MASTER THESIS

Strategic Optimisation of Parcel Locker Placement in Out-of-Home Delivery Networks: A Case Study at DHL eCommerce Benelux

Industrial Engineering and Management

Production and Logistics Management
Supply Chain and Transportation Management

This is the public version of the thesis. Certain data has been modified or omitted, some scales have been removed from figures, and some figures have been excluded entirely. A full confidential version may be requested from the author, subject to approval by DHL eCommerce Benelux.

Author:
M.N.G. (Max) Miedema

University Supervisors
Prof. Dr. Ir. M.R.K. Mes (Martijn)
Dr. D.R.J. Prak (Dennis)

Company Supervisors
Ir. I. Menger (Ilse)

July 2025

Management Summary

This report presents a study on optimising the placement and sizing of parcel lockers in the DHL eCommerce Benelux out-of-home delivery network. The research addresses the challenge of improving current locker planning decisions, which are currently based on intuition rather than data. This results in suboptimal network performance, including underutilised lockers and regional mismatches between locker locations, capacity, and demand. To support DHL's long-term strategy of expanding its parcel locker network cost-effectively while improving customer satisfaction, we formulate the central research question as follows:

> *"How can DHL eCommerce Benelux make data-driven strategic decisions regarding the placement and sizing of parcel lockers within its out-of-home network, to reduce operational costs and improve customer satisfaction?"*

To address this, the research begins with a context analysis that provides a structured understanding of DHL's out-of-home network, offering insights into its current operations, performance, and key implications for the design of a strategic optimisation framework.

These practical insights help shape the focus of the literature review, which explores related academic research. Although interest in out-of-home delivery is clearly growing, with a rise in publications on optimisation problems related to this in recent years, the absolute number of studies remains relatively limited. This highlights both the novelty of this underexplored research domain and the academic relevance of this study. This thesis proposes a framework that addresses key literature gaps through novel model features and dynamics. Moreover, it introduces a unique combination of novel and established modelling elements that, to the best of our knowledge, have not been integrated before into a single framework.

Building on these foundations, the research developed a Mixed-Integer Linear Programming (MILP) model, formalised as the Last Mile Capacitated Parcel Locker Location Problem (LMCPLLP). The model determines the optimal placement and sizing of parcel lockers while allocating demand across delivery modes in a cost-efficient manner. It accounts for real-world constraints such as heterogeneous locker types, fallback or alternative delivery options (e.g., service points and home delivery), historical pickup behaviour, and partial demand allocation. To increase realism, the model includes extensions such as robustness against local worst-case demand fluctuations and an adaptive pickup radius that scales with population density.

The model was operationalised using real DHL data and is implemented in a user-friendly graphical user interface that allows planners to configure input parameters and interactively visualise model outcomes. It is designed to support both greenfield network design and incremental expansion of DHL's network, based on decision areas specified as input, whether at city, regional, or national scale. To validate the model and provide managerial insights, seven experimental phases were conducted using real-world data from DHL's operational regions, each designed to evaluate a distinct aspect of the model.

Some key findings include:

- Parcel lockers emerge as the preferred delivery mode from a cost-optimisation perspective, with an average share of over 85% at current cost levels across the experiments, and remain dominant under moderate cost changes, confirming their strategic robustness.
- Cost reductions of 18–22% were achieved across test regions, demonstrating the model's potential in reducing operational costs through improved locker placement and sizing.
- The model solves efficiently for large-scale planning (0–4 minutes per DHL's RegioHub region) in DHL's preferred settings, supporting scalability across DHL's national network.
- As demand grows, or when accounting for local uncertainty through robust optimisation, locker locations remain relatively stable, with mainly locker sizes increasing. This indicates the model's long-term effectiveness in initial placement decisions.
- Robustness experiments suggest that addressing extreme, concentrated demand spikes tends to be more costly than mitigating multiple moderate, distributed fluctuations.

Based on these findings, some key research recommendations for DHL are to:

- Embed the optimisation model into its strategic network planning to replace intuition-based placement.

- Promote increased usage of parcel lockers, e.g., through incentives or default selection, to enhance cost-efficiency and reduce reliance on less scalable delivery modes.
- Integrate real-time locker availability into public interfaces (e.g., DHL website or app), so that senders can see availability before drop-off, reducing failed attempts and improving user satisfaction.
- Extend the dynamic capacity control system with real-time reallocation logic to improve overflow handling and network efficiency.
- Proactively place larger lockers at key locations to accommodate future demand and improve robustness if financially viable.
- Invest in maintaining service points within the network, as they serve as flexible and cost-effective buffers that enhance network resilience and help postpone abrupt investments in locker expansion.
- Prioritise fixed cost reductions of lockers to enable wider viability of parcel lockers in its out-of-home network, especially in low-density areas.
- Develop data-driven insight into the reasons behind parcel diversions to enable targeted improvements.

This thesis contributes to both academic literature and industry practice by introducing a scalable, data-driven optimisation framework for strategic parcel locker placement and sizing decisions in an out-of-home network. Together, these contributions provide DHL with a solid foundation for future network design and cost optimisation, supporting its long-term strategy of expanding the parcel locker network.

Preface

This thesis marks the conclusion of my Master's in Industrial Engineering and Management at the University of Twente, where I specialised in Production and Logistics, with a focus on Supply Chain and Transportation Management. I truly enjoyed applying operations research within the logistics sector during my studies. The completion of this thesis also marks the end of my time in Enschede, which has been a valuable period of personal, academic, and professional development. Over the past years, I have had the opportunity to meet a lot of great people and make new friends. I would like to take this moment to thank everyone who made my time in Enschede unforgettable.

To begin, I would like to thank DHL eCommerce Netherlands for giving me the opportunity to conduct my graduation project within the Operations Engineering & Innovation department in Utrecht. From day one, I was warmly welcomed, both by the team and the broader operations department. I had a great time working alongside kind and dedicated colleagues, and I learned a great deal from the many smart people and innovative ideas around me. A special thanks goes to my supervisor and manager at DHL, Ilse Menger. I was impressed by the time and effort you put into guiding me. Our meetings were not only helpful and insightful, but also very enjoyable. Your support played a major role in making this thesis possible and also contributed to a memorable time at DHL, where you made it possible for me to explore and be involved in many different areas of the operation.

Second, I would like to thank my university supervisors for both their valuable academic feedback and the more informal off-topic conversations that made the past six months more enjoyable. To start with, my first supervisor, Martijn Mes, thank you for the insightful input during our meetings. I often found myself quickly typing to keep up with all the interesting ideas and directions you came up with. I always left with new insights that helped me view the problem from fresh and smarter angles. Your input played a key role in shaping both the structure and the content of this thesis. I am also grateful for your help in finding this assignment, which aligned perfectly with my interests and made this entire period possible in the first place. I would also like to thank my second supervisor, Dennis Prak, whom I have known for quite some time. Besides your great support during this thesis, through valuable feedback and engaging discussions, I would also like to take this moment to thank you for the opportunity to work with you as a teaching assistant during my studies. I learned a lot from those teaching experiences and truly valued the time, guidance, good times, and laughter you consistently brought to our collaboration, making it a memorable and enjoyable part of my time at university.

Last but not least, I would like to thank all my friends, my girlfriend, and my family. Your support over the past months and throughout my student life has given me memories I can proudly look back on.

I hope you enjoy reading this thesis.

Max Miedema
Utrecht, 29th of June 2025

Acronyms

B2B Business to Business.

B2C Business to Consumer.

C2B Consumer to Business (returns).

C2C Consumer to Consumer.

CDF Cumulative Distribution Function.

CpP Cost per parcel.

ECDF Empirical Cumulative Distribution Function.

FLP Facility Location Problem.

FM First Mile.

GUI Graphical User Interface.

ILP Integer Linear Program.

LM Last Mile.

LMCPLLP Last Mile Capacitated Parcel Locker Location Problem.

MILP Mixed Integer Linear Program.

OOH Out of Home.

OOHD Out of Home Delivery.

PL Parcel Locker.

SP Service Point.

Contents

1 Introduction

This chapter provides an introduction to the research conducted at DHL eCommerce Benelux. Section 1.1 introduces DHL Group and outlines its operations. Section 1.2 explains the research context and highlights the relevance of the study within the logistics industry. Section 1.3 discusses the motivation behind initiating this research. Section 1.4 presents a problem analysis and identifies the core problem. Section 1.5 defines the research scope and process boundaries. Lastly, Section 1.6 formulates the main research question and outlines the design of this study.

1.1 Company Description

Section 1.1.1 introduces DHL Group, and Section 1.1.2 describes its regional division, DHL eCommerce Benelux.

1.1.1 DHL Group

DHL Group, formerly known as Deutsche Post DHL Group, is a global leader in the logistics industry. DHL Group provides a comprehensive range of international express deliveries, freight transport, e-commerce and supply chain management services. It operates in over 220 countries and territories, employs nearly 600,000 employees and generated €81.8 billion in revenue in 2023 (DHL Group, 2025b).

DHL was founded in 1969 in the United States. The company was named after the founders Adrian **D**alsey, Larry **H**illblom and Robert **L**ynn. DHL started off as DHL Worldwide Express as an international air express service, providing rapid transport of documents and cargo papers by plane. In 2002, Deutsche Post acquired DHL and combined its entire express and logistics business under the DHL brand (DHL Group, 2025b). Nowadays DHL Group consists of the following divisions;

- **DHL Express**: Specialises in international express deliveries, ensuring urgent documents and goods are transported reliably and on time from door to door, mainly using airplanes.
- **DHL Global Forwarding (DGF)/ Freight**: DGF offers international air and ocean freight services and manages complex global transportation solutions. Freight offers road and rail freight (intermodal transport) services across 50 countries and territories around the world.
- **DHL Supply Chain**: Specialises in contract logistics, including warehousing, transport, and value-added services that can be customised into full supply chain solutions.
- **DHL eCommerce**: Focuses on domestic and international parcel and pallet delivery, primarily serving the e-commerce sector.
- **Post & Paket Deutschland**: Operates under Deutsche Post, handling mail and parcel services within Germany.

1.1.2 DHL eCommerce Benelux

Within the Benelux, all divisions excluding the German 'Post & Paket Deutschland' are present. DHL began operations in the Netherlands in 1976 and expanded rapidly through strategic acquisitions. For example, the takeover of Van Gend & Loos in 1999 strengthened DHL's domestic distribution network, and the acquisition of Selektvracht in 2011 expanded DHL's e-commerce services. DHL's growth together with these strategic acquisitions positioned DHL as a leading logistics provider in the Netherlands. Meanwhile, DHL expanded its presence across the Benelux region. It has been operating in Belgium since 1978 and has created an extensive network across Belgium, the Netherlands, and Luxembourg (DHL Express Belgium, 2025). Today, DHL's e-commerce operations in the Netherlands are part of DHL eCommerce Benelux.

DHL eCommerce Benelux specialises in domestic and international parcel and pallet delivery within, to and from the Benelux countries. The division employs around 15,000 people and operates a network tailored to both individuals and businesses, covering Business to Consumer (B2C), Consumer to Business (returns) (C2B), Business to Business (B2B) and Consumer to Consumer (C2C) services (DHL eCommerce Netherlands, 2025; DHL Group, 2025a). The DHL eCommerce Benelux headquarters is located in Utrecht. In the remainder of this thesis, 'DHL' refers specifically to DHL eCommerce Benelux.

The division offers parcel and pallet shipping, both door-to-door and via Out of Home (OOH) points, such as Service Points (SPs) and the relatively new Parcel Lockers (PLs). Services are divided into:

- **Consumer Services**: Individuals can send parcels up to 20 kg domestically and across Europe, with shipments ranging from envelopes to parcels.
- **Business Services**: Businesses have tailored shipping options domestically and across Europe, with shipments ranging from parcel to pallet deliveries to customers or other businesses.

1.2 Research context

Section 1.2.1 highlights the increasing relevance of Out of Home Delivery (OOHD) in the logistics industry and Section 1.2.2 outlines the concept and formats of OOH points used by DHL.

1.2.1 Industry relevance

Parcel deliveries have increased significantly in recent years. In 2016, 64 billion parcels were shipped worldwide, growing to over 161 billion in 2022 and is expected to rise to 225 billion by 2028 (Pitney Bowes, 2023). Also, the COVID-19 pandemic has significantly accelerated the growth of e-commerce and parcel delivery services, particularly OOHD options (Reiffer et al., 2023). Moreover, in recent years, OOHD has emerged as a significant trend in last-mile logistics, offering potential solutions to challenges like labour shortages, rising costs, and environmental concerns (Janinhoff et al., 2024). OOHD could help logistics companies to reduce costs by potentially minimising failed deliveries and combining demand (Savelsbergh & Woensel, 2016; Song et al., 2009). Consumers are shifting towards OOH points for reasons such as fewer failed home deliveries, the flexibility to pick up packages at their convenience, sustainability considerations, or potentially lower costs.

1.2.2 OOH points

This research focuses on the processes involving OOH points. OOH points can be categorised into two types:

- **Parcel Locker (PL)**: Also known as DHL Lockers, these are self-service units enabling customers to send and receive parcels at their convenience, often accessible 24/7. Lockers are typically located at places such as shopping centres, supermarkets, gas stations, or sports clubs. These lockers operate with a limited capacity.
- **Service Point (SP)**: Staffed locations where customers can drop off or collect parcels. These are typically local businesses partnering with DHL to operate an official SP and in return receive compensation for handling each parcel. These locations, therefore, could be seen as having 'unlimited' or large capacity, since most locations have sufficient storage space and want to handle more parcels for more revenue. However, in peak season it may occur that a staffed location refuses parcels.

1.3 Research motivation

With the rise of e-commerce, DHL experienced a rapid increase in online orders starting in 2015/2016. This trend intensified during the COVID-19 pandemic (2020–2023), when online shopping became the primary purchasing method for many people. While the COVID-19 period led to extreme growth, DHL now expects a more stable increase in the coming years. To prepare for this, DHL rapidly expanded its OOH network to make sure that the necessary infrastructure is in place. Moreover, DHL observes that an increasing number of consumers are shifting towards OOHD options. Therefore, they also expect the use of OOH points to further increase compared to home delivery in the upcoming years, leading to even greater pressure on these points.

Until now, the focus has been on rapid rollout and speed, prioritising fast-paced expansion. However, as DHL enters a phase of more gradual growth, the company recognises the need for a strategic reassessment of its OOHD system. This includes evaluating and improving OOHD processes such as location planning, capacity allocation, delivery flows, operational policies, or user-facing interfaces related to OOH points. This forms the motivation for this research.

1.4 Problem Identification

This section identifies the core problem this research addresses regarding DHL's efforts to optimise its OOH operations. Section 1.4.1 presents the underlying causes and relationships through a structured problem cluster. Section 1.4.2 identifies the core problems, and Section 1.4.3 describes the selection of the core problem tackled in this study.

1.4.1 Problem cluster

Due to the broad scope of DHL's ambitions in optimising the complete OOH process, this section provides a structured analysis to identify a suitable core problem. After analysing DHL's OOH operations, observing the process in real life, conducting data analysis and having interviews with relevant stakeholders, a cause-effect relationship has been identified. This relationship is visualised in a problem cluster in Figure 1. This cluster serves as a structuring tool to better understand the problem context and determine the core problem. The action problems are highlighted in green, being the discrepancy between the norm and reality, as perceived by the problem owner (Heerkens & van Winden, 2017), in this research, DHL.

This research identified that DHL faces two main action problems related to the OOH processes: too high operating and transportation costs, and too low customer satisfaction. The high transport and operating costs stem from unnecessary kilometres travelled due to inefficiencies in route planning and capacity utilisation. Additionally, some PLs operate at a loss. The low customer satisfaction is primarily caused by capacity issues, as lockers have limited space. When lockers reach full capacity, receivers may face delays, rerouted deliveries, or be forced to collect their packages from a different location than originally selected, therefore negatively impacting the customer experience.

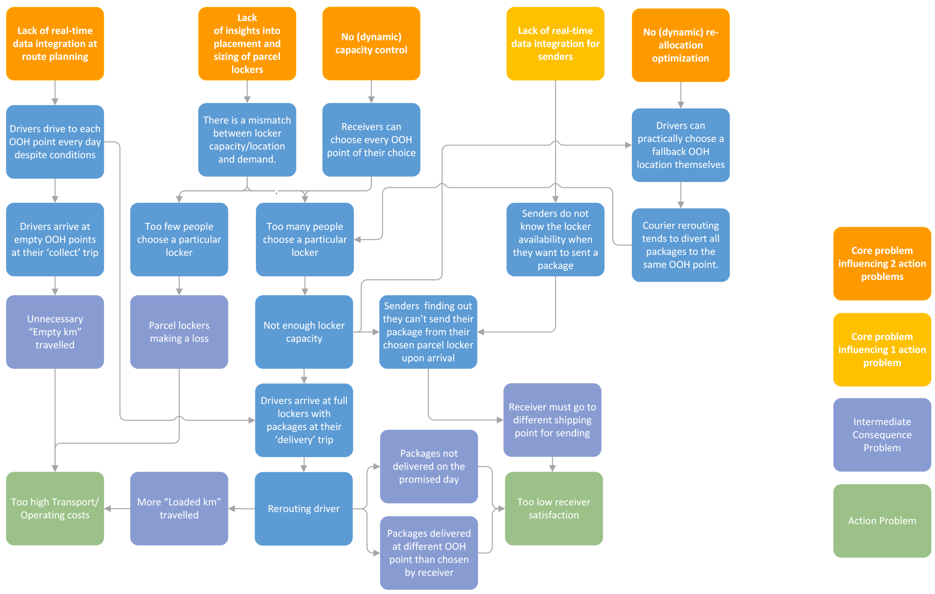

Figure 1: Problem cluster OOH processes DHL

1.4.2 Identification of core problems

Based on the problem cluster, four core problems were identified, each contributing to the action problems.

