 0.91. This can be explained by the fact that accepted first-mile parcels of delivery company 2 normally ‘reserve’ spots for the last-mile parcels. As the first-mile parcels of delivery company 2 cannot be collected by delivery company 1 and delivery company 2 can fill the wall with his last-mile parcels after collecting his first-mile parcels.

Table 11 - The results in service levels and utilization of the normal situation compared to the situation without first-mile parcels

		Normal situation	No first-mile parcels
Scaling factor = 1	LM1 percentage	0.99	1.00
	LM2 percentage	1.00	1.00
	Utilization	0.52	0.16
Scaling factor = 2	LM1 percentage	0.96	0.99
	LM2 percentage	0.97	0.96
	Utilization	0.80	0.23
Scaling factor = 3	LM1 percentage	0.90	0.99
	LM2 percentage	0.95	0.91
	Utilization	0.90	0.25

8.2 NO LAST-MILE PARCELS

In Figure 33, the service levels and the utilization are shown of the situation in which no last-mile parcels are accepted. The situation is shown for three different scaling factors. When the scaling factor is equal to 1, all first-mile parcels can be accepted and the average utilization is equal to 0.37. When the scaling factor is equal to two, both service levels are approximately 0.85. The average utilization increases to 0.69. Furthermore, when the scaling factor is equal to three, the service levels decrease below the 0.7 and the average utilization increases to 0.83. It can be seen that the first-mile parcel stream is bigger than the last-mile parcel stream as the average utilization is almost four times as high as it was in Figure 32. Furthermore, the service levels are more equal than in the situation in which no first-mile parcels were accepted. This can be explained by the fact that last-mile parcels of delivery company 1 always arrive first but first-mile parcels of both companies arrive equally spread during the day.

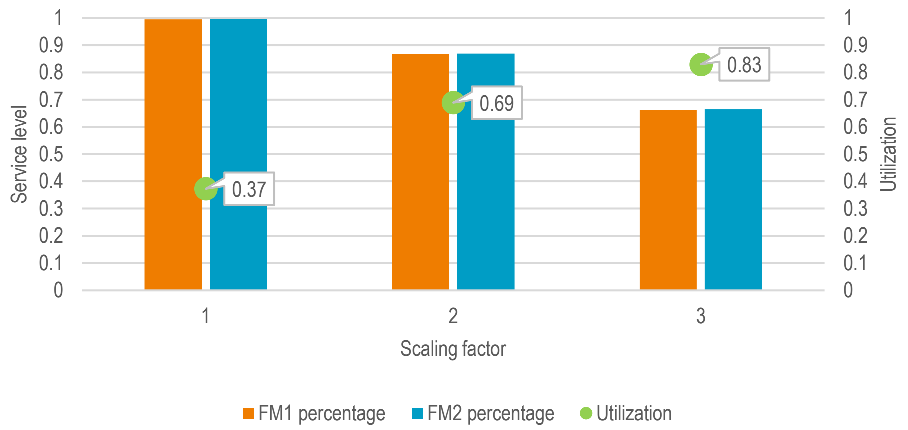

Figure 33 - Service levels and utilization of three situations in which no last-mile parcels are accepted

Table 12, the results are shown of the normal scenarios and the scenarios in which no last-mile parcels are accepted. It can be seen that the first-mile service levels are equal to 1.00 in the normal situation instead of 0.96. The utilization of the wall is 0.37 instead of 0.52. Furthermore, in the situation with a scaling factor of 2, a service level of 0.90 is almost obtained for both streams with a utilization of 0.69. In the last scenario (scaling factor = 3), the service levels are only slightly better and the utilization is 0.83. Compared to the situation we analysed in 8.1, in which no first-mile parcels were accepted, the effect on the service levels and utilization is smaller. Banning last-mile parcels from the wall will enable the company to obtain an acceptable service level in situations that are twice as crowded as normal. But when it is more crowded, the service levels decrease rapidly.

Table 12 - The results in service levels and utilization of the normal situation compared to the situation without last-mile parcels

		Normal situation	No last-mile parcels
Scaling factor = 1	FM1 percentage	0.96	1.00
	FM2 percentage	0.96	1.00
	Utilization	0.52	0.37
Scaling factor = 2	FM1 percentage	0.74	0.87
	FM2 percentage	0.74	0.87
	Utilization	0.80	0.69
Scaling factor = 3	FM1 percentage	0.56	0.66
	FM2 percentage	0.57	0.66
	Utilization	0.90	0.83

8.3 MORE "OTHER" DELIVERY COMPANIES

In Figure 34, the results in terms of service levels and average utilization are shown for six experiments. In the experiments, the situations are analysed in which the proportion of other delivery companies is bigger. This is simulated by making the share of delivery company 2 bigger. Normally, the distribution was 50/50 and now they are evaluated for a distribution of 0.3 and 0.1. This means that 70% and 90% of the last-mile and first-mile parcels will be from delivery company 2, respectively. Delivery company 2 can represent multiple different delivery companies that arrive later on a day than delivery company 1.

In the first and third experiment, the distributions are evaluated with the scaling factor equal to 1. The utilization is approximately 0.50 in both cases and the service levels are all above the 0.95. In the other experiments it can be observed that the service level of the last-mile parcels of delivery company 2 is really high whereas the other three service levels decline. This can be explained by the situation in which a lot of first-mile parcels of delivery company 2 are accepted and therefore last-mile parcels always fit after emptying the first-mile parcels. The service level of the first-mile parcel stream is lower because this stream is bigger in absolute numbers and therefore there are more customers that wish to visit the wall but find no locker available.

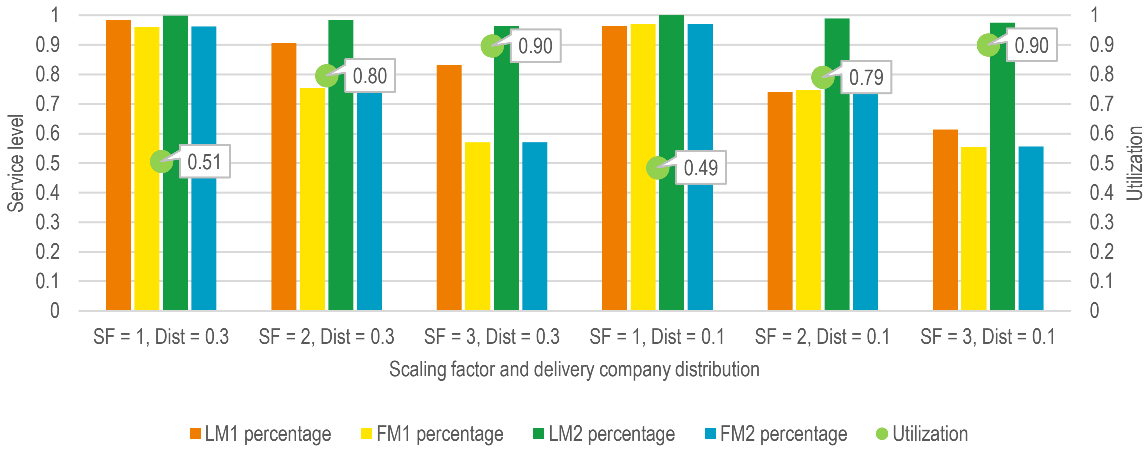

Figure 34 - Six experiments in which scenarios are simulated where the share of other delivery companies is bigger

8.4 DELIVERY COMPANIES SPREAD THROUGHOUT THE DAY

In this section, we will analyse what the effect would be of spreading the delivery companies more during the day. To investigate this, we will analyse six experiments. First, we will analyse the current situation with a scaling factor equal to 1 and 3 using the FCFS policy. In this situation, the first delivery company arrives at 10:00 and the second delivery company at 12. Thereafter, we will analyse what the effect would be if the second delivery company would arrive at 14:00, 16:00 and 18:00 for both scaling factors. The experimental settings are presented in Table 13.