Lack of real-time data integration in route planning

Although DHL has access to real-time data on the current state at OOH points, this data is not integrated into the operational routing logic. As a result, if there are deliveries for a certain OOH point, the location will automatically be added to the courier's route, regardless of whether it has sufficient capacity to handle the deliveries. Moreover, all PLs are also visited later in the day to collect parcels that are shipped from the OOH points, even when no parcels are ready for collection at those locations. These inefficiencies lead to unnecessary transport costs and even (late) rerouting of drivers, causing packages to be delivered to different OOH points than originally selected by the receiver, or resulting in delays. This, in turn, contributes to both action problems.

There is a lack of strategic, data-driven insights into the optimal placement and sizing of PLs within DHL's OOH network

Currently, the sales team determines PL placements based on intuition rather than data-driven insights. There is no clear strategy on where and how many PLs should be placed. This leads to overutilised lockers in high-demand locations while others remain underutilised. This results in financial losses and capacity issues. Moreover, capacity problems can also lead to reroutings, causing packages to be delivered to different OOH points than chosen by the receiver or even resulting in delays. Thus again contributing to both action problems.

No (dynamic) capacity control for OOH points
Currently, customers can choose any OOH point without any capacity-based restrictions at checkout. There is no forecasting model that predicts, e.g., locker inflow and outflow, which could be used to optimise capacity and limit the OOH selection of receivers at their checkout to evenly distribute demand across OOH points. The lack of dynamic capacity control again leads to overutilised lockers in high-demand locations, while others remain underutilised, and thus in the same way contributing to both action problems.

Lack of real-time data integration for senders
Although real-time data on locker availability is available within DHL's system, it is not integrated into the public user interface. As a result, senders do not know locker availability before selecting a locker drop-off point. This lack of transparency could cause failed drop-offs, requiring senders to travel to another location upon arrival at the full locker, negatively impacting customer experience and satisfaction.

No (dynamic) re-allocation optimisation
At the moment, drivers can choose an alternative drop-off point when facing full OOH points, but there is no automatic (dynamic) reallocation process that considers customer proximity and capacity distribution. As a result, courier rerouting often leads to parcels being diverted to the same locations or to easily accessible points, such as SPs where capacity is unlimited and delivery is quicker. This again causes overloaded OOH points, while others remain underutilised. This creates a circular dependency, further increasing rerouting costs and contributing to failed deliveries. Therefore, this issue again contributes to both action problems.

1.4.3 Core problem selection

For this research, a choice must be made regarding the core problem to be addressed. Since real-time data integration primarily involves practical adjustments, it is considered more suitable as a recommendation for DHL rather than the main focus of this research.

Moreover, DHL is currently starting the process of investigating the implementation of (dynamic) capacity control for PLs. This includes forecasting inflows and outflows to optimise OOH point selection at checkout. With this project they aim to only display the 'expected' or 'forecasted' available OOH points to the receiver at the checkout. With this, DHL aims to better distribute capacity and prevent rerouting due to full capacity. As a result of this initiative, dynamic reallocation optimisation has become less relevant for this research, as rerouting should naturally decrease after this integration.

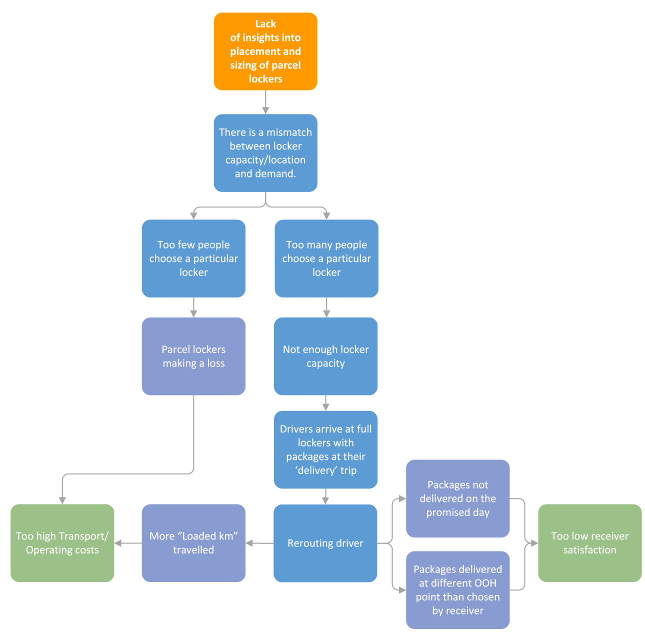

Figure 2: Problem cluster for this research

On the other hand, the lack of insight into strategic decisions regarding the placement and capacity sizing of PLs within DHL's OOH network remains a critical issue. As parcel lockers are a key focus for DHL's future expansion, DHL has indicated the importance of addressing this challenge. Given these considerations, this research focuses on the following core problem: *Lack of insights to support strategic decision-making on the placement and sizing of parcel lockers within DHL's OOH network*. Accordingly, the associated problem cluster can be refined to reflect this core focus and is visualised in Figure 2.

1.5 Research scope

DHL operates a complex logistics network, integrating multiple hubs with distinct functions to ensure efficient parcel and pallet delivery. The largest and most important of these networks is the B2C network. Figure 3 provides a simplified overview of this process, illustrating the main pathway for most e-commerce parcels.

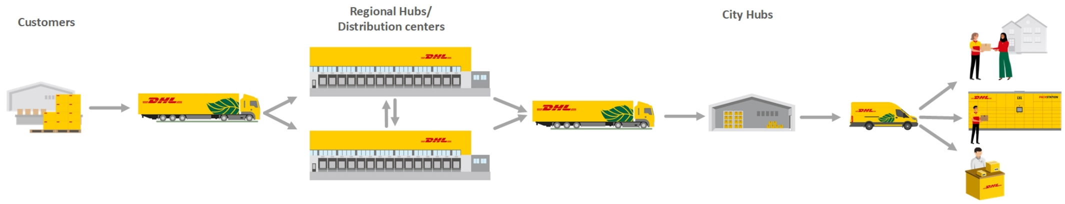

Figure 3: B2C network DHL

In this network, parcels are picked up from customers, such as large clothing retailers and warehouses, and follow a route through regional hubs and distribution centres before arriving at a city hub. From the city hub, they are distributed either via home delivery, PL, or SP.

This research focuses on the OOH process from the city hub to the final delivery point, with a particular focus on the strategic placement of PLs. However, within this process scope, two different parcel flow directions can be distinguished:

- **First Mile (FM)**: The process of sending a parcel via an OOH point, including returns (C2B), personal shipments (C2C) or business shipments (B2C). This is the initial stage of a parcel's journey before entering the distribution network of DHL.
- **Last Mile (LM)**: The process of delivering a parcel from the city hub to its final destination, for example, an OOH point. This applies to both B2C and C2C shipments.

Figure 4 shows the processes within the scope of this research and the distinction between FM and LM.

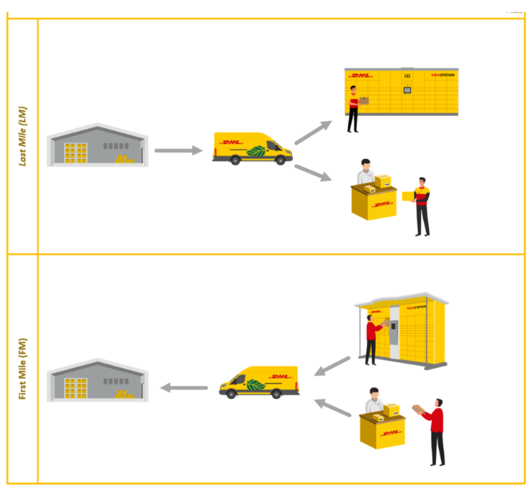

Figure 4: LM/FM research scope processes

1.6 Research design

To address the identified challenges, this section outlines the research approach. The main research question is derived from the problem analysis and is formulated as follows:

"How can DHL eCommerce Benelux make data-driven strategic decisions regarding the placement and sizing of PLs within its OOH network, to reduce operational costs and improve customer satisfaction?"

This research question addresses the core issue currently faced by DHL: *Lack of insights to support strategic decision-making on the placement and sizing of parcel lockers within DHL's OOH network.* While the OOH network has rapidly expanded in recent years, many placements have been made based on intuition or availability rather than quantitative performance criteria. With DHL planning to significantly grow its OOH network in the coming years, especially through a major expansion of PLs, this issue becomes even more relevant. Strategic, data-driven placement is essential to optimise the current network and align future expansion with demand and sustainable growth. To answer the main research question, we formulate a number of sub questions, each answered in a separate chapter.

Chapter 2, provides a structured analysis of DHL's current OOH network. It begins by examining the operational processes surrounding FM and LM activities. Then, it investigates the strategic, financial, and spatial factors influencing placement decisions. Finally, the chapter evaluates the network's current performance, focusing on geographical distribution, parcel turnover, customer behaviour, and capacity utilisation. Together, these insights form the foundation for identifying placement inefficiencies and informing a data-driven optimisation approach.

> **Chapter 2.** *What are the current operational processes, placement challenges, and performance characteristics of DHL's OOH network, and how can these insights inform a data-driven optimisation model?*
>
> ---
>
> Operational Context
> - What are the current FM and LM processes at OOH points for DHL?
>
> Placement Challenges
> - What strategic, operational, and spatial constraints affect the placement of PLs and SPs?
> - How do known volume thresholds and cost structures influence the financial viability of different OOH formats?
>
> Network Performance Insights
> - What is the current distribution of OOH points in terms of geography and capacity?
> - How do OOH locations perform in terms of turnover, customer pickup behaviour, and diversion patterns?
>
> Implications for Modelling
> - Which findings from this analysis have implications for the design of the optimisation model?

Building on the insights from the current situation analysis, Chapter 3 explores academic literature on facility location models for OOH delivery networks. The findings from Chapter 2 help define the practical requirements for a suitable optimisation approach. In turn, this chapter investigates how existing models from the literature can address these needs. It begins by positioning DHL's case within the broader field of Facility Location Problems (FLPs), focusing on models for PLs and SPs in both LM and FM contexts. An in-depth comparison of modelling characteristics is provided, which supports the selection of an appropriate modelling direction and reveals literature gaps which the model in Chapter 4 aims to address.

> **Chapter 3.** *Which methods for the strategic placement of OOH points are discussed in the literature?*
>
> ---
>
> - How can the OOH placement challenge be translated into a theoretical problem?
> - What is the current state-of-the-art regarding OOH network design?
> - What limitations regarding model realism, scalability, or applicability to DHL's OOH network remain in the literature, and how does this study contribute?

After identifying the challenges and current situation at DHL, the scientific gaps in the literature and insights from academic studies, Chapter 4 introduces a mathematical model for the strategic placement and capacity sizing of PLs in DHL's OOH network.

Chapter 4. *How can a mathematical model be developed to support strategic placement and sizing decisions for PLs in DHL's OOH network?*

- How can the facility location problem be formulated to incorporate the specific characteristics of DHL's PLs and SPs?
- Which model extensions can enhance the practical applicability and robustness of the proposed facility location model?

After developing the mathematical model, Chapter 5 describes how the model is implemented and prepared for experimentation. This includes the required data input and processing, parameter settings, and experimental design. Together, these elements form the foundation for evaluating the model in a realistic DHL context.

Chapter 5. *How should the developed model be implemented for real-world application, and which experimental setup should be used to evaluate it?*

- What data sources and preprocessing steps are required to support the model input?
- How should cost structures, demand estimates, and capacity settings be translated into parameters?
- How should the key model extensions be used within the model?
- How should test scenarios be designed to evaluate model performance, validate outcomes, assess robustness and sensitivity, and generate managerial insights regarding strategic PL placement?

Following the experimental setup, Chapter 6 presents the results of the conducted experiments. These results are analysed from several angles, including model validation, robustness, sensitivity, scalability, and their implications for DHL's strategic decision-making.

Chapter 6. *What are the outcomes of the experimental design?*

This chapter evaluates the experimental outcomes from multiple perspectives:
- Model validity and performance: How well does the model perform under realistic baseline scenarios, and are the results intuitive and in line with expectations?
- Robustness and sensitivity: How sensitive are the results to changes in parameters or settings?
- Scalability and applicability: How does the model scale and perform in different or larger configurations?
- Managerial relevance: What strategic insights can be derived for DHL regarding their PL placement and sizing decisions?

Deliverables After finishing the research and answering the main research question, the following deliverables are presented:

- A structured analysis of DHL's current OOH network, including operational processes, placement challenges, and performance analysis.
- A literature-based positioning of the PL placement problem within the Facility Location Problem domain.
- A mathematical model tailored to DHL's OOH context to support strategic decisions regarding the placement and sizing of PLs.
- An experimental framework to evaluate the model under various scenarios.
- Insights and recommendations for DHL on the strategic design of their OOH network, including the placement and sizing of PLs.
- An outline of the limitations of this research and directions for future research.

This concludes the research design and provides a roadmap for the remainder of this research. In the next chapter, we present a detailed context analysis of DHL's current OOH network, focusing on operational workflows, placement practices, and network performance. These insights serve as input for identifying inefficiencies and requirements for model development.

2 Context Analysis

The goal of this chapter is to gain a thorough understanding of DHL's current OOH network and its performance. In doing so, the chapter aims to generate insights that can support DHL in evaluating the effectiveness of its current placement and operational strategies. Furthermore, these insights are used to identify which aspects of the current processes should be reflected in the strategic optimisation model and its experimental design. This chapter answers the second set of research questions, which revolve around the following main question:

"What are the current operational processes, placement challenges, and performance characteristics of DHL's OOH network, and how can these insights inform a data-driven optimisation model?"

Section 2.1 examines the operational processes and delivery commitments surrounding FM and LM flows at OOH points. Section 2.2 discusses the practical, strategic, and financial challenges related to the operations and placement of PLs and SPs. Section 2.3 analyses the current state of the OOH network in terms of geographical spread and capacity. Finally, Section 2.4 evaluates the performance of the network, based on turnover, utilisation, customer pickup behaviour, parcel diversions, and the volume distribution between FM and LM within the OOH network.

2.1 Operational processes in FM and LM at OOH Points

This section outlines the operational processes at DHL's OOH points. Section 2.1.1 first describes the role of DHL's CityHubs and their involvement in OOH deliveries. Section 2.1.2 then explains the service commitments that dictate daily delivery and pickup schedules. Subsequently, the operational workflow at PLs and SPs is described at Section 2.1.3, followed by a discussion in Section 2.1.4 on the expected operational impact of the upcoming capacity control system.

2.1.1 CityHub operations and role in OOHD

DHL operates an extensive network of around 130 CityHubs within the Netherlands (see Figure 41 in Appendix A.0.1). These CityHubs serve as the distribution centres for all OOH and home deliveries. In these hubs, the routes are determined by external vehicle routing software. Some SPs and PLs are driven in dedicated routes, while others are delivered together with home deliveries. This depends on the volumes of these points and the expert opinion of the CityHub owner.

Currently, DHL is changing the PL interface to reduce handling time. In addition, the company is transitioning towards a specialised PL network, where PLs are serviced separately from other deliveries by dedicated personnel. These changes aim to enhance efficiency in the PL process.

2.1.2 FM and LM commitments at OOH points

DHL has made specific service promises to its receivers/senders regarding parcel handling at OOH points for both FM and LM processes.

- **LM Delivery Promise**: DHL guarantees that parcels delivered with day delivery to OOH points can be received after 12:00. This means that DHL must visit all OOH points before 12:00, referred to as the 'Morning Trip'.
- **FM Collection Promise**: DHL ensures that parcels dropped off at OOH points before 16:00, will be shipped the same day, meaning that DHL must visit all OOH points after 16:00 to ensure these shipments, referred to as the 'Afternoon Trip'.

2.1.3 OOH flow

Upon arriving at a PL, the driver must first collect outgoing parcels before placing new deliveries inside. This process is regulated by DHL's software system within the PLs. At SPs, this approach is also commonly followed, as business owners prefer to clear outgoing shipments as soon as possible to free up storage space. Figures 42 and 43 in Appendix A.0.2 and A.0.3 illustrate the morning and afternoon trip workflows from the perspective of the drivers.

2.1.4 Impact of future capacity control project

As previously explained, the responsibility for handling full OOH points currently lies with the driver. When a OOH point reaches full capacity, the driver must divert the package to the closest other OOH point. With the implementation of the capacity control project, as explained in Section 1.4.3, the frequency of diversions due to full OOH points is expected to decrease. The introduction of dynamic capacity control will prevent ‘expected’ future full lockers from appearing as selectable options at checkout (e.g., in webshops), redirecting demand to nearby available alternatives. This is expected to result in a more balanced distribution of parcel flows across the network. Therefore, the variability in demand for an OOH point can be reduced drastically and extreme capacity spikes at individual lockers become less frequent.