Table 13 - Experimental settings for the experiments focused on spreading the delivery companies throughout the day

Exp.	Scaling factor	Arrival time delivery company 1	Arrival time delivery company 2
1	1	10:00	12:00
2	1	10:00	14:00
3	1	10:00	16:00
4	1	10:00	18:00
5	3	10:00	12:00
6	3	10:00	14:00
7	3	10:00	16:00
8	3	10:00	18:00

In Figure 35, the results of the experiments in terms of total profit are shown. On the left, the results of the experiments with a scaling factor equal to 1 are shown. It can be seen, the total profit increases slightly but the values are really close to each other. The p-values of the T-tests between the experiments showed that experiment 1, 3, 5 and 7 do not have different means with 95% confidence. Therefore, the differences seen in the left graph of Figure 35 are not significant. However, the differences in mean between the experiments with a scaling factor of 3 (experiments 2, 4, 6 and 8) are different. In the right graph in Figure 35 we can see that the total profit increases if delivery company 2 arrives later until 16:00 and after that decreases again. When delivery company 2 arrives around 12:00, the profit is approximately equal to 6900. When the delivery company arrives four hours later, the total profit increases to almost 7500. The other two experiments also have a higher total profit as result.

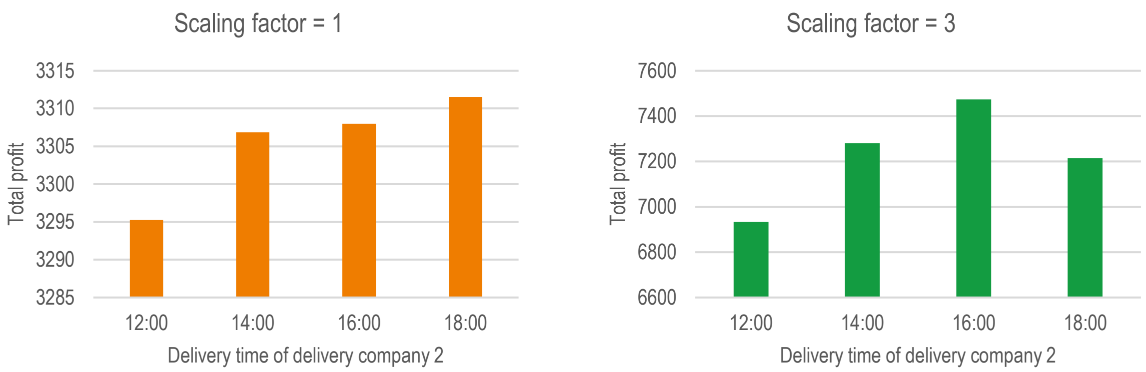

Figure 35 - Results in total profit of the experiments with other delivery times for delivery company 2

The total profit that is obtained by shifting the second delivery company to 16:00 in the afternoon is equal to 7473. The highest reward that was obtained by any of the implemented policies was equal to 6892. Shifting the second delivery time to 16:00 in the afternoon outperforms all policies and will also outperform the current situation.

In Figure 36, the service levels are presented of the four experiments with a scaling factor of 1. As can be seen, the service levels do not seem to vary. The small differences seem larger due to the small range on the y-axis. The lack in difference was to be expected as we already saw that there was no significant difference in total profit.

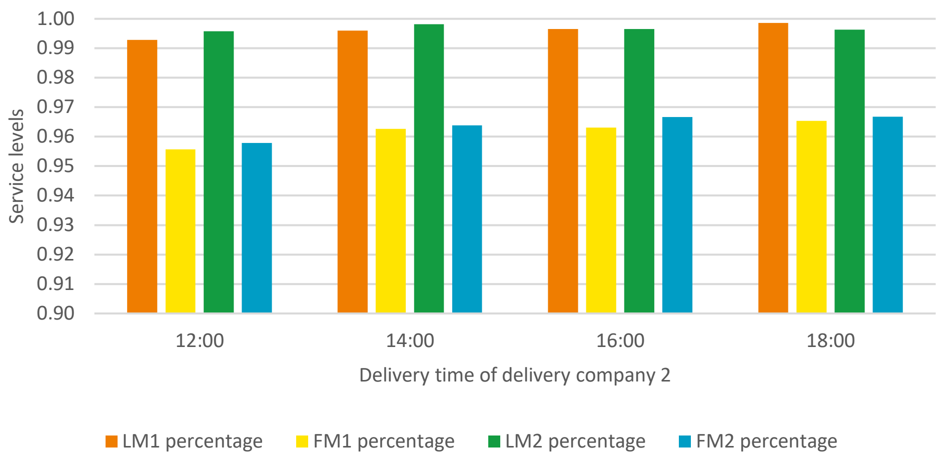

Figure 36 – The four service levels different delivery times of delivery company 2 with a scaling factor of 1

In Figure 37, the results of the experiments 2, 4, 6 and 8 are presented. In these experiments, the scaling factor was equal to 3. It can be seen that the last-mile service level of the delivery company 1 approaches 0.95 when delivery company 2 arrives later. This seems counterintuitive but can be explained by the fact that all first-mile parcels of delivery company that are delivered during the day are collected and therefore do not occupy lockers that can be used by delivery company 1.

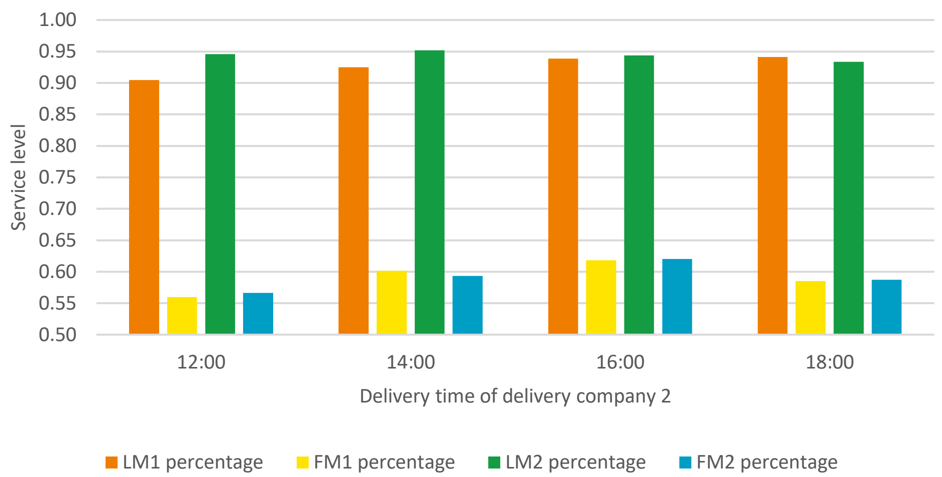

Figure 37 - The four service levels different delivery times of delivery company 2 with a scaling factor of 3

To make a comparison of the effect of the spreading of the delivery companies, we have put the total profit and utilization of the best-performing policies next to the results we obtained here. We analyse the results that are obtained in the scenario with a scaling factor of 3. The results are shown in Table 14. It can be seen that with the normal cost and income configuration, changing the second delivery time yields more profit than implementing any other policy.

Table 14 - The results in terms of total profit and utilization of the policies and the scenario with a later delivery time for delivery company 2

Policy	Total profit	Utilization
FCFS	6933	0.90
FCFS (delivery time DC2 16:00)	7473	0.86
Myopic [0.1] method	6598	0.87
MDP	6892	0.86

8.5 FROM WHICH UTILIZATION SHOULD EXPANSION BE CONSIDERED?

In this section, we will analyse the average utilization of the locker wall for different situations. We will first analyse the normal situation. This means the scaling factor will be equal to 1 and the FCFS policy will be used. Thereafter, we will analyse more crowded situations and increment the scaling factor with 0.5 each experiment. In Figure 38, the results of the experiments in terms of average utilization are presented. As can be seen, the utilization in the normal situation is equal to 0.46. Thereafter, it increases rapidly to 0.66 (scaling factor = 1.5) and 0.73 (scaling factor = 2). Subsequently, the average utilization increases less rapidly. Over the last four experiments it only increases with 0.04 in total.