2.2 OOH points selection challenges

This section discusses key challenges in the placement of OOH points. Section 2.2.1 outlines strategic and operational constraints. Section 2.2.2 addresses financial viability, including the current cost-efficiency threshold for PLs to match the cost of SPs.

2.2.1 Strategic and operational challenges

Currently, DHL operates approximately 3,500 SPs and 1,000 PLs across the Netherlands. In recent years, DHL has shifted its strategic focus toward expanding the PL network by setting ambitious targets for future growth. By the end of 2025, DHL aims to more than double its number of PLs to 2,500, and by 2030, the company plans to operate 6,000 PLs.

PLs can be placed on DHL-owned property, at third-party business locations, or municipal grounds. The lockers at DHL-owned property are typically located at hubs, distribution centres or office buildings. However, the majority of PLs are installed at third-party locations. These are locations such as supermarkets, gas stations, or sports clubs. At these locations the placement is arranged through contractual agreements. When PLs are placed at municipal grounds, approval from the local government is required.

SPs operate within independent businesses, such as retail stores. New SPs are generally recruited by the DHL sales team but business owners can also apply to become a DHL SP. However, in recent years, retaining existing SPs and attracting new ones has become increasingly challenging. This trend contributes to DHL’s strategic shift towards expanding the PL network, in order to ensure future delivery capacity through a more controllable infrastructure.

DHL does not own the majority of OOH locations. For that reason, the expansion is dependent on external agreements with third-party businesses, municipalities, and property owners. Therefore, precise placement requirements for these OOH points are not practical. Instead, a strategic methodology is needed to estimate capacity requirements for PLs at higher aggregation levels, such as region, postcode, or city level.

2.2.2 Financial challenges

The placement of OOH points involves different costs, depending on whether the location is a SP or PL. Opening a SP does not incur direct setup costs for DHL. However, each parcel handled at a SP has a Cost per parcel (CpP) for DHL of €[redacted] .

This figure has been removed due to confidentiality.

Figure 5: (Estimated) Cost per parcel (CpP) comparison between parcel lockers and service points, *Data source: Internal DHL data*

For PLs, CpP varies based on the throughput of the locker. The total costs are calculated per year based on an average locker type and include both operating expenses and investment costs, with DHL assuming an expected economic lifetime of 12 years. Operating expenses include, for example, location rental fees, cleaning services, internet connectivity, and energy consumption. These costs need to be compensated by a sufficient parcel volume.

Figure 5 visualises the CpP for both PLs and SPs. The figure shows that PL becomes more cost-effective than SP when processing at least [redacted] packages per day, according to this internal average calculation. If the volume is lower, the CpP increases sharply, making the PL financially less viable and significantly more expensive than a SP.

2.3 Current OOH network

This section provides an overview of the current structure of DHL's OOH network. Section 2.3.1 discusses the geographical distribution of PLs and SPs across the Netherlands. Section 2.3.2 then examines the capacity characteristics of these OOH points.

2.3.1 Geographical distribution

Figure 6 presents the geographical distribution of DHL's OOH points in the Netherlands. It shows a high concentration in urban areas, particularly in the Randstad region (Amsterdam, Rotterdam, The Hague, Utrecht). More rural areas, such as Friesland, Drenthe, Groningen, and Zeeland, have fewer OOH points. The network is dominated by SPs, as they are more widespread compared to PLs, which are still in an expansion phase.

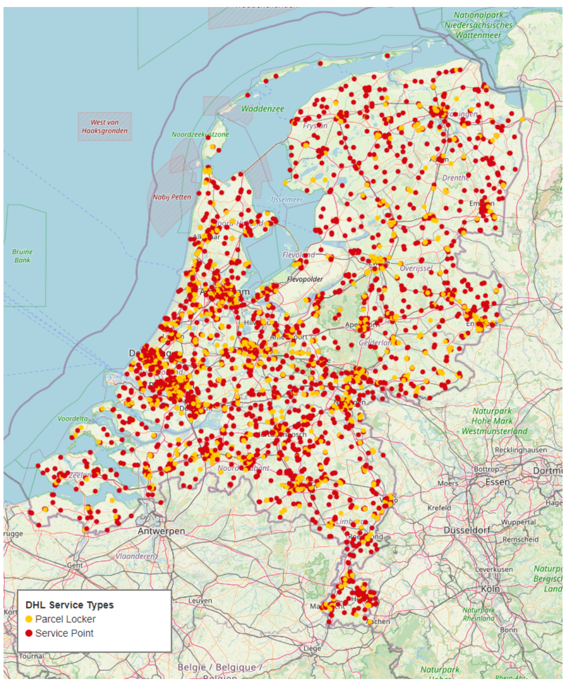

Figure 6: OOH network available for the full period from 01-11-24 to 27-03-25

This figure has been removed due to confidentiality.

Figure 7: Capacity distribution parcel locker network for the full period from 01-11-24 to 27-03-25

2.3.2 Capacity distribution

Service points
Although SPs are physical locations with limited storage space, DHL does not register or manage their exact capacity. DHL considers SPs to have no formal capacity limit, as most partnering stores are willing to handle large parcel volumes in exchange for compensation. However, during peak periods, some SPs may still refuse additional parcels due to space limitations, but this is not registered.

Parcel lockers
Figure 7 shows the distribution of capacities across all PLs. Most lockers fall within the 60 to 100 compartment

range, with a peak around 70 to 80. Only a few lockers exceed 150 compartments, and these are considered exceptions.

Figure 44 in Appendix A.0.4 shows the geographical distribution of these lockers and their corresponding capacities. The map categorises lockers into six capacity levels, ranging from 'Very Small' to 'Extra Large'. Larger lockers are mainly located in urban areas such as the Randstad and major cities in Brabant and Limburg. Smaller lockers are more commonly found in less densely populated regions, such as the northern and eastern parts of the Netherlands.

2.4 Performance

This section evaluates the performance of DHL's OOH network, based on data collected between 1 November 2024 and 27 March 2025. This period represents the longest continuous time frame during which data from all OOH points was consistently and correctly recorded. Prior to this period, not all SPs and PLs were systematically tracked, making earlier data incomplete or unreliable for performance evaluation. The analyses are based on a combination of internal data sources. The datasets were accessed, merged, filtered, and analysed using Python, with custom scripts developed specifically for this research to extract relevant insights.

Due to data access restrictions, a direct connection to DHL's Oracle-based database system was not permitted. As a result, the available data was limited in scope, and daily-level data for individual locations was not always accessible. However, in those cases, it was possible to calculate average values per or over all OOH locations over the analysis period.

Section 2.4.1 analyses the distribution of FM and LM parcel flows. Section 2.4.2 investigates customer pickup behaviour, while Section 2.4.3 explores parcel diversion patterns. Section 2.4.4 analyses turnover levels at different OOH locations, incorporating the financial feasibility threshold for PLs. Finally, Section 2.4.5 examines utilisation of PLs, both at the 'peak hours' and throughout the complete day.

2.4.1 Distribution FM and LM at OOH points

Table 1 presents the distribution of parcel flows between FM and LM at OOH points. Across the total network, LM parcel flows represent 44.5% of the activity, while FM flows account for a slightly higher 55.5%. However, this balance varies substantially between PLs and SPs.

PLs are predominantly used for LM deliveries, with 75.89% of their activity involving the delivery of parcels that can be collected by the receivers. In contrast, SPs show a more FM-focused usage profile, with 60.31% of flows being drop-offs from customers. This likely reflects the convenience of staffed SPs, where employees assist with label printing and returns. In contrast, customers using PLs typically need to print the label themselves prior to drop-off, which can form a barrier for some users. However, DHL has recently initiated a pilot for a printless return process at PLs, which may lower this threshold in the future and potentially shift the balance between the two OOH types.

Type	**LM (%)**	**FM (%)**
Parcel Locker	75.9	24.1
Service Point	39.7	60.3
Total (all OOH)	**44.5**	**55.5**

Table 1: Distribution of parcel flows in first mile and last mile per OOH point type, based on calculated average daily values on data from 01-11-24 to 27-03-25

This figure has been removed due to confidentiality.

Figure 8: ECDF of time till pickup, based on data from 01-11-24 to 27-03-25

2.4.2 Customer pickup behaviour

Figure 8 presents the Empirical Cumulative Distribution Function (ECDF) of the times till pickup for parcels collected from PLs and SPs by the receiver. The x-axis represents the number of days passed since the parcel became available for pickup, while the y-axis indicates the cumulative proportion of parcels collected by the receiver within these days. The analysis is based on aggregated pickup data across all OOH locations, as location-specific records were unavailable as explained in Section 2.4.

The steep initial rise in both distributions shows that a significant share of parcels are picked up within the first two days (around [redacted] %). Parcels from lockers tend to be retrieved slightly faster, as indicated by the higher ECDF values in the early days. This behaviour could be due to, in some cases, the increased opening times of PLs. Fast pickup is crucial, as it frees up compartment space for incoming parcels and reduces the likelihood of capacity-related rerouting or failed deliveries. DHL retrieves uncollected parcels from OOH points after a 7-day period to re-send them to the sender. The slight increase after 7 days suggests some delays in this retrieval process, which could be due to factors such as Sundays, when DHL does not collect parcels, or operational delays in recollecting unclaimed packages.

2.4.3 Parcel diversion

Table 2 summarises diversion activities within the LM of the OOH network. Out of more than [redacted] million LM deliveries during the analysis period, around [redacted] parcels were diverted. In the LM flow, a diversion represents a failed delivery to the initially selected OOH point. As such, the diversion rate can be interpreted as an estimation for the LMs service level: it reflects the proportion of parcels that could not be delivered to the receiver's preferred location. Based on this data, approximately [redacted] % of all LM deliveries were diverted, implying a service level of [redacted] % in terms of successful delivery to the intended OOH point.

Metric	Value
Total Last Mile (LM) Deliveries	[redacted]
Total Parcel Diversions	[redacted]
% of Deliveries Diverted	[redacted]

Table 2: Summary of parcel diversions within the last mile over the period 01-11-24 to 27-03-25.

From	To	Count	%
Parcel Locker	Parcel Locker	[redacted]	[redacted]
Parcel Locker	Service Point	[redacted]	[redacted]
Service Point	Parcel Locker	[redacted]	[redacted]
Service Point	Service Point	[redacted]	[redacted]

Table 3: Detailed diversion breakdown by origin and destination type, for the period 01-11-24 to 27-03-25.

Table 3 breaks down these diversions by OOH point type. It shows that [redacted] % of diversions originated from PLs, while only [redacted] % were redirected to lockers. In contrast, SPs accounted for [redacted] % of the diversions but absorbed nearly [redacted] % of all rerouted parcels. This imbalance highlights the key role of SPs in managing diversion, probably due to their higher capacity. These findings might also suggest that couriers prefer SPs as fallback locations, as they can more easily and faster accommodate all diverted parcels. Conversely, lockers, likely due to their fixed and limited capacity, are involved in more frequent rerouting. However, it is important to note that the exact reason for each diversion is not stored in the system. As a result, some diversions may not strictly reflect capacity limitations but could also stem from couriers deviating from protocol, for instance by choosing more convenient or familiar locations themselves. To further illustrate the geographic dynamics of these diversions, heat-maps of rerouting origins and destinations for both SPs and PLs are included in Appendix A.0.9 (Figures 46 to 49).

2.4.4 Turnover

OOH points

This analysis evaluates the turnover of both SPs and PLs. Turnover is defined as the average total parcel flow per day at an OOH point and is calculated by averaging the sum of the daily number of FM and LM packages.

Table 4, Figure 9 and Figure 10 show that SPs notably have a higher average turnover than PLs, with a mean of [redacted] compared to [redacted] parcels per day. What is particularly noteworthy is the difference between FM and LM turnover at PLs and SPs. The average LM turnover differs by only [redacted] parcels between lockers ([redacted]) and SPs ([redacted]). In contrast, the gap in FM turnover is significantly larger with [redacted] parcels. This indicates that PLs are substantially underutilised for FM activities. This underutilisation is responsible for around [redacted] % of the total turnover difference between PL ([redacted]) and SP ([redacted]) and therefore there is significant potential for improvement in promoting or enabling more FM usage at PLs to get the turnover rates more equal.

Table 32 in Appendix A.0.6 further shows the summary statistics of the turnover across the OOH points. PLs exhibit a more consistent performance, as indicated by their lower standard deviation ([redacted] vs. [redacted]). The interquartile range (Q3–Q1) supports this observation: turnover at SPs ranges from [redacted] to [redacted], a spread of [redacted], whereas turnover at PLs is more narrowly distributed between [redacted] and [redacted], yielding a spread of [redacted]. The maximum turnover at SPs reaches [redacted], more than five times the maximum observed at PLs ([redacted]). This indicates that while most lockers serve a 'more average' number of users, some SPs handle exceptionally high volumes. These outliers suggest that certain SPs, probably especially in high-demand areas, handle substantial volumes far beyond the average, making them particularly valuable nodes in the OOH network.

This figure has been removed due to confidentiality.

Figure 9: Boxplot of turnover by OOH point type, created using data from the period 01-11-24 to 27-03-25.

Figure 10: Supporting graph above table

This figure has been removed due to confidentiality.

Means	Parcel Locker	Service Point
FM Turnover	[redacted]	[redacted]
LM Turnover	[redacted]	[redacted]
Total Turnover	[redacted]	[redacted]

Table 4: Turnover by OOH Point Type, including FM/LM split data from the period 01-11-24 to 27-03-25.

Parcel lockers

Focusing only on PLs, a key financial benchmark for this analysis is the cost efficiency threshold of [redacted] packages per day, as explained in Section 2.2.2. A PL becomes only more or equally cost-efficient as a SP when this daily volume is met. However, the analysis reveals that the average daily turnover of a PL is just [redacted] parcels (see Table 4), a figure that falls significantly below the break-even point, and already signals structural inefficiency across the locker network mainly due to the lack of FM packages. Figure 11 further illustrates this: the distribution is heavily right-skewed, with most lockers clustered in the [redacted] to [redacted] parcel range and only a limited number exceeding the [redacted]-parcel benchmark. This means that most lockers are structurally underperforming and less financially viable than SPs. Table 5 quantifies this pattern: [redacted]% of the lockers do not even reach a daily turnover of [redacted] parcels, and only [redacted]% exceed the break-even threshold of [redacted]. That means that [redacted]% of the lockers are not financially viable according to the threshold. Alarmingly, [redacted]% of the lockers handle fewer than [redacted] parcels per day, and [redacted]% operate at less than [redacted] parcels per day, potentially making them up to ten times more expensive per parcel than SPs.

Figure 12 shows that even the best performing provinces, such as Utrecht and Noord-Brabant, do not reach the cost-efficiency threshold on average. In less urbanised provinces such as Drenthe and Fryslân, average turnover is as low as [redacted]–[redacted] parcels per day, being on average already twice as costly per parcel than SP. This geographical distribution of these turnover levels is further visualised in Figure 13, where each PL is colour-coded based on its average daily turnover. The map reveals clear regional differences, with clusters of lockers around or better than the threshold (green dots) concentrated in densely populated areas such as the Randstad, especially within the big cities. Many lockers in rural regions, particularly in the north and east, show almost only lower turnover levels (red and orange dots). These red dots, and especially dark red, indicate a significant financial burden. Their CpP can range from approximately twice as expensive as an SP to nearly ten times more expensive in cases where daily turnover drops to only [redacted] parcels. From an operational cost-efficiency standpoint only, these locations could be reevaluated for potential closure or integration into nearby higher performing points.

This figure has been removed due to confidentiality.

Figure 11: Distribution of parcel locker turnover across the network, based on data from 01-11-24 to 27-03-25.

This table has been removed due to confidentiality.

Table 5: Cumulative distribution of parcel lockers by daily turnover threshold, based on data from 01-11-24 to 27-03-25.

This figure has been removed due to confidentiality.

Figure 12: Average parcel locker turnover by province over the period 01-11-24 till 27-03-25.

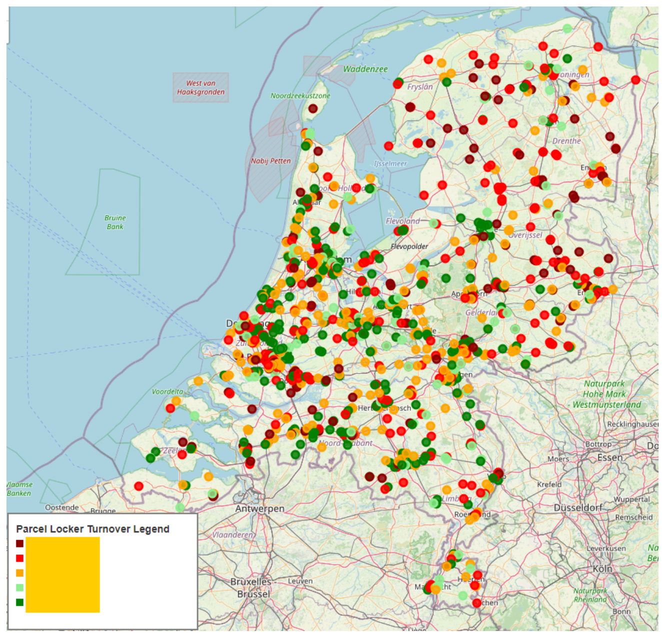

Figure 13: Geographical distribution of parcel locker turnover, based on data from 01-11-24 to 27-03-25.

To gain better insight into potentially well- or poorly-performing locker locations, Figure 14 shows the average daily turnover for PL by location type. Due to the difference in group sizes, it is particularly notable to look at groups with a high number of records. Supermarkets, DIY stores (e.g., Gamma, Praxis) and gas stations show above-average turnover, each exceeding the overall average of [redacted] parcels per day. Especially supermarkets and tank stations, which make up about [redacted] % of all lockers, perform significantly better than others, both averaging around [redacted] parcels per day. In contrast, gyms, sports clubs, and stadiums perform below average.

This figure has been removed due to confidentiality.

Figure 14: Average locker turnover per 'Group' over period 01-11-24 till 27-03-25

2.4.5 Utilisation

Utilisation per locker

Unlike PLs, SPs do not have a predefined or registered capacity within the DHL systems. As a result, it is not possible to calculate the utilisation rate of SPs. Therefore, this analysis focuses solely on PLs, for which the number of compartments is known. The utilisation rate per PL is calculated as the average number of parcels handled relative to the total capacity of the locker, providing information on how intensively the lockers are used. This is based on available data, where maximum inventory levels are recorded for each hour of the day.

Figure 15 shows the distribution of average utilisation rates for PLs, with the mean utilisation across the PL network being [redacted] %. Unlike turnover, there is no universally accepted threshold for what is a 'good' or 'bad' utilisation level. Interpreting utilisation is inherently more complex, as PLs are dynamic by nature: their compartments are filled and emptied throughout the day. For example, a locker might be completely full right after the morning delivery but nearly empty again in the evening and during the night as customers pick up their packages. Therefore, even lockers with seemingly low average utilisation can experience moments of full capacity. This limits conclusions from averages, though the distribution still gives useful insight into locker usage.

This figure has been removed due to confidentiality.

Figure 15: Distribution of average utilisation for parcel lockers, based on data from 01-11-24 to 27-03-25.

This figure has been removed due to confidentiality.

Figure 16: Parcel locker capacity vs average utilisation, based on data from 01-11-24 to 27-03-25.

When examining the actual full occupancy moments of PLs, a relevant indicator for FM sending success, PLs were found to be fully occupied only [redacted] % of the time. In total, only five lockers ever recorded any full occupancy with an average of [redacted] % of the time, with the highest being [redacted] % at a location in central Amsterdam. These results suggest that senders in the FM flow rarely encounter capacity limitations.

The geographical trend of low percentage utilisation follows a similar pattern to what was observed with low-turnover lockers, suggesting that lockers with higher turnover also tend to achieve higher percentage utilisation.

Figure 45 in Appendix A.0.9 visualises this trend in a scatter plot. This pattern is intuitive, as more frequent usage naturally results in better capacity usage over time.

Figure 16 explores another relationship by visualising average utilisation against locker capacity. A weak positive correlation can be observed, indicating that larger lockers tend to have higher utilisation. However, this analysis also reveals that a lot of high capacity PLs, ranging from 150 to even 300 compartments, are significantly underutilised, operating at less than [redacted] % capacity. Also, lockers with lower capacities, such as those with around 60-80 compartments, have enormous variation in usage. Some PLs operate under [redacted] % while others exceed [redacted] % utilisation, highlighting a mismatch between locker size and local demand.

Utilisation per hour

In addition to the locker-specific utilisation calculation, a second method is used to analyse utilisation over time. This approach examines the average utilisation across all PLs at each specific hour of the day. The average of the maximum hourly inventory levels (relative to capacity) are computed across all lockers. This results in a time-based utilisation profile, providing insights into how intensively lockers are used throughout the day.

This approach offers valuable insights for DHL. Around the expected peak, just before the LM delivery deadline at 12:00, the average locker utilisation reaches only about [redacted] % (Figure 17), [redacted] . To better assess peak-period capacity alignment, this section focuses on the hours 11:00–13:00, when deliveries have been made but most parcels have yet to be collected, as explained in Section 2.1. This window provides a meaningful snapshot of how well locker capacity aligns with actual demand at its anticipated peak. We only observe a peak driven by LM dynamics, as the FM still represents a relatively small share of total locker flows as shown in Section 2.4.4

This figure has been removed due to confidentiality.

Figure 17: % Utilization per hour for all parcel lockers, based on data from 01-11-24 to 27-03-25.

This table has been removed due to confidentiality.

Table 6: Cumulative share of parcel lockers by utilisation level (hours 11-13), based on data from 01-11-24 to 27-03-25.

Table 6 presents the cumulative distribution of PLs during the 'peak period'. It shows that half of all lockers operate below [redacted] % utilisation at this critical time of day, and only 2.25% exceed [redacted] % utilisation. This distribution highlights that high utilisation is the exception rather than the norm, even at the moment when occupancy is expected to be at its highest. Notably, one in three lockers does not even reach [redacted] % utilisation, indicating that these lockers are also significantly underused during the delivery peak.

Together, these findings reinforce the earlier suggestion that capacity constraints for FM usage are minimal in DHL's current OOH network, even during the morning peak window capacity. Moreover, due to the natural flow of operations, parcels are picked up throughout the day after deliveries, creating continuous availability for FM drop-offs.

Figure 50 in Appendix A.0.8 provides a geographical view of PL utilisation during the peak period. Each locker location is colour-coded based on its average percentage utilisation between 11:00 and 13:00. The geographical distribution reinforces earlier findings: a significant share of lockers across the country, particularly in the northern and eastern provinces, show low utilisation, indicated by the frequency of orange and red markers. These colours represent lockers operating well below [redacted] % capacity during the peak period. Only a limited number of locations, mostly in urban centres, display green dots, signifying high utilisation.

2.5 Summary & conclusion

This chapter investigated DHL's current OOH network and processes to answer the research question: *"What are the current operational processes, placement challenges, and performance characteristics of DHL's OOH network, and how can these insights inform a data-driven optimisation model?"* To answer this question, a structured analysis of operational flows, placement challenges, and network performance was conducted. The findings consist of three types of insights: (1) insights into how DHL's FM and LM processes are currently structured and executed, (2) insights for DHL regarding the actual performance of its OOH infrastructure, and (3) insights that can support the development of a strategic optimisation model.

Insights into DHL's OOH process
The DHL OOH network is coordinated by 130 CityHubs, which manage all parcel flows via routing software, although flexibility remains in how PLs and SPs are served in practice. The operational model revolves around two fixed trips per day: a morning delivery round (LM) before 12:00, and an afternoon pickup round (FM) after 16:00. PLs and SPs are handled accordingly, with outgoing FM parcels always cleared before new LM deliveries are placed. The network is skewed towards urban areas, especially the Randstad, while rural regions remain less covered. PLs vary in size and have known capacities, unlike SPs, which are more flexible but untracked in terms of capacity. Most PLs are located at third-party venues such as supermarkets, gas stations, or gyms. The current locker network is still in expansion, with plans to grow from 1,000 to 2,500 lockers by 2025 and to 6,000 by 2030.

Insights for DHL
Firstly, a clear distinction is observed between location types: PLs placed at supermarkets, gas stations, or DIY stores consistently show above average turnover, while lockers at sports clubs, gyms, and stadiums often underperform. On a geographic level, urban regions, particularly the Randstad and large cities show stronger performance, whereas rural provinces consistently fall short.

The turnover analysis shows that the average PL handles [redacted] parcels per day. The analysis further highlights that over [redacted] % of lockers do not exceed [redacted] parcels per day, and nearly [redacted] % fail to meet the cost-efficiency threshold of [redacted] parcels. This threshold indicates the volume at which a locker becomes more cost-effective than a SP, based on yearly costs spread over the expected lifetime of an 'average' PL. Although LM turnover is relatively consistent between PLs and SPs, FM activity at lockers remains structurally low, accounting for [redacted] % of the average turnover gap of [redacted] parcels between SPs and PLs.

The utilisation analysis reveals that PLs are rarely full. Even during peak hours (11:00–13:00), only 2.25% exceed [redacted] % utilisation, and just [redacted] % reach full occupancy in this window. Across all hours, this figure drops to only [redacted] %, suggesting a notably high service level for FM packages. Note that utilisation is measured on total capacity due to the data limitations. In practice, lockers may be 'full' for specific parcel sizes while still having overall space. This can result in censored demand for FM parcels, as users may abort drop-offs when no suitable compartment is available. In addition, locker space is frequently freed throughout the day due to continuous FM pickups, which ensures that LM drop-offs remain possible. This underlines a clear opportunity for DHL: since capacity pressure is mainly driven by LM deliveries, substantial locker space remains unused throughout the day, offering potential to grow FM flows, especially because this is where the biggest turnover gap with SPs lies.

Finally, the diversion analysis, which examines cases where parcels are rerouted from their originally selected OOH location, shows that SPs act as the primary buffer in the network. Although nearly [redacted] % of the diversions originate from lockers, most are redirected to SPs, indicating a heavy dependence on their flexibility. Despite a relatively high service level for LM deliveries ([redacted] %), defined as the parcels that are successfully delivered to the intended OOH point, the reasons behind diversions, whether due to capacity constraints, courier preferences, or other factors, remain unclear. Gaining better insight into these causes is essential for DHL to improve the network and minimise unnecessary rerouting.

Implications for the strategic optimisation model
First, the model should reflect DHL's strategic focus on expanding the PL network. The primary objective is to identify where new PLs should be placed to support long-term network growth, as DHL aims to scale up its locker infrastructure nearly sixfold by 2030. Although the existing SP network is not the primary focus, the model should still consider its role, alongside home delivery, as part of the overall OOH strategy. It can be used to explore how these alternatives interact with the future locker network, particularly to identify areas where lockers may not be viable and SPs or home delivery remain strategically relevant.

Second, the results show a clear mismatch between the placement of lockers and the actual demand. Some lockers handle fewer than 5 parcels per day, while others are heavily used. This is further supported by utilisation figures: many large lockers are underused compared to smaller ones, revealing a mismatch between locker sizing and local demand. These findings underline the need for a strategic model that aligns locker placement and sizing with actual demand. Rather than applying uniform growth strategies, the model should estimate regional demand and allocate capacity accordingly to prevent under- and overutilisation as well as inefficiencies.

Third, the model should support both the identification of new locker locations and the strategic evaluation of the current network. By comparing the existing network to a greenfield scenario ignoring the current network, where no lockers are yet placed and optimal locations are chosen purely based on demand and cost, it becomes possible to assess which current placements align with optimal design, and which do not. This allows DHL to justify well-placed lockers and identify underperforming locations that may be candidates for relocation or removal.

Fourth, while current cost comparisons between PLs and SPs are based on average figures, this approach is too simplistic for designing a OOH network. The model must instead incorporate the full cost structure across all delivery modes, including home delivery, different PL types, and their associated fixed and variable costs. Locker sizes and installation expenses vary significantly, yet such variation is not reflected and known in DHL's current decision-making or analysis. A more granular financial perspective would enable the strategic model to evaluate not just where capacity is needed, but also which type or size of OOH infrastructure, if any, is financially sustainable. In low-demand areas, the model may even conclude that home delivery is the only cost-effective solution, a critical insight given the increasing operational costs of DHL. This thesis addresses this gap by constructing a more detailed and realistic cost framework across all delivery modes and OOH facility types and sizes, making strategic decision-making possible within an optimization model.

Fifth, since this model takes a strategic perspective, it does not aim to address operational issues such as day-to-day locker overflows. However, with the introduction of the capacity control project, which prevents customers from selecting lockers 'predicted' to be full, such issues are expected to diminish, making them increasingly irrelevant for strategic decision-making. This system will steer demand toward available locker capacity using short-term forecasts, but this only works if sufficient structural capacity is in place within an acceptable pickup range from the demand points. As a result, the main function of the strategic model is to ensure that sufficient OOH capacity is available within each demand region, making sure that expected demand can be absorbed efficiently across DHL's network.

Last, the model should account for the distinct flow characteristics of FM and LM parcels, as these directly influence locker capacity and are therefore essential for strategic capacity planning. FM parcels are always collected before new LM deliveries are inserted, while LM parcels remain in the lockers until receivers collect them. As a result, FM parcels are only in the locker for a short time, typically less than a day, while LM parcels can remain much longer depending on when they are picked up. ECDF analysis of pickup behaviour shows that around [redacted] % of LM parcels are collected within one day, and all must be within seven days, due to DHL's policy. Incorporating this behaviour prevents underestimating capacity driven by parcel stay durations.

Together, these conclusions provide a clear foundation for the design of the strategic optimisation model, the associated data preparation and usage, and the supporting literature review in the next chapter, that examines academic approaches to optimising OOH networks.

3 Literature Review

To inform the development of a suitable strategic optimisation model, this chapter reviews academic literature on facility location methods applicable to OOH networks. This chapter addresses the following research question:

"Which methods for the strategic placement of OOH points are discussed in the literature?"

To answer this question, Section 3.1 outlines the scope of this literature review, summarising the growth of academic interest in OOHD and explaining the focus on Facility Location Problem (FLP) for PLs and SPs. Section 3.2 introduces the FLP. Section 3.3 reviews existing studies on FLP related to OOHD and systematically compares a selection of studies based on key modelling features and research characteristics. Finally, Section 3.4 identifies literature gaps and highlights how this thesis contributes to the literature.

3.1 Literature scope and selection

Research on OOH points has grown significantly with the rise of e-commerce and the increasing integration of OOH points. In the literature, three main types of models are used to support decision-making in OOHD: location models, routing models, and integrated location-routing models. Figure 18 illustrates the increasing number of publications of these model types within OOHD. In this figure, FLP represents Facility Location Problems, VRP covers Vehicle Routing Problems, and LRP & Mobile refers to Location-Routing Problems, which integrate both strategic location decisions and operational routing decisions, including the novel concept of mobile PLs. Despite this growth, the absolute number of studies remains relatively low, highlighting the novelty and scarcity of research in this specific domain. This research focusses on the literature concerning OOH FLPs.

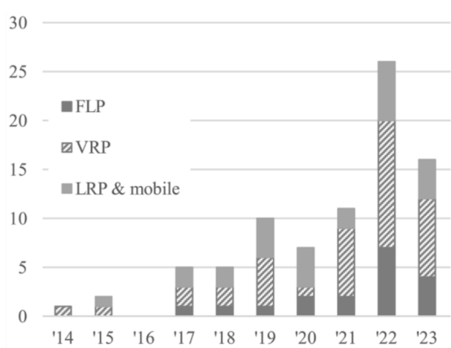

Figure 18: Number of publications per year (Janinhoff et al., 2024)

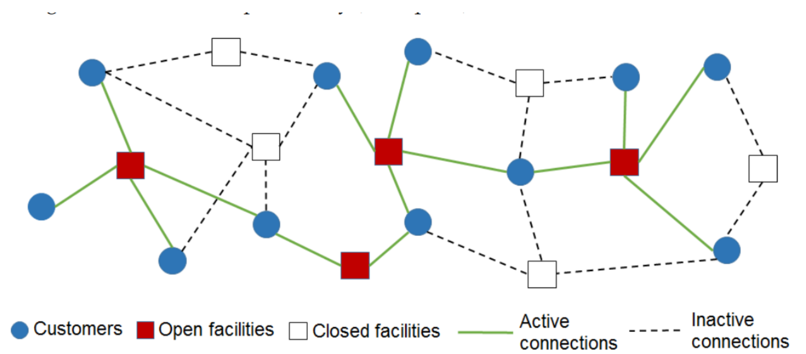

Figure 19: Illustrative example of a classical FLP from Rabe et al. (2021), based on de Armas et al. (2017)

3.2 Facility location problem

The FLP was introduced in the field of operations research in the 1960s and was initially referred to as the Plant Location Problem (Balinski, 1965). The problem involves determining the optimal location for one or more facilities to efficiently serve a set of demand points. It is widely applied in various industries, such as the placement of train stations, gas stations, stores, and airports.

In its basic formulation, the FLP consists of a set of potential facility locations and a set of demand points or customers that these locations must serve, where each facility has a fixed cost for opening and a variable cost associated with operation. The objective of the model is to determine which subset of facilities should be opened in order to serve all customers while minimising total costs. These costs include fixed facility costs, operational costs, and transportation costs, which are usually modelled based on distance (Northwestern University, 2022). Figure 19 provides a visualisation of the basic FLP concept. Various approaches have been used to solve FLPs, including genetic algorithms (Wadhwa & Garg, 2011), exact methods, and heuristic algorithms (Ulukan & Demircioğlu, 2015). Recent research explores the optimisation of PL networks as a solution to LM logistics challenges.

3.3 FLP regarding OOH points in LM and FM

All research found on FLP for OOH points focusses on LM logistics, or does not explicitly distinguish between LM and FM operations. This aligns with DHL's case, where capacity problems in the FM are not a significant issue compared to the LM.

To understand the differences between DHL's case study and those found in academic literature, a selection of 17 relevant papers covering FLP for OOH points were reviewed. These papers represent the subset of all academic work that explicitly combines parcel delivery, facility location modelling, and OOH infrastructure. Due to the lack of a widely accepted definition for OOHD, academic literature often does not clearly distinguish between SPs and PLs (Janinhoff et al., 2024). In this review, we classify the problems according to the definitions used in this research.