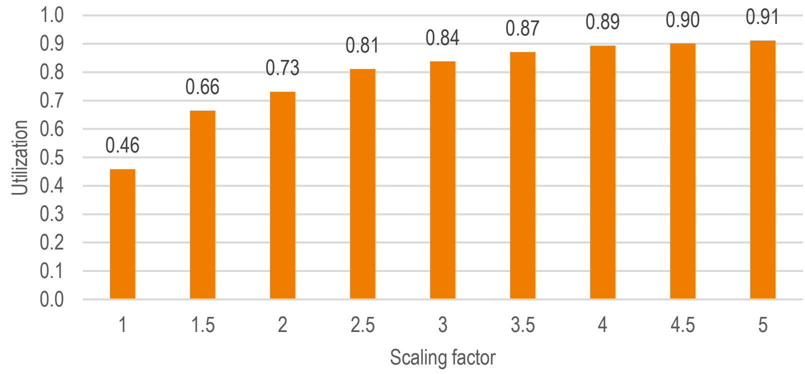

Figure 38 - The average utilization for situations with nine different scaling factors

In Figure 39, the results of the nine experiments are shown in terms of service levels. As can be seen, all service levels exceed the 0.9 in the normal situation. Evidently, the service levels decrease when the scaling factor increases. However, the first-mile service levels decrease more rapidly than the last-mile service levels. The last-mile service levels are equal to 0.9 or higher with a scaling factor between the 1 and 3.5. However, the service levels of the first-mile parcel streams are only higher than 0.9 in the first experiment, which represents the normal situation. Thereafter, the service level of the first-mile parcel streams decrease. When the scaling factor is higher than 3.5, the service levels do not exceed the 0.5 anymore.

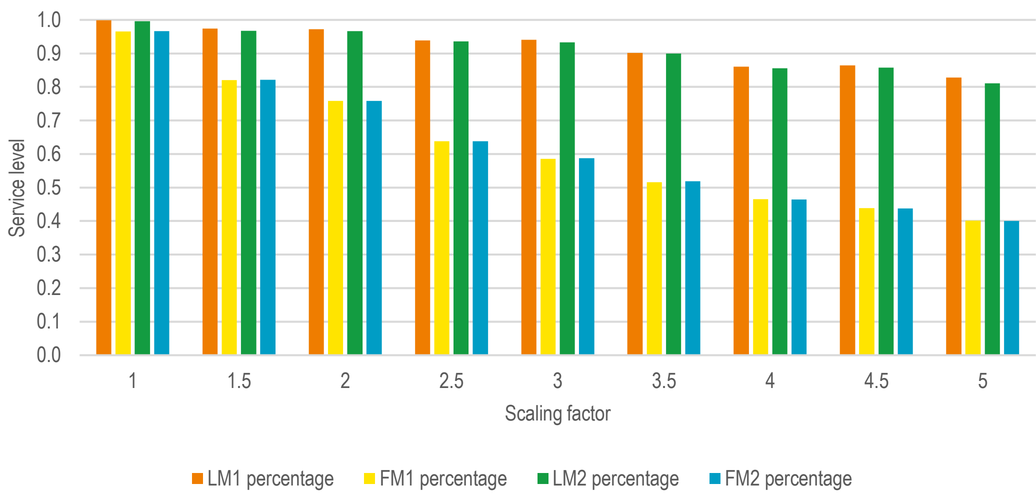

Figure 39 - The service levels of the four different parcel streams for situations with nine different scaling factors

Looking at Figure 38 and Figure 39, it can be seen that if the company aims to obtain a service level of at least 0.9 for all four parcels streams, the utilization should be somewhere between 0.46 and 0.66 maximum. To find out where the turning point is exactly, five more experiments are executed with scaling factors between 1 and 1.5 incrementing with a step of 0.1. The experiments and resulting service levels and average utilization are shown in Table 15.

Table 15 - The results in service levels and average utilization for six situations with different scaling factors between 1.0 and 1.5

Exp.	Scaling factor	LM1 service level	FM1 service level	LM2 service level	FM2 service level	Average utilization
1	1	1.00	0.97	1.00	0.97	0.46
2	1.1	1.00	0.96	1.00	0.96	0.46
3	1.2	0.99	0.93	0.99	0.93	0.54
4	1.3	0.99	0.93	0.99	0.93	0.54
5	1.4	0.99	0.88	0.98	0.88	0.60
6	1.5	0.97	0.82	0.97	0.82	0.66

In Table 15, it can be seen that the last time all service levels are above the 0.9 is in experiment 4. In this experiment, the scaling factor is equal to 1.3 and the average utilization is equal to 0.53. This means that when the average utilization of the locker wall is between 0.55 and 0.60, a service level of 0.9 cannot be guaranteed in these circumstances. Therefore, from an average utilization between 0.55 and 0.60 expansion of the locker wall should be considered.

9. CONCLUSION AND DISCUSSION

In this chapter, we will reflect on the results and discuss their meaning. First, we will look back on the research objective and questions (Section 1.4 and 1.5) and draw conclusions in Section 9.1. Subsequently, we will discuss the results and conclusions in Section 9.2. Thereafter, we will discuss interesting research directions in this area in Section 9.3.

9.1 CONCLUSION

We will first look back at the first chapters of our research and conclude these findings in Section 9.1.1. Thereafter, we will look back at our results and interpret them in Section 9.1.2.

9.1.1 Conclusion of research environment

In this research, we have found, developed, implemented and analysed different policies to allocate lockers of a locker wall used for parcel distribution. First, we discovered that the parcel locker environment consists of different actors (Figure 2) and there are four different parcel streams. Namely, the last-mile parcel delivery, the last-mile parcel pickup, the first-mile parcel delivery and the first-mile parcel pickup. There are two important actors regarding the four parcel streams: the customer and the delivery company. The delivery company delivers the last-mile parcels (between 11:00 and 13:00) and picks up the first-mile parcels that are dedicated to him. The customer picks up the last-mile parcels (throughout the day) and brings away first-mile parcels that must be picked up by the delivery company. Because the delivery companies all arrive relatively short after each other, the locker wall can become full. When a parcel cannot be delivered to the locker wall, a direct income is missed because the parcel cannot be accepted. Furthermore, goodwill is lost on two sides. First, the customer needs to pick up his parcel on another location than he wished. When this happens too often, customers may not choose the locker wall as a pickup location the next time. Second, the delivery company must visit another drop-off point to deliver the parcels that do not fit.

Therefore, we have investigated if there are policies that optimize the capacity of the lockers in the locker wall. Because it is not always possible to expand locker walls that are placed at busy locations, it may be valuable to have a method other than First-Come First-Serve that allocates the scarce resources.

To investigate this, we devised multiple research questions. First, we focused on the current situation and background knowledge of the locker wall environment. The first research question is answered in a number of key points.

What does the locker wall environment look like?

- The locker walls are located at locations that are convenient for customers and delivery companies. On the one hand, they should be near locations customer already pass (e.g. supermarkets, petrol stations, carpooling spots). On the other hand, they should be easily reachable by car so deliverers do not lose a lot of time visiting the locker walls.
- The locker wall consists of lockers in different sizes to accommodate all types of parcels. The number of lockers and configuration of the locker walls can be tailored to the specific needs or expectations of a location.
- There are four main actors. Namely, the delivery companies, the customers, the locker wall owners and the location partner. They all have different relationships with each other. If everything goes well, the delivery companies and the customers should mostly interact with each other. The owner of the locker wall company offers service if anything goes wrong and the location partner provides the location.