Section 3.3.1 categorises the reviewed studies based on whether they optimise only location, or both location and capacity decisions, and gives a brief summary of each paper. From there, Section 3.3.2 evaluates a subset of ten key studies in more detail by systematically comparing modelling and research characteristics.

3.3.1 Research selection and categorisation

The first step in this review is to analyse the objectives of existing research and what these studies aim to determine in order to make an initial selection of relevant literature. A more detailed modelling comparison of the most relevant studies follows in the next section. The key objective, from DHL's perspective, is to identify high-potential locations and, importantly, determine the required capacity at each location. Therefore, we assessed whether each study: only optimises location placement, or determines both location and capacity requirements. This distinction is crucial for DHL, as the company cannot freely select exact OOH locations but instead seeks a tool to evaluate location suitability and determine necessary capacity. Based on this, a further sub-selection of research is made. Table 7 provides an overview of all the studies reviewed and their key objectives.

Much of the reviewed research focusses on maximising total coverage and ensuring acceptable customer travel distances under various conditions. For example, Lin et al. (2020) and Tadic et al. (2023) optimise facility locations to maximise demand coverage. Similarly, Lin et al. (2020) employ discrete choice modelling to analyse how OOH network design influences customer behaviour. Additionally, Luo et al. (2022) highlight the impact of locker location type, showing that lockers closely tied to shopping malls or metro stations significantly influence customer usage. Next to that, for example Faugère and Montreuil (2018), do not address location selection but instead focusses on facility layout optimisation based on pre-determined locations. Notably, the study by Xu et al. (2021) is the only study found that considers both SPs and PLs. However, it develops two separate models for each type and does not combine them in one network.

Table 7: Summary of research on OOH location problems

Work	Determines	Summary
Lin et al. (2020)	Location	Presents a model to optimally place PLs in a shared last-mile delivery network, maximizing service levels while incorporating customer choice through a multinomial logit framework.
Luo et al. (2022)	Location, Capacity	Focusses on a multi-objective PL network design problem that aims to optimise the total cost of the network and the accessibility of customers to PL stations.
Lyu and Teo (2022)	Location	Presents a model to optimally place parcel lockers (PLs) in a shared last-mile delivery network in Singapore, maximizing utilization and reducing congestion in the central business district.
Mancini et al. (2023)	Location	Addresses the problem of locating locker boxes in the last-mile delivery context under uncertainty in demand and capacity, modelling it as an extension of the capacitated facility location problem, and proposing a stochastic mathematical model as well as three matheuristics to solve it. However the capacity of the lockers in the paper is fixed and homogeneous across all locker stations.
Deutsch and Golany (2018)	Number, Location	Presents an approach to determine the optimal number and locations of PL facilities as a solution to the Logistics Last Mile Problem, with the objective of maximizing the total profit for the consignment company.

Continued on the next page...

Table 7: (continued)

Work	Determines	Summary
Faugère and Montreuil (2018)	Layout, Design	Presents optimization-based methods for designing smart locker banks in the context of omni-channel business-to-consumer logistics and supply chains. It aims to maximise expected profit considering ergonomic, acquisition, and implementation costs.
Kahr (2022)	Layout, Location, Capacity	Proposes and studies the stochastic multi-compartment locker location problem (SMCLLP) to determine optimal locations and layouts for PLs to support supply chain viability at the last mile with the goal of maximise expected profit given budget constraints using Benders decomposition.
Lee et al. (2019)	Location	Presents a decision-making system for selecting the optimal locations to install unmanned PLs in residential areas, which integrates the processes of finding potential locations, determining the number of locations, and selecting the optimal locations using location set-covering and p-median models to maximise demand coverage of PLs.
Ottaviani et al. (2023)	Number, Location, Module Capacity	Proposes a solution to the problem of optimally locating Automated PLs (APLs) to serve customer demand, combining mixed-integer linear programming and greedy heuristic algorithms, and testing the approach on real customer demand data from Turin with the goal of minimising cost while covering 90% of estimated demand.
Rabe et al. (2021)	Number, Location	Proposes an integrated simulation-optimization approach that combines system dynamics, facility location optimization, and Monte Carlo simulation to determine the optimal number and location of automated PL systems that minimises cost for last-mile distribution in the city of Dortmund, Germany.
Raviv (2023)	Number, Location, Capacity	Presents a mixed-integer linear programming (MILP) model to simultaneously optimise the location and capacity of a network of service points equipped with automatic PLs to facilitate last-mile parcel delivery, with the objective of minimizing the total setup cost of the SPs and the expected number of rejected parcels.
Sawik et al. (2022)	Number, Location	Presents a multi-criteria simulation-optimization analysis to determine the optimal location and number of automated PLs in the city of Poznan, Poland, in order to maximise profits and minimise costs for last-mile delivery.
Sweidan et al. (2022)	Location, Capacity	Presents a mixed-integer linear programming model to optimise the placement and usage of PL stations as a solution to the last-mile delivery problem in e-commerce, with a case study in Doha, Qatar, with the goal to maximise profitability while maintaining high service levels.
Tadic et al. (2023)	Location	Introduces a novel hybrid model for locating collection and delivery points (CDPs) in users' households using a combination of heuristic and meta heuristic algorithms, with the goal of minimizing the sum of distances between users, the nearest CDPs and maximizing total coverage.
Wang et al. (2020)	Location, Capacity	Proposes a robust optimization approach to determine the number and locations of movable PL units to minimise the operating cost under stochastic demands in order to minimizing operating costs.
Wang et al. (2022)	Location, Number	Presents a robust optimization method for determining the optimal locations and number of PLs under uncertain delivery demands, considering both large and small parcels to be received and sent by customers.
Xu et al. (2021)	Location	Proposes a data-driven method to optimise the locations of collection and delivery points (CDPs) for online retailers by estimating customer purchase probabilities and then optimizing the locations of attended and unattended CDPs, but both in separate models.

In the context of DHL's operations and objective, literature that considers both location and capacity (or number) decisions is relevant; we denote this as the Last Mile Parcel Capacitated OOH Facility Location Problem. However, there are no studies related to the combination of both PL and SPs in a capacitated facility location problem. Therefore, literature is reviewed based on what we denote as the Last Mile Capacitated Parcel Locker Location Problem (LMCPLLP) which results in ten relevant studies that are reviewed in detail in the next section.

3.3.2 Last Mile Capacitated Parcel Locker Location Problem

This section classifies the selected LMCPLLP literature based on the distinctive characteristics of the problems and solution methods. To systematically compare the various approaches in the literature, Table 8 presents an overview of ten relevant studies. Each column in the table represents a distinct modelling characteristic, allowing for cross-paper insights:

- **Locker Placement Set**: Indicates whether locker locations are freely chosen, fixed to demand nodes, or restricted to a predefined candidate set.

- **Demand Location**: Describes how detailed the demand data is and at what level it is grouped (e.g. per customer, per area, or per district).
- **Demand Modelling**: Captures whether demand is modelled deterministically, stochastically, or via robust formulations, as well as how it is generated (e.g., historical-data, scenario-based, or simulation-driven).
- **Demand Allocation**: Describes how demand is assigned to lockers, such as through proximity-based rules, layered willingness-to-travel structures, or accessibility functions, and indicate whether a demand location must be fully served by a single locker or whether fractional assignment across multiple lockers is possible.
- **Demand Fulfilment**: Reflects whether all demand must be met, if partial fulfilment is allowed, or if fulfilment depends on customer willingness or system constraints.
- **Parcel Overflow Handling / Divert Location**: Specifies whether overflow at locker level is allowed, and, if so, how it is managed, through rejection, redirection, postponement, or rerouting.
- **Capacity**: States whether lockers are capacitated or not.
- **Solution Method**: Identifies the primary computational technique used, ranging from exact (e.g., Mixed Integer Linear Program (MILP)) to hybrid or heuristic approaches.
- **Objective**: Summarises the optimisation goal, such as minimising cost, maximising profit, or balancing multiple criteria.
- **Extra**: Provides any noteworthy modelling extensions, such as multi-period structures, mobile lockers, pickup behaviour modelling, replenishment logic, or if the study includes details on compartment types.

This comparison facilitates the identification of common patterns, novel contributions, and gaps in current modelling practices.

Table 8: Comparison of locker optimization approaches

Study	Locker Placement Set	Demand Location	Demand Modelling	Demand Allocation	Demand Fulfilment	Parcel Overflow Handling/ Divert Location	Capacity	Solution Method	Obj.	Extra
Deutsch and Golany (2018)	Same as demand nodes	Aggregated nodes (city areas or neighbourhoods)	*Deterministically* aggregated per node	Decreases with distance (& layered based)	Variable	None, due to unlimited capacity	Unlimited	MILP	Max Profit	
Ottaviani et al. (2023)	Predefined nodes (centroids postcodes)	Aggregated nodes (synthetic per 100 customers)	*Stochastic*, generated from exponential and beta distributions	To closest, within radius	Constraint, predefined minimum share of total demand	None, lost demand	Limited	MILP & Greedy Heuristic	Min Costs	
Rabe et al. (2021)	Same as demand nodes (aggregated per district)	Aggregated nodes (one per district)	*Stochastic*, generated from SDSM forecasts and sampled from uniform, triangular, and log-normal distributions in simulation.	Distance-Based Allocation (higher costs for further)	Full Fulfilment (in model), Partial (in testing)	None, lost	Limited	Hybrid: MILP + Simulation (SDSM + Monte Carlo)	Min Costs	Multi-Period
Wang et al. (2020)	Predefined nodes (Restricted Locations)	Aggregated nodes ((sub)-community level)	*Stochastic* Demand (Robust Optimisation)	To closest, within radius	Variable (Unmet not penalised)	Overflow Redirects (depot return or fixed locker)	Limited	Exact Optimisation (ILP with Robust Modelling)	Min Costs	Movable Mobile PLs
Luo et al. (2022)	Same as demand nodes	Individual customer Locations	*Deterministic* Demand	Distance-based (& Accessibility-based)	Partial Fulfilment (System-Wide, some customers no locker)	None (demand assigned must be fulfilled)	Limited	Meta heuristic (ALPEA Evolutionary Algorithm)	Multi-Obj (Costs, Accessibility)	Fresh/ Normal goods lockers
Kahr (2022)	Predefined nodes	Aggregated nodes (per grid cell)	*Stochastic* Demand (Scenario-Based)	Cover radius only	Variable (only with coverage)	None	Limited	Exact Optimisation (ILP + Benders Decomposition)	Max profit (under budget constr.)	Multi-Compartment, Deterministic Replenishment rate (pickup behaviour)

Continued on next page

Table continued from previous page

Study	Locker Placement Set	Demand Location	Demand Modelling	Demand Allocation	Demand Fulfilment	Parcel Overflow Handling/ Divert Location	Capacity	Solution Method	Obj.	Extra
Raviv (2023)	Predefined Nodes	Aggregated nodes (certain area)	Demand is *stochastic*, and based on the area characteristics: parcel arrivals follow a Poisson distribution, and pickup durations are geometric	To closest, within radius	Full Fulfilment (Strict Assignment Required of demand points)	Overflow Handling (Rejection/-Postponement) (penalised)	Limited	DTMC & MILP	Min costs	Stochasticity in pickup, geometric distribution
Sawik et al. (2022)	Same as demand nodes	Aggregated Nodes (per district)	*Stochastic* Demand (Simulation, scenario-driven)	Within district	Full coverage	None (demand assigned must be fulfilled)	Limited	Hybrid: MILP + Agent-based Simulation	Min Cost (Multi-Criteria)	Multi-Period (updating)
Sweidan et al. (2022)	Same as demand nodes	Aggregated nodes (per district)	Demand is *deterministic*, based on population data and fixed ordering frequencies for a single day	Distance Layer-based Assignment (with perc.)	Variable (if layer available)	None	Limited	Exact Optimisation (MILP)	Max Profit	
Wang et al. (2022)	Same as demand nodes	Aggregated nodes (per area)	Demand is *stochastic* but bounded. Uses Bertsimas, Sim robust optimisation to protect against worst-case deviations.	To closest, within radius	Variable	None, lost	Limited	Exact Optimisation (MILP with Robust Formulation)	Min costs	Large/Small Compartments & Deterministic dwell time (pickup behaviour)

While Table 8 presents a structured overview of modelling characteristics across the selected studies, this section discusses broader trends, key modelling innovations, and recurring limitations. It highlights how various approaches differ in terms of assumptions, objectives, and techniques, and connects these findings to identify opportunities for this research.

Foundational model

A fundamental contribution to the locker location problem is made by Deutsch and Golany (2018), who formulates an uncapacitated facility location model. Their objective is to maximise total profit, defined as revenue from customer use minus fixed and operational costs of lockers. Lockers are located at predefined aggregated nodes, with demand deterministically modelled per node based on population and order frequency. Customers are assumed to choose the nearest locker, following a distance-minimisation logic. To reflect diminishing willingness to travel, the model includes distance 'layers', each with decreasing percentages of customers willing to collect parcels from lockers further away. Partial demand fulfilment is allowed; customers unwilling to travel are considered lost demand. The model is solved as an Integer Linear Program (ILP) and shown to be equivalent to the classical Uncapacitated Facility Location Problem (UFLP). It serves as a conceptual uncapacitated basis for many capacitated extensions in later work.

Capacitated extensions under deterministic or stochastic demand

Sweidan et al. (2022) extend the uncapacitated framework by incorporating capacity constraints per locker station while maintaining the layered distance assignment from Deutsch and Golany (2018). Their model maximises profit by accounting for lost demand caused by customers' unwillingness to travel, along with operational costs and incentives linked to travel distance. Locker overflow is not explicitly handled and capacity constraints must be satisfied. If too much demand is allocated to a locker, the solution is infeasible. All customer demand must be assigned, and there is no allowance for unmet or rejected demand, unless due to behavioural unwillingness to travel (which is built into the model via participation ratios in the layers). Demand allocation follows this layered structure, prioritising nearby lockers (starting at Layer 0) and shifting demand outward if no closer lockers are available. Demand is deterministically aggregated at the zone level, based on survey and statistical data for a single-day planning, and the model is implemented as a single-period MILP.

Ottaviani et al. (2023) further advance the capacitated modelling by considering variation (stochasticity) in both demand volume and individual willingness to travel based on survey results. Lockers can be modularly expanded, and the model includes a fulfilment of minimum share of total demand through a service level constraint. The lock sites are fixed to 33 postcode centroids in Turin, with stochastically generated demand at 1,020 synthetic user clusters. Parcel volumes and travel tolerances are drawn from exponential and beta distributions, respectively. Customers are assigned to their nearest locker within their individual travel range; excess demand is not fulfilled. The authors propose a hybrid solution approach: first, they solve a MILP; thereafter, two greedy heuristics (Algorithm 1 and 2) are developed to efficiently generate 'good' quality solutions for larger instances with lower computation time. The methods are applied separately, there is no integration between the MILP and heuristics. Both are tested on real-world demand data from Turin, validating practical applicability.

Dynamic and simulation-based models

Where previous studies proposed single-period approaches, Rabe et al. (2021) take a multi-period perspective by proposing an integrated simulation–optimisation framework for strategic planning of PL networks under demand uncertainty. Locker placement is restricted to 62 predefined districts, with demand generated using forecasted trends and simulation-based stochastic scenarios derived from demographic and e-commerce data in a real-world case study of Dortmund, Germany. Demand allocation follows a cost-minimisation principle based on Euclidean distance between district centroids, assuming customers are assigned to the nearest locker in terms of cost efficiency. Once installed, lockers remain operational throughout the 36-month planning horizon and must meet minimum utilisation thresholds to avoid underuse. Although the model allows lockers to be opened at different points within the planning horizon, these openings are strategically planned based on forecast demand and cost considerations. This multiperiod structure does not reduce the model's strategic nature; instead, it enhances realism and flexibility by reflecting the gradual rollout of locker networks, as also observed in practice. The models aim to minimise total cost, including installation and service costs, while ensuring capacity constraints are respected. Overflow is not permitted; unmet demand contributes to a reliability measure evaluated via Monte Carlo simulation. The model is solved as a MILP, and results demonstrate trade-offs between cost and service reliability in locker network design.

Sawik et al. (2022) also use simulation by developing a dynamic simulation–optimisation framework for large-scale PL network planning. Their approach integrates a MILP model with agent-based simulation to evaluate evolving parcel demand over a three-year horizon and to dynamically respond to changes in population, e-commerce adoption, and user behaviour. The framework is tested using real-world data from Poznań, Poland. Lockers are capacitated and placed at district level, with full demand coverage required in each simulation period. In each period, parameters of the mathematical model are treated as deterministic, although the dynamic nature of the problem is addressed through simulation. Demand allocation and modelling are based on district-level assignment rather than precise distance minimisation; while location costs are distance related, no Euclidean distance is used. Locker overflow is not explicitly modelled, but capacity constraints are hard, meaning any assignment exceeding capacity makes the solution infeasible. Each PL has a defined number of compartments, and the total parcel demand per district must not exceed this capacity. The objective is to minimise total system costs, including setup (opening) costs, decommissioning (closing) costs, maintenance costs, and customer assignment costs. The model follows a multi-objective, cost-oriented optimisation. Lockers can be installed or removed in city districts (not individual addresses). As in the study of Rabe et al. (2021), the model is multi-period, covering a three-year planning horizon. Locker deployment decisions are reviewed monthly, and network reconfiguration is driven by dynamic changes in population, parcel demand, and customer behaviour. This enables adaptive and time-responsive optimisation through iterative coupling of simulation and mathematical programming.

Pickup modelling and advanced locker management

Raviv (2023) presents a stochastic optimisation model that uniquely integrates customer pickup behaviour using a discrete-time Markov chain (DTMC). Parcel arrivals (demand) follow Poisson distributions, and pickup durations are geometrically distributed. Two overflow handling strategies are considered: (1) rejection, where excess parcels are redirected (e.g., home delivery), and (2) postponement, where undelivered parcels are stored at the depot and delivered in the next cycle. These strategies are incorporated into the objective function through penalty and holding cost terms. The model designs a network of 'SPs' equipped with PLs to facilitate LM parcel delivery. These 'SPs' are placed at a finite set of candidate locations and must cover all demand points within a maximum walking distance. Their model and experiments are fully focused on unattended, limited-capacity lockers, where the 'SPs' purely refer to candidate locker locations, not to attended facilities. Demand is modelled at the level of aggregated zones, each with Poisson distributed parcel arrivals, and assigned to its nearest PL only. Partial fulfilment is not allowed in the base model, though the authors note that this assumption can be relaxed. The problem is formulated as a MILP using preprocessed piecewise-linear service-level functions. Unlike scenario-based approaches, stochasticity is embedded directly in the objective by incorporating expected rejection values derived from a Markov chain analysis. These expected values are approximated with convex piecewise-linear functions, allowing them to be captured within a linear objective. Numerical results demonstrate this piecewise linear model outperforms both deterministic and scenario-based benchmarks in scalability and cost-efficiency.

Kahr (2022) also incorporate locker-level dynamics by introducing modular, multi-compartment lockers and pickup modelling by a deterministic replenishment rate to simulate delayed pickups. The model jointly optimises locker location and internal layout, incorporating compartments tailored to different parcel sizes, an aspect largely neglected in prior literature. Demand is disaggregated at individual customer grid level and modelled stochastically across pandemic-driven scenarios using Austrian case data. Lockers can be installed at a discrete set of candidate locations, and each customer is assigned to at most one locker within a maximum coverage distance δ, with full flexibility in selecting lockers within this range. The model allows partial demand fulfilment system-wide: only customers within the distance threshold are served, but full fulfilment is required per assigned customer. Capacity constraints are strictly enforced; demand exceeding locker capacity is not redirected or delayed, and no overflow handling is allowed. The objective is to maximise expected returns from covered demand, subject to a budget constraint on locker installation. The solution method is based on an ILP formulation, supported by a scalable Benders decomposition approach for solving large-scale problem instances.

Robust optimisation and movable locker models under demand uncertainty

Like the study of Kahr (2022), Wang et al. (2022) also incorporate locker compartments into their optimisation framework, distinguishing between small and large parcels. However, Wang et al. (2022) use a robust optimisation model that jointly determines the location, number, and type of lockers per site. Lockers can only be installed at predefined demand sites, with demand aggregated by area rather than individual customers. Assignment is based on acceptable walking distances, and each customer is allocated to the nearest locker within a threshold radius.

Demand is modelled as uncertain but bounded. For each site, the number of incoming and outgoing parcels by parcel size is defined using known averages and maximum deviations. The model adopts the Bertsimas and Sim (2004) robust optimisation approach, using a robustness parameter (Γ) to control the level of protection against demand deviations. Overflow is not explicitly handled during operations; if lockers are insufficient, parcels are marked as unsatisfied demand and tracked. This allows the model to evaluate trade-offs between cost and service levels under constrained capacity. pickup behaviour is incorporated implicitly through deterministic locker dwell times. Parcels not picked up within a defined period (for example, 48 hours) are assumed to occupy locker space, influencing capacity planning. The formulation is solved as a MILP.

Wang et al. (2020) also utilise robust optimisation but focus on movable PLs that are deployed and returned daily. Lockers could be placed at predefined selected community points, with constraints reflecting local feasibility. Demand is modelled at the community or sub-community level, aggregated per point, and customers are assigned to the nearest locker within a predefined walking distance, based on Euclidean distances. The objective is to minimise total operational costs, including location-specific land rent, investment and maintenance costs, and travel expenses between the depot and locker deployment sites. To address demand uncertainty, the model also adopts the Bertsimas and Sim (2004) robust optimisation framework, using symmetric deviation intervals to manage variability. Overflow is handled indirectly: if local capacity is exceeded and no nearby fixed locker is available, parcels are returned to the depot for next-day redelivery. Although this does not trigger penalties in the objective function, it introduces realistic operational constraints. Only mobile PLs are considered, and partial demand fulfilment is permitted. Parcels that cannot be delivered due to capacity or distance constraints are excluded from the solution without penalty. The model is formulated as an ILP.

Behaviour and accessibility-focused approaches

Luo et al. (2022) formulate a multi-objective PL network design problem that jointly minimises total system cost and maximises customer accessibility. Accessibility is defined via an exponential distance-decay function applied to Euclidean travel distances, penalising long trips and reflecting diminishing customer willingness to walk. Capacity must match assigned demand, and no overflow or rerouting is allowed, if capacity is insufficient, the solution is infeasible. Locker placement is restricted to a discrete set of candidate nodes, which also serve as demand locations. Like in Deutsch and Golany (2018) and Sweidan et al. (2022), demand is modelled deterministically. It is defined per customer node and includes separate product types (general and fresh goods), based on population size and product type specific demand rates per customer location. Customers are assigned to lockers within a specified travel threshold ε, but accessibility preferences may override default recommendations. In terms of demand fulfilment, all demand assigned to lockers must be served. However, customers located beyond the distance threshold or without suitable recommendations may remain unassigned, allowing partial fulfilment across the network. To solve the model, the authors propose the Active Learning Pareto Evolutionary Algorithm (ALPEA), which jointly minimises total system cost and maximises customer accessibility to effectively navigate the trade-offs between these conflicting objectives. The algorithm is tested on 70 benchmark instances and a real-world case from Shenyang, China.

3.4 Contribution statement

While the literature on PL location models has grown significantly in recent years, key limitations remain when compared to the requirements of DHL's OOH network. Table 9 provides an overview of key model features across studies, including those introduced in this research to address the identified gaps.

First, no study integrates both PLs and SPs within a single capacitated facility PL location network model. SPs, despite their operational relevance, are either excluded, or treated separately. In contrast, DHL's network relies on the flexible and complementary use of both location types.

Second, although two studies address overflow or demand rejection, they rarely offer structured fallback mechanisms. Our research is the first to model two different fallback/delivery options, SPs and home delivery, alongside parcel locker delivery. When PLs exceed capacity, or when alternative delivery options are more cost-efficient, parcels may be delivered to nearby SPs or sent via home delivery. This overflow design reflects real-world decision-making and thereby improves the robustness of the delivery network and enhances cost control.

Third, demand allocation in most existing studies is modelled as a binary assignment: where each demand location is served by exactly one facility. Although Raviv (2023) allows fractional allocation only in cases of equidistant lockers, and Sweidan et al. (2022) apply predefined percentage splits across distance-based layers, this research is the first to introduce fully flexible fractional allocation across all feasible facilities, driven by cost and capacity optimisation.

Fourth, while some studies use uniform travel distances or layered zones, none support region-adaptive pickup radius. This research incorporates location-specific accessibility/travel willingness thresholds based on urban density or context, increasing the realism of demand allocation. This adaptive allocation is novel in the context of the placement of PLs.

Fifth, although robust optimisation has been applied before, it has not been explored in this DHL specific OOH context. Existing models assume binary demand assignment, and this research extends robustness to a setting with fractional allocation and a heterogeneous mix of OOH points.

In summary, this study contributes a novel MILP-based framework that:

- integrates PLs and SPs in a capacitated FLP for PLs;
- models overflow via fallback/delivery options to SPs and home delivery;
- allows full fractional (non-binary) demand allocation within a radius;
- supports region-adaptive pickup radii for geographic pickup realism;
- integrates robust optimisation for fractional demand and mixed facility types, while maintaining linear solvability.

The inclusion of novel features introduces elements not present in existing models. In addition, the combined integration of both new and established modelling features, as can be seen in comparison Table 9, into a single framework results in a unique approach not explored in previous studies. Together, this model contributes to key literature gaps and supports DHL's ambition of data-driven decision support in locker placement within its existing OOH network.

3.5 Summary & conclusion

In this chapter, we systematically reviewed academic literature relevant to the strategic placement of PLs, focussing on OOH networks that include both PLs and SPs in the LM, to answer the research question: *"What methods for strategic placement of OOH points are discussed in the literature?"* The goal was to identify state-of-the-art modelling approaches and assess how existing research aligns with the operational context of DHL.

The literature shows increasing academic interest in OOHD, yet the absolute number of studies remains relatively limited. Most studies focus solely on PLs and omit SPs, or model them in isolation without integrating both types in a single network. Among the studies reviewed, ten were identified as particularly relevant to DHL's core problem: determining both location and capacity for PLs. These studies were systematically compared on distinct modelling characteristics. This comparison uncovered several gaps in the existing literature.

First, no research considers both PLs and SPs within a capacitated PL location framework. Second, existing approaches mostly assume binary demand allocation, such that all demand from one demand location must be fulfilled by strictly one locker. Third, no study integrates a region-specific pickup radius. Fourth, while some studies consider overflow/rejection mechanisms, structured fallback/ alternative options such as delivery to SPs or home are not incorporated. Lastly, although robust optimisation has been applied, it is not yet extended to models with fractional allocation or different facility types.

Based on these findings, this thesis proposes a novel MILP-based framework that closes these identified gaps. Moreover, it introduces a unique combination of both novel and established modelling features, which have not yet been integrated into a single framework in existing literature. This combined approach is clearly visualised in Table 9, which highlights the distinctive contribution of this study. Together, these contributions provide both academic relevance and novelty, as well as practical value for DHL. The next chapter presents the mathematical model that addresses these gaps.

Table 9: Comparison of features across studies and this research

Feature Category / Criteria	**This Research**	Deutsch and Golany (2018)	Ottaviani et al. (2023)	Rabe et al. (2021)	Wang et al. (2020)	Luo et al. (2022)	Kahr (2022)	Raviv (2023)	Sawik et al. (2022)	Sweidan et al. (2022)	Wang et al. (2022)
OOH types											
Parcel lockers	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Service points	✓										
Locker placement											
Fixed predefined locations	✓		✓		✓		✓	✓			
Same as demand locations		✓		✓		✓			✓	✓	✓
Demand											
Aggregated demand	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
Stochastic demand	✓		✓	✓	✓	✓	✓	✓	✓		✓
Deterministic demand	✓	✓								✓	
Demand allocation											
Within radius/ area	✓		✓		✓		✓	✓	✓		✓
Linked to closest point	✓		✓		✓			✓			✓
Distance based		✓		✓		✓				✓	
Layered based		✓								✓	
Accessibility based						✓					
Region adaptive	✓										
Binary demand allocation		✓	✓	✓	✓	✓	✓		✓		✓
Fractional demand allocation	✓										
Fulfilment of demand											
Full	✓			✓				✓	✓		
Full within coverage						✓	✓			✓	
Constraint portion			✓								
Variable		✓			✓						✓
Locker capacity											
Limited	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓
Overflow handling											
None		✓	✓	✓		✓	✓		✓	✓	✓
Rejection/ postponement					✓			✓			
To different locker	✓				✓						
To service points	✓										
Home delivery	✓										
Solution method(s)											
Exact (MILP/ ILP)	✓	✓	✓		✓		✓	✓		✓	✓
Robust optimisation	✓				✓						✓
(Meta)Heuristics			✓			✓					
Hybrid exact/Simulation–optimisation				✓					✓		
Objective											
Minimise costs	✓		✓	✓	✓			✓	✓		✓
Maximise profit		✓					✓			✓	
Multi-criteria					✓			✓			

4 Model Formulation

This chapter presents the modelling approach used to support placement and sizing decisions for PLs in DHL's OOH network. The following research question is central in this chapter:

"How can a mathematical model be developed to support strategic placement and sizing decisions for PLs in DHL's OOH network?"

To answer this question, Section 4.1 first presents the strategic modelling approach and design assumptions, based on the findings and requirements identified in Chapter 2. Subsequently, Section 4.2 presents the base MILP formulation developed to address the strategic design objectives. Section 4.3 introduces several model extensions such as fixed existing locker infrastructure, region-adaptive pickup behaviour, and placement limitations. Section 4.3 also presents a robust optimisation extension. Finally, Section 4.4 integrates all components into a full MILP model, offering a unified formulation from which alternative model configurations can also be derived.

4.1 Modelling approach

The following list summarises and explains the core modelling principles and strategic design assumptions that guided the model design.

- **Strategic design objective:** The model is designed to support long-term placement and sizing decisions of PL within DHL's OOH network, aiming to prevent current mismatches between locker placement and demand, and to guide the planned sixfold strategic locker network expansion by 2030. While the primary focus lies on PLs, the model also considers the roles of SPs and home delivery, identifying where these alternatives may remain the most strategic or cost-effective solution without placing lockers. To enable such decisions, the model incorporates a detailed cost structure across all delivery modes, including different locker sizes, SPs, and home delivery, which are currently not used and unknown in DHL's decision-making processes. Together, this modelling aligns with the strategic objectives outlined in Chapter 2.
- **Strategic benchmark scenario:** As outlined in Chapter 2, another key requirement is to assess the strategic quality of existing placements by comparing them to an idealised greenfield solution. To enable this, a first model is created without considering existing infrastructure, providing a benchmark to evaluate which current lockers are well-placed and which are not. In a subsequent step, an extension is introduced that allows existing infrastructure to be fixed, supporting realistic expansion planning.
- **Demand representation:** Aligned with the model's strategic focus, demand is represented using average daily volumes per demand point, indicating how many FM and LM parcels are placed into lockers each day. Fine-grained demand distributions are not available, and DHL's upcoming dynamic capacity control system is expected to eliminate day-to-day demand variation across different demand points by dynamically redistributing parcels across nearby lockers, as explained in detail in Chapter 2. Therefore average daily demand offers a practical and robust input for long-term planning. Moreover, demand is defined as the number of parcels assigned to a locker over the course of a day. Due to the static nature of the optimisation model, all daily demand, both FM and LM, is assumed to be present simultaneously within the capacity calculations. This can significantly overstate peak occupancy, as it ignores natural fluctuations from parcel drop-off and pickup throughout the day (see Chapter 2). Nevertheless, this approach aligns with DHL's operational and financial perspective: what matters is not the number of parcels physically present at any moment, but the total number of parcels assigned to a locker over the day. As such, it provides a realistic input for strategic planning while also implicitly incorporating a buffer against intra-day variation.
- **Robust optimisation:** Despite the considerations above, a residual risk remains that multiple nearby demand points may simultaneously experience large, unexpected short-term peaks. To address this, the model includes a robust optimisation extension against worst-case demand scenarios in surrounding demand points, where it is uncertain which demand points will be affected, ensuring the network remains resilient under such worst-case scenarios if desired. Even though exact demand distributions are unknown, DHL can estimate critical peak loads using historical patterns and expert knowledge, which can be incorporated into the robust formulation. In addition to robustness, Chapter 5 uses experiments to explore how the model responds to structural demand growth. In these experiments, overall demand is scaled by multiple factors. This differs fundamentally from robustness: rather than guarding against local worst-case uncertainty, where it is unclear in advance where excess demand may occur, these scenarios simulate systematic volume increases across the

entire network. As such, they provide insight not only into long-term growth planning, but also into how the network can be adapted to handle sustained periods of higher demand such as longer seasonal peaks.

- **Parcel behaviour correction (LM/FM):** As outlined in Chapter 2, FM and LM parcels differ in their residence times, with LM parcels often remaining in lockers for multiple days. This behavioural difference impacts locker occupancy and is incorporated into the model by adjusting the effective locker capacity to account for longer residence times of parcels. This correction is further discussed in Chapter 5.
- **Realism-enhancing model extensions:** In addition to these core design assumptions, the model includes two extra extensions to better reflect real-world decision environments. A placement limit is introduced to simulate budgetary constraints and to prioritise locker locations with the highest strategic value. Moreover, region-specific pickup radius is implemented to account for varying willingness to travel in different geographic settings. These extensions are detailed in Section 4.3.

At a high level, these assumptions and modelling choices translate into a base model that operates on a single planning day, using fixed average daily parcel volumes per demand point $i \in I$. The model considers only one unified parcel demand type across all demand points. The set I represents a fixed, finite set of demand points. The model includes both a fixed set of existing and candidate locker locations, as well as a finite set of existing SP locations, together denoted by $j \in J$. SPs have fixed capacities and are activated at existing locations only when used in the solution, whereas each candidate locker site can have at most one locker of a predefined type $s \in S$, with associated capacity and installation cost. Parcel demand at each point can be distributed across multiple delivery modes and assigned to nearby facilities, including PLs, SPs, and home delivery. Lockers and SPs are only available when the distance to these facilities is within the allowable pickup radius, where δ_{ij} represents the distance between demand point i and facility j, while home delivery is always possible. All facility types and delivery modes are modelled with distinct fixed and variable costs. The resulting base model determines the optimal placement and sizing of PLs, as well as the optimal delivery mode allocation for each demand point.