- The first-mile parcels are delivered by customers throughout the day. This follows the same pattern as the pickup process of last-mile parcels.
- The last-mile parcels are delivered by the delivery companies in the late morning and early afternoon. They also collect the first-mile parcels that are dedicated to them.
- The utilization of the locker wall increases steadily during the day. When the delivery company arrives, the first-mile parcels dedicated to him are emptied and the new last-mile parcels are delivered. Therefore, a V-shape is observed in the utilization of a locker wall at the arrival of a delivery company.

Subsequently, we investigated what is already known about using locker walls in last-mile parcel distribution. Furthermore, we investigated if there already exist other policies that cope with this challenge. This research question was answered and answered in some most important key points.

What literature exists concerning locker walls for parcel distribution and which relevant policies are described?

- The main strategical challenges in the locker wall industry concerns the placing and configuration of the locker walls. The location of the locker wall is very important as we already explained at the last question that it should be easy to use for customers and easy to reach for delivery companies. On a tactical level, a business model and policy should be chosen. This means the choice should be made which customers may use the wall in what capacity. Finally, on an operational level, the decision must be made to accept or reject a customer when he arrives at the locker wall. Another operational challenge can be determining the price for which a locker can be used.
- The problem of parcel locker allocation borders on more traditional revenue-management areas in the literature in which a lot of research has been done. Examples of this are plane seat reservation and car rental problems. What companies in these area tend to do is making a distinguishment in customer classes. Subsequently, one can make different agreements with different customers. The customer that wishes to receive the highest service levels pays a premium but penalty costs are incurred if this level is not met.
- Based on this customer classes system, a Markov Decision Process model was developed in the literature by Gans & Savin (2007). The model captures a system in which two types of customers arrive to a rental system with a fixed capacity and uncertain rental durations.
- A myopic allocation policy is also proposed in which contract customers are always accepted and walk-in customers are offered a fee that maximizes the discounted revenue from his rental. This is expected to work well if the demand and capacity are balanced.
- Other policies that are discussed are the so-called threshold policies. Threshold policies look at the state of the wall in terms of what percentage is full. If that percentage exceeds a certain threshold, no more parcels of specific types of customers are not accepted anymore.

Thereafter, we developed our own policies based on the ones that were found in the literature. A threshold policy was proposed. A policy that was based on the myopic allocation policy was proposed and the MDP policy was applied to this situation. Like Pazour & Roy (2015), we used a discrete-event simulation environment to analyse our policies in a locker wall environment. We implemented a simplified representation of a locker wall environment with 10 locker in the same size. The arrival rates of first-mile and last-mile customers were based on historical data of the locker wall company. Furthermore, the assumption was made that two equally big delivery companies deliver to the locker wall. One arrives at 10am in the morning and the other arrives at 12 in the afternoon. The performance of the simulation was measured based on the total profit and the service levels of the last-mile and first-mile parcel streams of both delivery companies.

9.1.2 Conclusion of results

The results of applying the different policies to the current situation showed that the current policy (FCFS) works best. This can be explained by the fact that the capacity of the locker wall is big enough for all the demand at this point. The myopic allocation policies perform equally well. This can be explained by the fact that all customers of delivery company 1 are accepted and the customers of delivery company 2 if it will not lead to not having enough space for the premium delivery company in the future with a prespecified probability. The threshold policies were less intelligent and perform worse. It can be seen, the threshold policy with the highest value performed best. This can be explained by the fact that the threshold policy with the highest value approached the FCFS policy. We will now analyse the results of the myopic allocation policies and the MDP policy. Finally, we will look back at the extra analyses that were done to see what the effects of other settings would be.

Myopic allocation policy

The myopic allocation policies focus on the service level of delivery company 1. During the first experimental settings, the income per parcel of the walk-in customers was higher than the income per parcel of delivery company 1. However, when there are not enough lockers available for the last-mile parcels of delivery company 1, a penalty is incurred. In the most crowded situation (a scaling factor equal to 3), the myopic [0.7] policy slightly underperformed compared to the FCFS policy. The results of the experiments are shown again in Table 16. The service levels of delivery company 1 are a lot better using the myopic allocation policy. However, the total profit of the other policy is still higher. This is explained by the fact that the penalties that are prevented using the myopic allocation policy do not outweigh the missed income of the denied customers of delivery company 2.

Table 16 - Results of the FCFS policy and the Myopic [0.7] policy in terms of total profit and service levels

	FCFS policy	Myopic [0.7] policy
Total profit	6933	6631
LM1 service level	0.90	0.98
FM1 service level	0.56	0.64
LM2 service level	0.95	0.91
FM2 service level	0.57	0.40

As we saw in Section 7.2.1, the myopic allocation policy will still not outperform the FCFS policy if the income per parcel is equal for both companies. However, if the parcels of delivery company 1 yield more than the parcels of delivery company 2, the policy starts to outperform the FCFS policy. A small difference was already made implementing the myopic [0.1] policy. When the income per parcel of delivery company 1 is equal to 0.7 and the income per parcel of delivery company 2 equal to 0.5, the myopic allocation policies all outperform the FCFS policy and especially the policies that build in more certainty (with a probability of 0.5 or higher). Using these settings with a myopic [0.7] policy with a lower penalty (equal to 0.4 per parcel that does not fit) leads to a result in total profit that is equal to when the FCFS policy is used. Therefore, we can conclude that with a difference of 0.2 in income per parcel between the delivery companies, it is valuable to implement the myopic allocation policy when the penalty is higher than 0.4. This was confirmed when the experiment with a penalty of 1.2 showed that the difference grew larger.

The policies focus on ensuring that the possibility that there are enough lockers available for delivery company 1 in the morning is equal to a prespecified probability. In order to ensure this, first-mile parcels of delivery company 2 may be rejected. The penalty that is incurred for not having enough space in the morning is equal to 0.8 per last-mile parcel. Furthermore, the income per parcel of delivery company 2 is equal to 0.7 and the income per parcel of delivery company 1 equal to 0.5. When first-mile parcels of delivery company 2 are rejected, a direct income of 0.7 is rejected. When the probability that is focused on is really high, e.g. equal to 1.0, a lot of last-mile parcels of delivery company 2 will be rejected to ensure no last-mile parcels of delivery company 1 have to be rejected. When the proportions are like this, rejecting parcels of delivery company 2 is costly. Namely, when a last-mile parcel of delivery company 1 does not fit in the morning, a direct income of 0.5 is missed and a penalty of 0.8 is incurred. So not having enough space costs 1.3 per parcel. If two last-mile parcels of delivery company 2 are rejected, this costs 1.4. Therefore, if, on average, more than 2 first-mile parcels are rejected to ensure there is space for a last-mile parcel, the myopic allocation policy will cost money. When the penalty becomes higher, or the difference in income smaller (or bigger in favour of delivery company 1), this ratio changes. It will become more worthwhile to reject first-mile parcels of delivery company 2 to save space for last-mile parcels of delivery company 1.

MDP policy

The MDP policy did not outperform the FCFS policy in the normal situation and the situation with a scaling factor equal to two. However, from a scaling factor equal to three it started to perform equally well and better. As we have seen in Section 7.2.2, the MDP policy starts to outperform the FCFS policy when the locker wall is really crowded. In the most crowded situation, no more first-mile parcels of delivery company 2 are accepted. It can be seen that this results in a service level of 0.99 for the last-mile parcels of delivery company 1. It is interesting to see that denying first-mile parcels of the second delivery company in really crowded situations has such a positive effect on the service levels of the first delivery company.

The solution of the MDP policy is strict. This means is does not allow a lot of first-mile parcels of delivery company 2 are accepted. For example, it could be seen in Figure 25 (scenario with scaling factor = 2) and Figure 27 (scenario with scaling factor = 3) that the last-mile service level of delivery company 1 was higher using the MDP policy than in the scenario that used the FCFS policy. However, the amount of missed direct income of rejected first-mile parcels of delivery company 2 does not outweigh the prevented penalties.