4.2 Mathematical formulation

While various formulations have been proposed in the literature and discussed in Chapter 3 to support OOHD networks, many focus either on simplified assumptions or specific applications that do not fully reflect DHL's operational context. Therefore, this section introduces a novel and tailored mathematical formulation that explicitly considers the characteristics of DHL's OOH network, such as heterogeneous locker types, capacity constraints, partial demand allocation across multiple delivery options, and fallback/alternative delivery mechanisms such as home delivery and existing SPs. The formulation is presented below, starting with a structured overview of all sets, parameters, and decision variables involved. It presents the mathematical formulation of the Last Mile Capacitated Parcel Locker Location Problem (LMCPLLP), as introduced in Chapter 3, tailored to DHL's context.

Set and indices	Definition
$i \in I$	Demand locations
$j \in J$	All (possible) OOH facility locations
$j \in J^{\mathbf{L}} \subseteq J$	PL locations
$j \in J^{\mathbf{SP}} \subseteq J$	SP locations
$s \in S$	Locker types (e.g., small, medium, large)
$N(i) := \{j \in J \mid \delta_{ij} \leq R\}$	Facilities j within radius R of demand point i
$A(j) := \{i \in I \mid j \in N(i)\}$	Demand locations i assignable to facility j

Parameter	Definition
d_i	Parcel demand at demand location i
$f_s^{\mathbf{L}}$	Fixed cost of installing a PL of type s
$f^{\mathbf{SP}}$	Fixed cost of operating a SP
$c^{\mathbf{SP}}$	Variable cost per parcel assigned to a SP
$c^{\mathbf{L}}$	Variable cost per parcel assigned to a PL
$c^{\mathbf{H}}$	Variable cost per parcel assigned to home delivery
δ_{ij}	Distance between demand location i and facility j
c_s	Capacity of a PL of type s
c_j	Capacity of SP $j \in J^{\mathrm{SP}}$
R	Maximum allowable pickup radius
ϵ	Weighting factor for distance-based cost penalty
M	A sufficiently large constant used to enable conditional constraints (big-M method)

Variable	Definition
$x_{js} \in \{0, 1\}$	Binary variable; 1 if a locker of type s is installed at location $j \in J^{\mathrm{L}}$
$z_j \in \{0, 1\}$	Binary variable; 1 if SP $j \in J^{\mathrm{SP}}$ is activated for use
$q_{ij} \in \mathbb{Z}_+$	Number of parcels from demand location i assigned to facility $j \in N(i)$
$y_i^{\mathbf{H}} \in \mathbb{Z}_+$	Number of parcels from demand location i assigned to home delivery

$$
\begin{aligned}
\min\ z = & \sum_{i \in I} \sum_{j \in N(i) \cap J^{\mathrm{SP}}} c^{\mathrm{SP}} \cdot q_{ij} + \sum_{i \in I} \sum_{j \in N(i) \cap J^{\mathrm{L}}} c^{\mathrm{L}} \cdot q_{ij} + \sum_{i \in I} c^{\mathrm{H}} \cdot y_i^{\mathrm{H}} \\
& + \sum_{j \in J^{\mathrm{L}}} \sum_{s \in S} f_s^{\mathrm{L}} \cdot x_{js} + \sum_{j \in J^{\mathrm{SP}}} f^{\mathrm{SP}} \cdot z_j \\
& + \epsilon \cdot \sum_{i \in I} \sum_{j \in N(i)} \delta_{ij} \cdot q_{ij} \qquad (1)
\end{aligned}
$$

s.t.

$$\sum_{j \in N(i)} q_{ij} + y_i^{\mathrm{H}} = d_i \qquad \forall i \in I \qquad (2)$$

$$\sum_{i \in A(j)} q_{ij} \le M \cdot \sum_{s \in S} x_{js} \qquad \forall j \in J^{\mathrm{L}} \qquad (3)$$

$$\sum_{i \in A(j)} q_{ij} \le M \cdot z_j \qquad \forall j \in J^{\mathrm{SP}} \qquad (4)$$

$$\sum_{i \in A(j)} q_{ij} \le \sum_{s \in S} c_s\, x_{js} \qquad \forall j \in J^{\mathrm{L}} \qquad (5)$$

$$\sum_{i \in A(j)} q_{ij} \le c_j \qquad \forall j \in J^{\mathrm{SP}} \qquad (6)$$

$$\sum_{s \in S} x_{js} \le 1 \qquad \forall j \in J^{\mathrm{L}} \qquad (7)$$

$$q_{ij} \ge 0 \qquad \forall i \in I,\ j \in N(i) \qquad (8)$$

$$y_i^{\mathrm{H}} \ge 0 \qquad \forall i \in I \qquad (9)$$

$$x_{js} \in \{0, 1\} \qquad \forall j \in J^{\mathrm{L}},\ s \in S \qquad (10)$$

$$z_j \in \{0, 1\} \qquad \forall j \in J^{\mathrm{SP}} \qquad (11)$$

The objective function (1) minimises the total cost of the system, including variable delivery costs for each option, fixed installation costs for lockers and usage costs of SPs, and a distance-weighted penalty (soft constraint) to encourage the assignment to nearby facilities, thereby promoting accessibility. Constraint (2) ensures that the entire parcel demand at each location is either assigned to an OOH facility or fulfilled via home delivery. Constraint (3) ensures that lockers can only be used if they are installed, and constraint (4) ensures that SPs can only be used when activated. Constraints (5) and (6) enforce capacity limits at each facility, using either locker-type capacity or the known capacity of SPs. Constraint (7) ensures that at most one locker type can be installed at each location. Constraints (8) and (9) define non-negativity for parcel OOH facility assignments and home delivery, respectively. Finally, Constraints (10) and (11) impose binary restrictions on locker and SP installation/usage decisions, completing the MILP formulation.

4.3 Model extensions

While the base formulation in Section 4.2 captures DHL's core operational logic, several extensions are modelled to reflect real-world constraints and increase practical applicability. This section introduces four such enhancements. Section 4.3.1 models the fixed placement of existing lockers. Section 4.3.2 imposes a limit on the number of new lockers that can be installed. Section 4.3.3 introduces the region-adaptive pickup radius to better reflect geographical differences. Finally, Section 4.3.4 presents a robust optimisation formulation that protects locker locations against worst-case demand scenarios, where multiple nearby demand points may simultaneously experience large, unexpected short-term peaks.

4.3.1 Fixed placement of existing lockers

This extension incorporates the current network of PLs by distinguishing between existing infrastructure and candidate locations. Existing lockers are considered fixed: they cannot be relocated, removed, or resized, and their capacity is pre-assigned.

Updated sets, parameters and constraints

Set	Definition
$J^{\mathbf{L}}$	Set of all PL locations (existing and candidate)
$J^{\mathbf{L}}_{\text{new}} \subseteq J^{\mathbf{L}}$	Set of candidate locker locations
$J^{\mathbf{L}}_{\text{exist}} \subseteq J^{\mathbf{L}}$	Set of existing lockers with fixed placement and capacity
S	Set of locker types, including the special type `EXIST`

Parameter	Definition
c_s	Capacity of locker type $s \in S$ (with $c_{\texttt{EXIST}} = 0$)
c_j^{exist}	Fixed capacity of existing locker $j \in J^{\mathrm{L}}_{\text{exist}}$
f_s	Fixed cost of installing locker type $s \in S$ (with $f_{\texttt{EXIST}} = 0$)

The capacity constraint is updated as follows to distinguish between new, existing, and SP locations:

$$\sum_{i \in A(j)} q_{ij} \leq \begin{cases} \sum_{s \in S} c_s \cdot x_{js} & \text{if } j \in J^{\mathrm{L}}_{\text{new}} \\ c_j^{\text{exist}} & \text{if } j \in J^{\mathrm{L}}_{\text{exist}} \\ c_j & \text{if } j \in J^{\mathrm{SP}} \end{cases} \tag{12}$$

To ensure correct assignment of locker types, we explicitly fix the `EXIST` type to all predefined lockers and prohibit its use at candidate sites:

$$x_{j,\texttt{EXIST}} = 1 \qquad \forall j \in J^{\mathrm{L}}_{\text{exist}} \tag{13}$$

$$x_{j,\texttt{EXIST}} = 0 \qquad \forall j \in J^{\mathrm{L}}_{\text{new}} \tag{14}$$

4.3.2 Maximum number of lockers to place

This extension reflects real-world constraints, such as investment budgets or spatial limitations, that can restrict the number of PLs that can be installed. To enforce an upper bound on the total number of lockers in the network, the following constraint is introduced:

$$\sum_{j \in J^{\mathrm{L}}} \sum_{s \in S} x_{js} \leq U \tag{15}$$

Here, U denotes the maximum number of lockers that can be installed. If this extension is combined with the fixed-locker extension, where a subset of lockers is already pre-installed, the value of U must be adjusted accordingly to ensure that the upper bound only applies to newly placed lockers. The adjusted value then becomes:

$$U^{\text{adjusted}} = U + |J^{\mathrm{L}}_{\text{exist}}| \tag{16}$$

4.3.3 Region-adaptive pickup radius

In the model formulation in Section 4.2, a fixed radius R determines which facilities can be assigned to each demand location. However, a single global threshold may not accurately reflect local differences in urban density, infrastructure, or consumer preferences.

This extension replaces the global radius R with a location-specific value r_i, allowing more granular control over assignment feasibility. Let r_i denote the maximum pickup radius for demand location $i \in I$. The neighbourhood set $N(i)$ is then redefined as:

$$N(i) := \{j \in J \mid \delta_{ij} \leq r_i\} \tag{17}$$

4.3.4 Robust optimisation

To account for uncertain parcel demand at locker locations, this extension incorporates a robust optimisation approach based on the budgeted uncertainty model by Bertsimas and Sim (2004). Among the reviewed studies that incorporate demand uncertainty, a few adopt robust optimisation (Wang et al., 2020, 2022). Both studies employ the budgeted uncertainty framework introduced by Bertsimas and Sim (2004).

The Bertsimas and Sim (2004) model offers a number of advantages over alternative robust and stochastic methods. In contrast to purely scenario-based stochastic approaches, it does not require a full probability distribution of uncertain parameters, which are also unavailable in this study. Moreover, stochastic approaches often rely on independence assumptions or require known correlation structures between uncertain parameters, which are difficult to estimate in practice. In contrast, robust optimisation guarantees feasibility under worst-case conditions, whereas stochastic methods optimise for expected performance and may underperform in tail-risk scenarios, which is relevant for DHL due to potential demand spikes during peak moments. In addition, a key advantage of this robust approach is the interpretability and tunability of robustness through the Γ parameter. For each locker location j, this parameter specifies the uncertainty budget, which bounds the total intensity of demand deviations that may occur across all demand points within the travel radius of the locker. When $\Gamma_j = 0$, the model behaves fully deterministically (i.e., no robustness). When Γ_j equals the number of demand points within the travel radius of the locker, the model is fully conservative, as it protects against the case where all demand locations within the range of locker j simultaneously experience their maximum deviation. This allows for a tunable trade-off between cost-efficiency and protection. In the context of DHL, this tunability is particularly relevant: the company must safeguard against localised demand spikes, but also keep costs under control. By adjusting Γ, the planner can explore a range of risk scenarios and evaluate robustness-versus-cost trade-offs in a transparent and controlled way. Moreover, the Bertsimas and Sim (2004) formulation keeps the problem computationally manageable: the robust capacity constraints remain linear and can be solved as part of a MILP, whereas many alternative robust formulations introduce non-linearities that are difficult to solve. Similarly, stochastic models often require solving the model across many scenarios, which significantly increases computational effort and reduces scalability. This linearity ensures that the model remains solvable for realistically sized instances and enables efficient exploration of different robustness levels, allowing planners to assess these trade-offs.

For these reasons, this study adopts the Bertsimas and Sim (2004) framework, which has also been applied in earlier work such as Wang et al. (2022). However, whereas Wang et al. (2022) applied the framework in a simplified setting, using a binary demand assignment model and only two locker types, this study extends its application to a heterogeneous locker network with multiple facility types and partial demand allocation. This extension enables robust optimisation in settings where parcel demand is flexibly distributed across multiple nearby locations and sufficient capacity must be ensured to be protected against uncertain, localised peaks.

Uncertainty is introduced specifically in the capacity constraint, as defined in this chapter, and forms the basis of the new Constraint (18), since locker capacity is one of the model's key decision variables. Let d_i denote the nominal parcel demand at demand location $i \in I$, and let Δ_i represent the maximum possible (worst-case) upward deviation (in %) from this value. The true demand at location i may thus reach up to $d_i + \Delta_i$. To account for such deviations, auxiliary variables $k_i \in [0, 1]$ are introduced in the inner maximisation problem, representing the fraction of the maximum demand deviation Δ_i that may occur at demand location $i \in A(j)$. These variables are bounded by an uncertainty budget Γ_j, which limits the total intensity of deviation that locker $j \in J^{\mathrm{L}}$ must be protected against. This is expressed as $\sum_{i \in A(j)} k_i \leq \Gamma_j$, allowing the model to account for worst-case demand increases distributed across multiple nearby demand locations, without assuming that all deviations occur simultaneously.

$$\sum_{i \in A(j)} q_{ij} + \max_{\substack{k_i \in [0,1] \\ \sum_{i \in A(j)} k_i \leq \Gamma_j}} \sum_{i \in A(j)} \Delta_i k_i q_{ij} \leq \sum_{s \in S} c_s \cdot x_{js} \qquad \forall j \in J^{\mathrm{L}} \tag{18}$$

Constraint (18) ensures that, for each locker location, the model protects against the worst-case combination of demand deviations by allowing fractional deviations $k_i \in [0, 1]$ from each nearby demand point $i \in A(j)$, with a total deviation intensity bounded by Γ_j. This allows the model to account for a distributed overload risk across multiple demand locations, without assuming that all of them fully deviate simultaneously. The cumulative impact of these deviations is incorporated directly into the locker's capacity constraint, thereby ensuring robustness. The parameter Γ_j controls the level of protection: higher values safeguard against more or larger combined deviations, while lower values assume only limited deviation pressure. The magnitude of the robust adjustment is proportional to the quantity assigned from each demand location, meaning that demand points with more assigned parcels contribute more to the locker's potential overload if deviating.

Linear reformulation via duality

The robust locker capacity constraint in Equation (18) includes a maximisation term. This formulation introduces non-linearity due to the inner maximisation problem. To enable integration into a linear MILP framework, we apply duality theory to reformulate this maximisation as a set of linear constraints.

The inner maximisation problem seeks the combination of demand deviations (subject to the uncertainty budget Γ_j) that leads to the largest cumulative overload at locker j. Since this is a linear program over auxiliary variables

$k_i \in [0,1]$, with a linear objective and linear constraints, it satisfies the conditions for strong duality.

$$\max_{\substack{k_i \in [0,1] \\ \sum_{i \in A(j)} k_i \leq \Gamma_j}} \sum_{i \in A(j)} \Delta_i k_i q_{ij} \qquad \forall j \in J^{\mathrm{L}} \tag{19}$$

Due to strong duality, the inner maximisation of equation (19) can be replaced by its dual formulation, preserving the optimal value and ensuring linear compatibility with the MILP model.

Updated sets, parameters and constraints

Parameter	Definition
Γ_j	Uncertainty budget for locker $j \in J^{\mathrm{L}}$
Δ_i	Fractional deviation parameter for demand location i (e.g., 0.2 = 20% increase)

Variable	Definition
$\pi_j \geq 0$	Dual variable associated with constraint $\sum k_i \leq \Gamma_j$; marginal value of allowing one additional unit of total deviation
$\theta_{ij} \geq 0$	Dual variable associated with constraint $k_i \leq 1$; marginal value of forcing location i to fully deviate

The dual reformulation of the robust locker capacity constraint is:

$$\sum_{i \in A(j)} q_{ij} + \Gamma_j \cdot \pi_j + \sum_{i \in A(j)} \theta_{ij} \leq \sum_{s \in S} c_s \cdot x_{js} \qquad \forall j \in J^{\mathrm{L}} \tag{20}$$

To ensure the feasibility of the dual formulation, the following supporting constraints are added:

$$\pi_j + \theta_{ij} \geq \Delta_i q_{ij} \qquad \forall j \in J^{\mathrm{L}},\ i \in A(j) \tag{21}$$

$$\pi_j \geq 0, \quad \theta_{ij} \geq 0 \qquad \forall j \in J^{\mathrm{L}},\ i \in A(j) \tag{22}$$

These constraints form the linear reformulation of the robust locker capacity constraint. Constraint (20) ensures that each locker has enough capacity to handle both the nominal demand and a worst-case deviation, captured via the dual variables. Constraints (21) and (22) are supporting constraints that define how large the dual variables must be to protect against the allowed deviations in demand (within budget Γ_j). Constraint (21) ensures that the total compensation provided by the dual variables is sufficient to cover the worst-case deviation of each demand point $i \in A(j)$. Constraint (22) enforces non-negativity of these variables. Together, they ensure that the overload adjustment in (20) is large enough to protect the locker against demand deviations created by the uncertainty budget Γ_j. The three constraints replace the non-linear max-term and allow the full model to remain linear.