The MDP policy only focuses on maximizing the total profit based on the income per parcel of delivery company 1 and 2 and the penalty that are given. The MDP policy does not focus on obtaining a given service level. What can be seen is that the MDP policy is so strict in comparison with the FCFS policy that the utilization is much lower. In Table 17, the results of the experiments with the FCFS and MDP policies are shown for all three different scaling factors. First of all, it was seen that the total profit of the FCFS policy is higher but the difference decreases as the situation becomes more crowded. We saw in Section 7.2.2, that the MDP policy starts to outperform the FCFS policy in really crowded situations. In the three scenarios in Table 17 it can be seen the utilization using the MDP policy is approximately 0.04 lower. The difference in service level of the last-mile parcel stream of delivery company 1 is higher in all scenarios. The difference in the service level of first-mile parcels of delivery company 2 is bigger. In the normal situation, it is still equal to 0.77 but in the most crowded scenario it is only 0.46 (against 0.57 using the FCFS policy). This means that using the MDP policy will results in a lower utilization of the locker wall. It focuses on keeping lockers free for last-mile parcels of delivery company 1.

Table 17 - Results of the experiments in terms of total profit, utilization and the LM1 and FM2 service levels

Scaling factor	Policy	Total profit	Utilization	LM1 percentage	FM2 percentage
1	FCFS	3295	0.52	0.99	0.96
	MDP	3036	0.48	1.00	0.77
2	FCFS	5610	0.80	0.96	0.74
	MDP	5314	0.75	0.99	0.58
3	FCFS	6933	0.90	0.90	0.57
	MDP	6892	0.86	0.96	0.46

Conclusions of extra analyses

We will now discuss the extra experiments that were executed to see what the effect of certain choices or differences in environment. We will discuss them per point.

- *No more first-mile parcels:* without first-mile parcels the average utilization of the locker wall would be relatively low (0.16 normal situation and 0.25 with a scaling factor equal to 3). The service levels of the last-mile parcels would be equal to 1 in the normal situation. In more crowded situations, the service level of the delivery company that does not arrive first decreases. Thus, implementing this rule would lead to slightly higher service levels but also to a low utilization. At least as long as the delivery companies arrive at approximately the same time. The utilization and service level of delivery company 2 would probably increase if the second delivery company arrives in the late afternoon.
- *No more last-mile parcels:* in the normal situation this would lead to a utilization of 0.37 and both service levels would be equal to 1. In a more crowded situation (scaling factor = 2), this would already lead to service levels below the 0.90 and a utilization of 0.69. The average utilization in this situation is higher than the previous situation because the first-mile stream is bigger and the parcels are delivered throughout the day.
- *More "other" delivery companies:* in this situation, the proportion of parcels from the second delivery company was larger. If the proportion is 30/70 in the normal situation (scaling factor = 1), the service levels would be still between 0.95 and 1.00 for all parcel streams. However, if it becomes more crowded or the proportion moves to 10/90, the service level of the first delivery company decreases. This can be explained by the fact more first-mile parcels of the other delivery companies arrive and are accepted at the locker wall. This leads to lockers not being available when the first delivery company arrives and therefore it happens more often not all parcels fit in the locker wall.
- *Delivery companies more spread throughout the day:* in the normal situation (scaling factor = 1), the total profit increased when the arrival time of the second delivery company was later. No significant difference was found between different times (14:00, 16:00 and 18:00). However, in a more crowded situation (scaling factor = 3) the results in total profit were different. It could be seen that the total profit was highest when the arrival time of the second delivery company was 16:00. What is interesting to see is that when we look at the lambdas of the arrival rate in Section 4.2, the sum of the lambdas from 10:00 to 16:00 are approximately equal to the sum of the lambdas from 16:00 to 10:00 for the first-mile arrival rate as the last-mile arrival rate. This could clarify why this would be the optimal time to let the other delivery company arrive. Assuming that customers will behave the same after changing the time, they would be spread equally. Furthermore, we also saw that changing the delivery time of delivery company 2 would yield more extra profit than implementing any other policy. It is less costly to implement but may turn out to be more difficult. However, implementing this rule depends on external delivery companies and their routes. Explaining to them that their own service level will likely also be improved if they switch their arrival time may help.
- *From which utilization should expansion be considered?:* the utilization in the normal situation was equal to 0.46. The last-mile service levels are approximately equal to 1.00 and the service levels of the first-mile streams are also above the 0.95. When the scaling factor increases to 1.5, the utilization increases to 0.66. The service levels of the last-mile parcel streams are still above 0.95. However, the first-mile parcel stream service levels are only just above 0.80. From a scaling factor of approximately 3.5, the utilization does not increase fast anymore. A utilization around 0.90 seems the maximum. With a scaling factor of 3.5, the last-mile parcel

stream service levels are still 0.90 but the first-mile service levels are just above 0.50. In the extra experiments that were executed, it was seen the first-mile parcel streams are first below 0.90 when the utilization is equal to 0.60. We would therefore advise to consider expansion when the utilization is around that value. However, when the delivery companies are spread better a higher utilization can be used as reference.

9.1.3 General conclusion

During this research we have analysed the current situation and created different policies based on relevant literature to see if we could improve the current situation. The performance indicators we considered where the total profit, service levels and the utilization of the locker wall. The policies were evaluated in a simulation environment that represented a simplified version of reality. We have seen that in the normal situation, the implemented FCFS policy outperformed all other policies that were experimented with. All other policies focused on obtaining a high service level for delivery company 1 by rejecting parcels of delivery company 2. However, in all cases the income that was missed of rejecting first-mile parcels of delivery company 2 did not outweigh the penalties that were prevented for not having enough space for delivery company 1. In Section 7.2.1, we saw that the myopic allocation policies will outperform the FCFS policy if the income per parcel of the premium delivery company is higher than the income of the other delivery company. Therefore, it is good to ensure that the service that is promised to the premium delivery company outweighs the parcels that may be missed because of the service level that must be obtained. In crowded areas, it might be a good option to give one company a premium position and implement the myopic allocation policy. This will ensure that the service level of that company is sufficient and other companies and customers can still use the locker wall if there is space available. As said, this will only be profitable if the premium that is paid by the premium delivery company outweighs the missed income of the other delivery companies of which the parcels are rejected.

We also saw that in the current situation, changing the delivery time of the companies will outperform any policy that is focused on delivery company 1. As explained in the previous section, this will make sure that the mean value function of the arrival rates between the delivery companies is approximately equal to each other. This means that the number of first-mile and last-mile customer that arrive between the first and second delivery company is approximately equal to the number of customers that arrive between the second and the first delivery company. Therefore, this would balance the demand of the lockers and ensure a higher total profit can be obtained and higher service levels with the same number of lockers. The applicability of this solution depends on external parties, but it could be an interesting direction to further explore.

9.2 DISCUSSION

In this section, we will discuss our results and the conclusions. We will do that in two parts. First, we will discuss the usability of the results and conclusions of our research. Thereafter, we will discuss the limitations of our research.

9.2.1 Applicability of research

In this research, we aimed to find alternative policies to the First-Come First-Serve policy that is currently used. For this reason, we have developed and researched multiple policies. There are multiple outcomes from this research that may be useful in reality. First, the myopic allocation policies can be implemented in crowded situations in which the first-mile parcel stream of one delivery company should get priority over the other company. Based on the crowdedness and the importance of the first company, the strictness of the policy can be adjusted. In really crowded situations or with a major difference in income, a higher probability will lead to better results. The same holds if the difference in income is minimal and the penalty (in money or goodwill) of not having enough lockers available for the first company is high.

The Markov Decision Process policy that was developed did not outperform the FCFS policy in the situations with a scaling factor between 1 and 3. However, when the situation got even more crowded, it started to perform better relative to the FCFS policy. In the most crowded situation, no more first-mile parcels of delivery company 2 were accepted. However, this turned out to work quite well for the first delivery company. Therefore, it could be interesting to close off the possibility to deliver first-mile parcels for other delivery companies in really crowded walls while still allowing last-mile parcel delivery.

The conclusions of the extra analyses that were presented in Section 9.1.2 are based on the results from an abstract discrete-event simulation model. The effects of implementing these ideas can have the same effect in reality. However, before implementing or using them there should be thought about the differences between the simulation model and reality and the effect that these differences could have on the outcome of implementing these ideas. Two things that are mainly important to think about is the effect of having different sizes of lockers for different sizes of parcels. Another key point is that more than two delivery companies can use the white-label parcel locker walls. The effects could stay the same, but more delivery companies will also bring more first-mile parcels of different delivery companies to the wall. This could lead to a higher proportion of parcels in the wall that cannot be collected by individual delivery companies.

9.2.2 Limitations of research

The first limitation of this research is that it was difficult to get the data for a standard locker wall. Every locker wall stands on a unique location and a lot of factors play a role. Therefore, the arrival rates for every wall vary a lot. The proportions in which delivery companies visit locker walls also heavily depend on the location of the wall. Therefore, the simulation model and the lambdas that were used may not represent all locker walls throughout the Netherlands. However, the ideas may be helpful and implementable on comparable locations. In other situations, with walls that are not well represented by the model that was created in this research, the ideas still may be effective after adjustments are made. For this reason, we also showed results for situations with other delivery company proportions or more and less crowded situations. By showing how the service levels, utilization and total profit react to these kind of changes we have provided an idea of how the proposed solutions can be implemented in different situations.

One of the most things that may have influenced the results of this research, are the differences in the parameters of the different models and the simulation. An overview of the parameters of the models and simulation model are shown in Table 18. Every locker wall is different. The arrival rates of the walls that were studied at the start of the research represent the most standard situation. The inhomogeneous Poisson distribution makes it possible to cope with varying arrival rates during the day. However, correctly translating the empirical arrival rates into an inhomogeneous Poisson distribution is complex. First of all, the probabilities that are calculated in the myopic allocation policies are based on individual parcels. The assumption was made that the inhomogeneous Poisson distribution could be translated to calculate the probability that a specific parcel would be collected before the next morning. To this end, the mean value function in Equation 11 and the probability calculation in Equation 12 were used. The lambdas were multiplied with a factor of four to cope with the fact that normally only a quarter of the wall is filled with last-mile parcels. The MDP policy and the simulation are also based on inhomogeneous Poisson arrival rates. The arrival rates are scaled to lockers instead of parcels. For example, if there are four last-mile parcels in the locker wall with ten lockers, the arrival rates per locker are multiplied by ten and this is used to calculate the probability that a last-mile customer will arrive. The last-mile customer then randomly picks a last-mile parcel in the wall and leaves. If a last-mile customer arrives and no last-mile parcels are in the wall, he also leaves. In reality, this will never happen. Because last-mile parcel customers will only come to the wall if they are notified that their parcel is delivered. These differences make it difficult to compare the reality, policies and the discrete-event simulation. Furthermore, the decision moment of the myopic allocation policies and the discrete-event simulation differed from the decision moment of the MDP policy. To cope with and include the inhomogeneous Poisson arrival rates of the last-mile delivery companies and the first-mile customers, the decision had to be taken every hour. However, by decoupling the decision from the transition to another state, the MDP turned out to be very hard to solve. The idea behind it was that a chosen action could lead to the transition to multiple states, defined by the transition probabilities. So, an action did not ensure going to another state but ruled out going to a number of states. For example, choosing to only allow one first-mile customer the coming hour would set the transition probability of going to a state with more than one extra first-mile parcel to zero. Another limitation was that the program that was used solved the MDP really slowly. Because of this, few iterations were possible. Furthermore, it was not possible to implement a lot of different actions.

Table 18 - The differences between the input parameters of the models, discrete-event simulation and reality

	Focus	LM arrivals	FM arrivals	Number of LM parcels per batch	Decision moment
Reality	-	Empirical	Empirical	0.25 per locker on average	On event
Myopic allocation policy	Service levels	Inhomogeneous Poisson distribution scaled to parcels	Inhomogeneous Poisson distribution scaled to lockers	3 (10 lockers)	On event
MDP policy	Total profit	Inhomogeneous Poisson distribution scaled to lockers	Inhomogeneous Poisson distribution scaled to lockers	3 (10 lockers)	Hourly
Simulation	-	Inhomogeneous Poisson distribution scaled to lockers	Inhomogeneous Poisson distribution scaled to lockers	~U[0.75*2.5*scaling factor; 1.25*2.5*scaling factor]	On event

Another reason the MDP policy may behaved too strictly and thus rejected too many first-mile parcels of delivery company 2, is the cost configuration. The income of parcel of delivery company 1 is equal to 0.5 per parcel. The income per parcel of delivery company 2 is equal to 0.7. The amount that was lost in number of first-mile parcels that were rejected is not compensated by the penalties that were prevented. The MDP was focused on the policy that optimized the total profit. First-mile parcels that were accepted yielded an income of 0.5 and 0.7 for delivery company 1 and 2, respectively. Furthermore, when a terminal state is reached in which the available lockers (empty + first-mile parcels of delivery company 1) are smaller than the expected number of parcels, a penalty is given per locker short. The penalty is equal to 0.8 and consists of 0.5 missed income (parcel of delivery company 1) and 0.3 lost goodwill (0.25 of the customer and 0.05 of the delivery company). Because the horizon of the MDP setting ends when delivery company 1 arrives, the 0.5 income is included as missed income. However, the penalty cost is also used in the simulation. A penalty of 0.8 is incurred if a last-mile parcel of delivery company 1 cannot be delivered. However, the horizon of the simulation is 1000 days. Therefore, the parcel cannot be delivered and the time continues. Which means the penalty is incurred and the income is missed. Therefore, it may be argued that the penalty of missed income is incurred twice per parcel using the MDP policy. This could be a reason it is too strict.

Another point is the probability that was the central part of the myopic allocation policy. First, the policy would focus on a certain service level and the probabilities that were calculated with, were based on the service level. However, a service level is a long-term average. Given that in a lot of scenarios, all parcels can be accepted, the service level will always be higher on average. The myopic allocation policies now focused on obtaining a certain minimal service level. It was seen in the results, that the service level was almost always a lot higher than the probability that was focused on. Therefore, the myopic allocation policies could be adjusted so they focus on a service level on average instead of a minimum service level. This would allow the policy to be less strict and accept more first-mile parcels of delivery company 2. This would have a positive effect on the total profit and the service level of delivery company 2, while on the other hand the average service level to delivery company 1 can still be lived up to.

9.3 RECOMMENDATIONS FOR FUTURE RESEARCH

In this section we will discuss multiple directions for future research that we encountered during this research. First of all, it would be interesting to expand the current model and implement different locker and parcel sizes. This could, for example, show that the configurations of the locker walls could be adjusted to obtain a higher throughput on the same locker wall. For example, by changing bigger lockers for smaller ones. Besides this, research could be done into the seasonality of the demand at the locker walls. We have showed results for multiple scaling factors, by including the seasonality the policy that is used at the locker walls could include this. This would even make it possible to switch between different policies during the year. For example, during low season weeks/months the FCFS policy could be used whereas a myopic allocation policy could be implemented in the busier months to make sure the biggest clients still obtain a high service level.

Another point that is interesting to look at is extending the admission policy models with a dynamic pricing part. The same models can be used for determining the optimal pricing strategy. However, a price-response function should be devised. Based on the price-response function, the penalty per locker that is short, income per parcel and the probabilities that we used in the myopic allocation model, a dynamic pricing model can be applied to the locker wall environment. The revenue management model Xu & Li (2012) devised for computing capacity problems, a revenue management model could be made for the locker wall situation. The utilization can be used as a state and based on a price-response function, the price that maximizes the discounted profit can be offered. It is, however, difficult to say whether customers will visit the wall not knowing what it is going to cost. Therefore, it is important to not only research how dynamic pricing models could be designed for this purpose, but also in what way they should be implemented. Another way to overcome the price-response function that must be devised, is to use a reservation price like Farias & Roy (2010). That is simply enough a price that a customer from a certain customer class is willing to pay. If the price is beneath or equal to the reservation price, the offer is accepted and otherwise the customer leaves.

A third interesting direction to investigate would be the forecasting of the demand. A good forecasting method can serve as a basis for admission and dynamic pricing policies. Currently, the number of parcels that are missed due to locker walls that are too full is not kept track of. This could be investigated using censored demand estimation. Knowing this would give a better image on which bottlenecks to focus on.

An important part of the admission policy is the moment at which customers are accepted or denied. In this research, we assumed all decision are made at the locker wall. However, it may be desirable to already know at home if it will be worthwhile to drive to a locker wall, visit another parcel station, or drop off point. This can for example be done by giving customers insight into the real-time availability of the lockers. This can also be done by letting them indicate what type of parcel they wish to deliver and then say whether it is possible or not, without giving them all the information. This would also make it possible to implement admission policies as it will be difficult to say to a customer that he cannot use the wall after he saw there are lockers available.

10. BIBLIOGRAPHY

Altman, E., Jiménez, T., & Koole, G. (2001). On optimal call admission control in a resource-sharing system. *IEEE Transactions on Communications*, *49*(9), 1659–1668. https://doi.org/10.1109/26.950352

Angelopoulos, A., Gavalas, D., Konstantopoulos, C., Kypriadis, D., & Pantziou, G. (2016). An optimization model for the strategic design of a bicycle sharing system: A case study in the city of Athens. *ACM International Conference Proceeding Series*. https://doi.org/10.1145/3003733.3003784

Banks, J., Nelson, B. L., Carson, J. S., & Nicol, D. M. (2009). *Discrete-Event System Simulation*. Retrieved from https://www.researchgate.net/publication/293639861_Discrete-Event_System_Simulation

Biesinger, B., Hu, B., Stubenschrott, M., Ritzinger, U., & Prandtstetter, M. (2017). Optimizing charging station locations for electric car-sharing systems. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, *10197 LNCS*, 157–172. https://doi.org/10.1007/978-3-319-55453-2_11/TABLES/3

Bitran, G., & Caldentey, R. (2003). An Overview of Pricing Models for Revenue Management. *Manufacturing and Service Operations Management*, *5*(3), 203–229. https://doi.org/10.1287/MSOM.5.3.203.16031

Carroll, W. J., & Grimes, R. C. (1995). Evolutionary Change in Product Management: Experiences in the Car Rental Industry. *Interfaces*, *25*(5), 84–104.

Davydenko, I. Y., & Hopman, W. M. M. (2020). *Effect of pick-up points and returns on CO2 emissions in last mile parcel delivery networks*. https://doi.org/10.13140/RG.2.2.17054.56648

Den Boer, A. V. (2015). Dynamic pricing and learning: Historical origins, current research, and new directions. *Surveys in Operations Research and Management Science*, *20*(1), 1–18. https://doi.org/10.1016/J.SORMS.2015.03.001

Deutsch, Y., & Golany, B. (2018). A parcel locker network as a solution to the logistics last mile problem. *International Journal of Production Research*, *56*(2), 251–261. https://doi.org/10.1080/00207543.2017.1395490

Elmaghraby, W., & Keskinocak, P. (2003). Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science*, *49*(10), 1287–1309. https://doi.org/10.1287/MNSC.49.10.1287.17315

eMarketer. (2017). • Digital buyers worldwide 2021 | Statista. Retrieved April 26, 2022, from https://www.statista.com/statistics/251666/number-of-digital-buyers-worldwide/

Farias, V. F., & Roy, B. Van. (2010). Dynamic Pricing with a Prior on Market Response. *OPERATIONS RESEARCH*, *58*(1), 16–29. https://doi.org/10.1287/opre.1090.0729

Gallego, G., & Van Ryzin, G. (1994). Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons. *Management Science*, *40*(8), 999–1020. https://doi.org/10.1287/MNSC.40.8.999

Gans, N., & Savin, S. (2007). Pricing and capacity rationing for rentals with uncertain durations. *Management Science*, *53*(3), 390–407. https://doi.org/10.1287/MNSC.1060.0651/SUPPL_FILE/MNSC.1060.0651-SM-EC.PDF

Gosavi, A., Bandla, N., & Das, T. K. (2002). A reinforcement learning approach to a single leg airline revenue management problem with multiple fare classes and overbooking. *IIE Transactions (Institute of Industrial Engineers)*, *34*(9), 729–742. https://doi.org/10.1023/A:1015583703449

Heerkens, H., & Van Winden, A. (2017). *Solving Managerial Problems Systematically* (1st ed.). Noordhoff.

Huseby, A. B. (2021). STK2130 [Powerpoint slides]. https://www.uio.no/studier/emner/matnat/math/STK2130/v21/forelesningsnotater/lecture10_part2.pdf

Iwan, S., Kijewska, K., & Lemke, J. (2016). Analysis of Parcel Lockers' Efficiency as the Last Mile Delivery Solution – The Results of the Research in Poland. *Transportation Research Procedia*, *12*, 644–655. https://doi.org/10.1016/J.TRPRO.2016.02.018

Jain, A., Moinzadeh, K., & Dumrongsiri, A. (2015). Priority Allocation in a Rental Model with Decreasing Demand. *Manufacturing & Service Operations Management*, *17*(2), 236–248. https://doi.org/10.1287/MSOM.2014.0515

Kaufman, J. S. (1981). Blocking in a Shared Resource Environment. *IEEE Transactions on Communications*, *29*(10), 1474–1481. https://doi.org/10.1109/TCOM.1981.1094894

Kissell, R., & Poserina, J. (2017). Advanced Math and Statistics. In *Optimal Sports Math, Statistics, and Fantasy* (pp. 103–135). Academic Press. https://doi.org/10.1016/B978-0-12-805163-4.00004-9

Lazov, I. (2017). Profit management of car rental companies. *European Journal of Operational Research*, *258*(1), 307–314. https://doi.org/10.1016/J.EJOR.2016.08.064

Liu, C., Wang, Q., & Susilo, Y. O. (2019). Assessing the impacts of collection-delivery points to individual's activity-travel patterns: A greener last mile alternative? *Transportation Research Part E: Logistics and Transportation Review*, *121*, 84–99. https://doi.org/10.1016/J.TRE.2017.08.007

Macioszek, E. (2018). First and last mile delivery - problems and issues. *Advances in Intelligent Systems and Computing*, *631*, 147–154. https://doi.org/10.1007/978-3-319-62316-0_12

Molin, E., Kosicki, M., & Van Duin, R. (2022). Consumer preferences for parcel delivery methods. *European Journal of Transport and Infrastructure Research*, *22*(2), 183–200. https://doi.org/10.18757/EJTIR.2022.22.2.6427

Örmeci, E. L., & Burnetas, A. (2004). Admission control with batch arrivals. *Operations Research Letters*, *32*(5), 448–454. https://doi.org/10.1016/j.orl.2004.01.001

Örmeci, E. L., & Burnetas, A. (2005). Dynamic admission control for loss systems with batch arrivals. *Advances in Applied Probability*, *37*(4), 915–937. https://doi.org/10.1239/AAP/1134587747

Papier, F., & Thonemann, U. W. (2008). Queuing Models for Sizing and Structuring Rental Fleets. *Transport Science*, *42*(3), 302–317. https://doi.org/10.1287/TRSC.1070.0225

Papier, F., & Thonemann, U. W. (2010a). Capacity Rationing in Rental Systems with Two Customer Classes and Batch Arrivals. *SSRN Electronic Journal*, 73–85. https://doi.org/10.2139/ssrn.1567605

Papier, F., & Thonemann, U. W. (2010b). Capacity rationing in stochastic rental systems with advance demand information. *Operations Research*, *58*(2), 274–288. https://doi.org/10.1287/OPRE.1090.0745/SUPPL_FILE/OPRE.1090.0745-SM-EC.PDF

Pashenkova, E., Rish, I., & Dechter, R. (1996). Value iteration and policy iteration algorithms for Markov decision problem.

Pazour, J. A., & Roy, D. (2015). Analyzing rental vehicle threshold policies that consider expected waiting times for two customer classes. *Computers & Industrial Engineering*, *80*, 80–96. https://doi.org/10.1016/J.CIE.2014.10.030

Puterman, M. L. (1990). Chapter 8 Markov decision processes. In *Handbooks in Operations Research and Management Science* (Vol. 2, pp. 331–434). Elsevier. https://doi.org/10.1016/S0927-0507(05)80172-0

Robinson, S. (2014). *Simulation; The Practice of Model Development and Use*.

Rohmer, S., & Gendron, B. (2020). A Guide to Parcel Lockers in Last Mile Distribution-Highlighting Challenges and Opportunities from an OR Perspective.

Satır, B., Erenay, F. S., & Bookbinder, J. H. (2018). Production, Manufacturing and Logistics Shipment consolidation with two demand classes: Rationing the dispatch capacity. *European Journal of Operational Research*, *270*, 171–184. https://doi.org/10.1016/j.ejor.2018.03.016

Savin, S. V., Cohen, M. A., Gans, N., & Katalan, Z. (2005). Capacity management in rental businesses with two customer bases. *Operations Research*, *53*(4), 617–631. https://doi.org/10.1287/OPRE.1040.0193/SUPPL_FILE/OPRE.53.4.617-SM.PDF

Taboga, Marco (2021). "Combinations", Lectures on probability theory and mathematical statistics. https://www.statlect.com/mathematical-tools/combinations.

Vakulenko, Y., Hellström, D., & Hjort, K. (2018). What's in the parcel locker? Exploring customer value in e-commerce last mile delivery. Journal of Business Research, 88, 421–427. https://doi.org/10.1016/J.JBUSRES.2017.11.033

Wang, W. H., & Pinsky, E. (1989). An asymptotic analysis of complete sharing policy. *IEEE INFOCOM '89, Proceedings of the Eighth Annual Joint Conference of the IEEE Computer and Communications Societies*.

Whitt, W. (1991). The Pointwise Stationary Approximation for Mt/Mt/s Queues Is Asymptotically Correct As the Rates Increase. *Management Science*, *37*(3), 307–314. https://doi.org/10.1287/MNSC.37.3.307

Wu, P., Hartman, J. C., & Wilson, G. R. (2005). An integrated model and solution approach for fleet sizing with heterogeneous assets. *Transportation Science*, *39*(1), 87–103. https://doi.org/10.1287/TRSC.1030.0050

Xu, H., & Li, B. (2012). Maximizing revenue with dynamic cloud pricing: The infinite horizon case. *IEEE International Conference on Communications*, 2929–2933. https://doi.org/10.1109/ICC.2012.6364013

Yang, T. H., Lin, J. R., & Chang, Y. C. (2010). Strategic design of public bicycle sharing systems incorporating with bicycle stocks considerations. *40th International Conference on Computers and Industrial Engineering: Soft Computing Techniques for Advanced Manufacturing and Service Systems, CIE40 2010*. https://doi.org/10.1109/ICCIE.2010.5668312

Yuen, K. F., Wang, X., Ma, F., & Wong, Y. D. (2019). The determinants of customers' intention to use smart lockers for last-mile deliveries. *Journal of Retailing and Consumer Services*, *49*, 316–326. https://doi.org/10.1016/J.JRETCONSER.2019.03.022

APPENDIX A – CONFIDENCE INTERVAL METHOD

LM1							
N	**KPI**	**Average**	**St. Dev.**	**T-value**	**Delta**	**Error**	**Runs**
1	0.87942	0.879424	0	0	0	0	-
2	0.87495	0.877188	0.003162	12.7062	0.028411	0.032389	Sufficient
3	0.87713	0.877169	0.002236	4.302653	0.005555	0.006333	Sufficient
4	0.87215	0.875916	0.003102	3.182446	0.004936	0.005635	Sufficient
5	0.86970	0.874673	0.003865	2.776445	0.0048	0.005487	Sufficient

LM2							
N	**KPI**	**Average**	**St. Dev.**	**T-value**	**Delta**	**Error**	**Runs**
1	0.96117	0.961169	0	0	0	0	-
2	0.96602	0.96359	0.003427	12.7062	0.030792	0.031956	Sufficient
3	0.96677	0.96465	0.003039	4.302653	0.007549	0.007826	Sufficient
4	0.94657	0.96013	0.009376	3.182446	0.014919	0.015539	Sufficient
5	0.96720	0.96154	0.008714	2.776445	0.01082	0.011252	Sufficient

FM1							
N	**KPI**	**Average**	**StDev**	**Tstatistic**	**Delta**	**Error**	**Runs**
1	0.82964	0.829644	0	0	0	0	-
2	0.80928	0.81946	0.014403	12.7062	0.129404	0.157914	-
3	0.81490	0.81794	0.01052	4.302653	0.026132	0.031948	Sufficient
4	0.81686	0.81767	0.008606	3.182446	0.013694	0.016748	Sufficient
5	0.79444	0.81302	0.012784	2.776445	0.015873	0.019523	Sufficient

FM2							
N	**KPI**	**Average**	**StDev**	**Tstatistic**	**Delta**	**Error**	**Runs**
1	0.83719	0.837186	0	0	0	0	-
2	0.81080	0.82399	0.018656	12.7062	0.167619	0.203422	-
3	0.80265	0.81688	0.018053	4.302653	0.044846	0.054899	-
4	0.80926	0.81498	0.015224	3.182446	0.024225	0.029724	Sufficient
5	0.79682	0.81134	0.015483	2.776445	0.019224	0.023694	Sufficient

APPENDIX B – EXPERIMENT SETTINGS

Exp.	Scaling factor	Policy	# Exp.	Scaling factor	Policy
1	1	FCFS	19	2	Myopic policy [0.1]
2	1	Threshold [0.1]	20	2	Myopic policy [0.3]
3	1	Threshold [0.3]	21	2	Myopic policy [0.5]
4	1	Threshold [0.5]	22	2	Myopic policy [0.7]
5	1	Threshold [0.7]	23	2	Myopic policy [0.9]
6	1	Threshold [0.9]	24	2	MDP
7	1	Myopic policy [0.1]	25	3	FCFS
8	1	Myopic policy [0.3]	26	3	Threshold [0.1]
9	1	Myopic policy [0.5]	27	3	Threshold [0.3]
10	1	Myopic policy [0.7]	28	3	Threshold [0.5]
11	1	Myopic policy [0.9]	29	3	Threshold [0.7]
12	1	MDP	30	3	Threshold [0.9]
13	2	FCFS	31	3	Myopic policy [0.1]
14	2	Threshold [0.1]	32	3	Myopic policy [0.3]
15	2	Threshold [0.3]	33	3	Myopic policy [0.5]
16	2	Threshold [0.5]	34	3	Myopic policy [0.7]
17	2	Threshold [0.7]	35	3	Myopic policy [0.9]
18	2	Threshold [0.9]	36	3	MDP

APPENDIX C – SIMULATION PROGRAM