4.4 Full model integration

This section shows a full model that integrates all extensions into a single robust MILP framework. This includes fixed existing lockers, a maximum new locker installation limit, region-adaptive pickup radii, and robust capacity constraints under demand uncertainty. The model is designed in a modular way, allowing alternative versions, where only a subset of the extensions is activated, to be derived directly from the base formulation in Section 4.2, the explanations in Section 4.3, and this formulation. It is often more convenient to start from the base model when using only one or two extensions. However, for more substantial changes, such as robust optimisation or fixing existing lockers, it is more practical to work directly with the full model due to their greater impact on the overall structure, while still allowing easy removal of parts you do not wish to use. Moreover, a key advantage of this modularity is that it facilitates flexible implementation, where different model options can be enabled or disabled through simple conditional statements in programming code.

In this formulation, robust capacity constraints are only applied to new locker locations (J^{L}new), as existing lockers (J^{L}exist) have fixed capacities and cannot be resized. While this slightly reduces overall robustness, it reflects the real-world constraint that the existing infrastructure is not adjustable. If existing lockers were modelled as flexible, the robust formulation could be extended to cover the entire locker set.

$$
\begin{aligned}
\min\ z = &\sum_{i\in I}\sum_{j\in N(i)\cap J^{\mathrm{SP}}} c^{\mathrm{SP}}\cdot q_{ij} + \sum_{i\in I}\sum_{j\in N(i)\cap J^{\mathrm{L}}} c^{\mathrm{L}}\cdot q_{ij} + \sum_{i\in I} c^{\mathrm{H}}\cdot y_i^{\mathrm{H}} \\
&+ \sum_{j\in J^{\mathrm{L}}}\sum_{s\in S} f_s^{\mathrm{L}}\cdot x_{js} + \sum_{j\in J^{\mathrm{SP}}} f^{\mathrm{SP}}\cdot z_j \\
&+ \epsilon\cdot\sum_{i\in I}\sum_{j\in N(i)} \delta_{ij}\cdot q_{ij} && (23)
\end{aligned}
$$

s.t.

$$
\begin{aligned}
\sum_{j\in N(i)} q_{ij} + y_i^{\mathrm{H}} &= d_i && \forall i\in I && (24)\\
\sum_{i\in A(j)} q_{ij} &\le M\cdot\sum_{s\in S} x_{js} && \forall j\in J^{\mathrm{L}}_{\mathrm{new}} && (25)\\
\sum_{i\in A(j)} q_{ij} &\le M\cdot z_j && \forall j\in J^{\mathrm{SP}} && (26)\\
\sum_{i\in A(j)} q_{ij} &\le \sum_{s\in S} c_s\, x_{js} && \forall j\in J^{\mathrm{L}}_{\mathrm{new}} && (27)\\
\sum_{i\in A(j)} q_{ij} &\le c_j^{\mathrm{exist}} && \forall j\in J^{\mathrm{L}}_{\mathrm{exist}} && (28)\\
\sum_{i\in A(j)} q_{ij} &\le c_j && \forall j\in J^{\mathrm{SP}} && (29)\\
\sum_{s\in S} x_{js} &\le 1 && \forall j\in J^{\mathrm{L}} && (30)\\
\sum_{j\in J^{\mathrm{L}}}\sum_{s\in S} x_{js} &\le U + |J^{\mathrm{L}}_{\mathrm{exist}}| && && (31)\\
\sum_{i\in A(j)} q_{ij} + \Gamma_j\cdot\pi_j + \sum_{i\in A(j)} \theta_{ij} &\le \sum_{s\in S} c_s\cdot x_{js} && \forall j\in J^{\mathrm{L}}_{\mathrm{new}} && (32)\\
\pi_j + \theta_{ij} &\ge \Delta_i q_{ij} && \forall j\in J^{\mathrm{L}}_{\mathrm{new}},\ i\in A(j) && (33)\\
\theta_{ij} &\ge 0 && \forall j\in J^{\mathrm{L}}_{\mathrm{new}},\ i\in A(j) && (34)\\
\pi_j &\ge 0 && \forall j\in J^{\mathrm{L}}_{\mathrm{new}} && (35)\\
x_{j,\mathrm{exist}} &= 1 && \forall j\in J^{\mathrm{L}}_{\mathrm{exist}} && (36)\\
x_{j,\mathrm{exist}} &= 0 && \forall j\in J^{\mathrm{L}}_{\mathrm{new}} && (37)\\
q_{ij} &\ge 0 && \forall i\in I,\ j\in N(i) && (38)\\
y_i^{\mathrm{H}} &\ge 0 && \forall i\in I && (39)\\
x_{js} &\in \{0,1\} && \forall j\in J^{\mathrm{L}},\ s\in S && (40)\\
z_j &\in \{0,1\} && \forall j\in J^{\mathrm{SP}} && (41)
\end{aligned}
$$

4.5 Summary & conclusion

This chapter addressed the research question: *"How can a mathematical model be developed to support strategic placement and sizing decisions for parcel lockers in DHL's OOH network?"*

Within this chapter, a novel, modular MILP-based optimisation model was developed that aligns with DHL's strategic planning needs. The base formulation incorporates key real-world complexities relevant to DHL's operations, such as heterogeneous locker types, partial demand allocation, capacity constraints, and alternative delivery modes.

In addition, several extensions were introduced to reflect real-world constraints and operational complexities more accurately, including fixed infrastructure, placement limits, region-adaptive pickup radii, and robustness against local worst-case demand uncertainty. These features were integrated into a unified modular model that supports both greenfield design and incremental expansion, making it suitable as a long-term decision support tool.

Chapter 5 puts the model into practice by describing how it is applied to real-world data. It outlines the data preparation, parameter settings, and findings, introduces the input and solution GUI, and presents the experimental design used to evaluate the behaviour of the model in various scenarios.

5 Experimental Settings

This chapter answers the research question:

"How should the developed model be implemented for real-world application, and which experimental setup should be used to evaluate it?"

To answer this question, we first provide a complete overview of how the optimisation model is constructed, prepared, and applied. We explain how the required input data and parameters are collected, preprocessed, and defined. Furthermore, we describe how the model can be operated via a Graphical User Interface (GUI), and how the interactive output is presented to the user. Finally, we introduce the experimental design used to evaluate the performance of the model under different strategic scenarios.

Section 5.1 explains how the model is coded, solved, and configured within Python. Section 5.2 outlines the preprocessing steps and calculations applied to spatial data, demand data, facility data/assignments, and other relevant input parameters. Section 5.3 sets out the cost structures and capacity values defined during this research and used during the experiments. Section 5.4 presents the interactive tool through which DHL can use, run, and interpret model results. Section 5.5 then introduces the full experimental setup consisting of seven targeted phases.

5.1 Model implementation and solver environment

The MILP model is implemented in Python using the `Pyomo` package. A modular architecture is adopted to define sets, parameters, variables, constraints, and the objective function in clearly structured blocks. This design enables flexible integration of extensions such as robust optimisation, maximum number of new PLs, fixed existing infrastructure, and region-adaptive pickup behaviour.

All instances are solved using the `Gurobi Optimizer` with academic licence. Computations are performed on a Lenovo-27080 laptop equipped with an Intel® Core™ i7-9750H processor (6 physical cores, 12 threads, base frequency 2.60GHz) and 16 GB of RAM.

5.2 Input data & preprocessing

This section outlines the sequence of data-related and programming-related steps taken to operationalise the optimisation model. Rather than detailing every implementation aspect, we focus on those components that are essential for understanding how the model was applied to real-world data, and for ensuring transparency and reproducibility of the results.

First, Section 5.2.1 introduces the demand estimation method. Second, Section 5.2.2 explains how spatial reference data and hub mappings are combined to define the modelling geographic area scope. Section 5.2.3 presents the approach to calculate distances and generate feasible assignments across demand points and candidate OOH facility locations, while Section 5.2.4 accounts for the effect of parcel pickup behaviour on locker capacity. Section 5.2.5 discusses the generation of candidate OOH facility locations. Finally, Section 5.2.6 and Section 5.2.7 describe how key elements of two model extensions, robust demand formulation and region-adaptive pickup radii, are operationalised in practice, including how their input data and parameters are preprocessed, calculated and integrated into the experimental workflow.

5.2.1 Demand estimation and spatial allocation

While the model requires an average daily parcel volume per demand point $i \in I$, the implemented demand allocation method starts from an estimate of the total national parcel volume expected for the upcoming year, provided by the user along with the expected share of OOH deliveries. The resulting OOH volume serves as the basis for the demand allocation process.

Because OOH shipments are not linked to the receiver's home address in the available DHL data, demand allocation is based on home delivery parcel distributions. The total OOH demand is allocated to postcode areas using the historical distribution of home delivery parcel volumes. Specifically, internal shipment data from November 2024 to March 2025 is used to determine the relative share of home deliveries per postcode area. These fractions are then used to allocate the total OOH volume across the postcode areas, under the assumption that areas with higher home delivery activity also exhibit higher OOH demand.

The user can select how specific they want to model the demand by choosing the desired postcode level: `PC4`, `PC5`, or `PC6`. These levels represent different granularities in the Dutch postcode system, with `PC4` covering larger regions, and `PC5` and `PC6` providing smaller areas. Once a level is selected, the corresponding postcode geometries are converted to centroid points, which are then used as demand locations $i \in I$ in the mathematical model. Each centroid receives an integer demand value d_i, representing the expected number of OOH parcels per day for that specific postcode area.

5.2.2 Geographical reference data

The spatial foundation of the model is constructed from two open-access datasets published by Centraal Bureau voor de Statistiek (2023). Geometries are sourced from the 2023 shapefiles (`.gpkg`), which provide polygon boundaries for all Dutch postcode levels: `PC6`, `PC5`, and `PC4`, e.g., 8000AA, 8000A, and 8000 respectively. In parallel, demographic and housing characteristics per postcode are extracted from tabular data published in Excel format (e.g., population and household counts as of 1 January 2023). After import, spatial and statistical sources are merged by postcode identifier, such that each geometry is annotated with the available population figures. Based on this setup, the user can specify the desired postcode level (`PC4`, `PC5`, or `PC6`) as an input, determining the level at which demand is modelled and aggregated. This enables flexible control over the spatial granularity used in the optimisation process.

DHL-specific modelling scope definition To tailor the model to DHL's operational structure, internal lookup tables are used to associate each `PC4` area with a specific CityHub or RegioHub area (areas defined by DHL). These mappings are joined with the CBS postcode geometries to enable geographic filtering of the model input. The user may specify a DHL CityHub or RegioHub area as the modelling scope. Based on this selection, the model automatically determines which postcode areas are included in the optimisation run.

5.2.3 Distance calculation and facility assignment

To efficiently identify feasible assignments between demand points and OOH locations, all longitude and latitude coordinates are projected onto a flat Cartesian plane using a simple flat-Earth approximation. Given the relatively small spatial extent of the assignment problem (typically within a few kilometres), Earth curvature effects are neglected, which allows faster 2D distance calculations.

A `cKDTree` structure from the `scipy.spatial` module is used to quickly find all facilities within the pickup radius of each demand point. This results in a filtered subset of candidate facilities per location, without having to calculate all combinations of distances between facilities and demand points. Finally, the exact Euclidean distances between all valid demand–facility pairs within the radius are computed and passed to the optimisation model as the basis for assignment feasibility.

5.2.4 Capacity adjustment based on pickup behaviour

To account for realistic locker occupancy levels, we adjust the available capacity based on the observed pickup behaviour of the LM parcels. Since some parcels remain uncollected for multiple days, part of the locker capacity may already be occupied at the time of new deliveries. This effect is captured using the historical pickup distribution introduced in Chapter 2.

Let T denote the pickup time in hours since delivery. The Cumulative Distribution Function (CDF), $F(\tau) = P[T \leq \tau]$, gives the proportion of parcels picked up within τ hours. The corresponding survival function then represents the share of parcels still present at time τ:

$$S(\tau) = P[T > \tau] = 1 - F(\tau) \tag{42}$$

Let $k \in \{0, 1, \dots, K-1\}$ denote the number of days since delivery, with $k = 0$ representing the delivery day itself. The share of parcels from that delivery still present at the start of day k is then given by:

$$r_k = S(k), \qquad k = 0, 1, \dots, K-1 \tag{43}$$

Since DHL removes parcels after $K = 7$ days, as explained in Chapter 2, the average time a parcel resides in a location (in days) is approximated by:

$$\Phi = \sum_{k=0}^{K-1} r_k, \qquad \text{with } r_0 = 1 \tag{44}$$

This formulation uses a discrete-time structure, with one step per day, rather than a continuous model (integral). This reflects DHL's operational setup: new parcels are delivered once per day, typically around 12:00. Since this delivery moment is fixed, a discrete-time measurement is chosen.

Now, suppose that the demand location i has an average daily inflow of d_i LM parcels. The average number of parcels still occupying space from earlier deliveries can be estimated by:

$$\text{Expected leftover occupancy} = d_i \cdot \sum_{k=1}^{6} r_k \tag{45}$$

At the moment of delivery, location space is partially occupied by parcels delivered by DHL one to six days earlier that have not yet been picked up. Parcels from exactly seven days ago are removed, and today's deliveries are just arriving. This correction is implemented by adjusting the available location capacity. For a location with capacity c_j, the effective capacity becomes:

$$c_j^{\text{eff}} = c_j - d_i \cdot \sum_{k=1}^{6} r_k \tag{46}$$

This approach is closely related to Little's law, which states that the expected number of items in a system equals the arrival rate times the expected time spent in the system. In our setting, this translates to the expected number of parcels in a location being $d_i \cdot E[T]$, where d_i is the average daily inflow, and $E[T]$ is the average time a parcel spends in the location (in days), which in discrete form is approximated by $E[T] \approx r_1 + r_2 + \cdots + r_6$.

5.2.5 OOH locations

Candidate PL locations are generated using a user-defined input on a rectangular grid with a specified spacing in the east–west (x-axis) and north–south (y-axis) directions. Smaller grid spacing leads to a denser set of candidate locations, allowing for more granular placement. These candidate locations are generated within the geographical boundaries defined by the selected modelling area (e.g., one or multiple CityHub areas), as specified by the user. For example, a spacing of 500 metres in both the x- and y-axes places a candidate point every 500 metres along each axis, creating a uniform grid across the defined area. Existing PLs and SPs operated by DHL are included based on a merged dataset that we constructed from multiple internal sources. This combined dataset integrates different aspects such as longitude and latitude information, to identify the locations of all known PLs and SPs.

5.2.6 Robustness extension: Γ formulation

To ensure interpretability and scalability, the Γ parameter is implemented using a percentage-based specification. Rather than assigning a fixed absolute uncertainty budget Γ_j to each locker individually, we define it as a fraction of the number of demand locations within the pickup radius of the locker j, resulting in a dynamically computed value Γ_j, which represents the total deviation budget allowed across demand points within the reach of the locker.

In all subsequent experiments, references to Γ_j reflect this percentage-based interpretation: the user provides a percentage input, which is translated into a deviation budget based on the number of demand points around each locker, and used directly in the robust formulation described in Chapter 4. As a result, lockers in densely populated areas are assigned proportionally larger deviation budgets, whereas those in rural areas receive smaller budgets. This approach improves scalability, ensures consistency across locations, and provides a simple and intuitive way to control robustness in practice.

5.2.7 Region-adaptive pickup radius

When using the region-adaptive pickup radius, each demand location is linked to its corresponding `PC4` area and assigned a pickup radius based on the population density class of that area. This design choice mitigates local anomalies, such as a single high-rise building in an otherwise low-density area, that could distort radius assignment if evaluated at finer level (e.g., `PC6`). Demand locations are therefore first linked to their respective `PC4` area and then matched to a preprocessed-`PC4`-population density dataset to determine the radius.

The following population density thresholds are used in the experimental model setup, ranging from ultra-dense urban areas (≥10,000 inh./km^2) to very rural regions (<90 inh./km^2). Each class is mapped to a corresponding pickup radius, as shown in Table 10. These thresholds were developed in collaboration with and validated by DHL domain specialists, based on a national analysis of population density distributions to ensure their practical applicability. This aims to improve realism by maintaining granularity in urban centres while extending reach in rural areas, without the need to run multiple fixed scenarios manually.

Table 10: Mapping of population density to pickup radius

Population Density d (inh./km^2)	Radius r (m)
$d \geq 10{,}000$	400
$5{,}000 \leq d < 10{,}000$	600
$2{,}213 \leq d < 5{,}000$	800
$324 \leq d < 2{,}213$	1,200
$90 \leq d < 324$	2,000
$d < 90$	3,500

Table 11: Cost parameters per delivery method

Type	Fixed	Per-parcel
Locker S	€[redacted]	€[redacted]
Locker M	€[redacted]	€[redacted]
Locker L	€[redacted]	€[redacted]
Locker XL	€[redacted]	€[redacted]
Locker XXL	€[redacted]	€[redacted]
Service Pt	€[redacted]	€[redacted]
Home Dlv.	—	€[redacted]

Table 12: Defined locker types and their capacities

Type	Capacity
Locker S	60
Locker M	90
Locker L	120
Locker XL	160
Locker XXL	220
Service Pt	150

5.3 Cost and capacity parameters settings

As previously highlighted in Chapter 2, DHL's current decision-making does not incorporate a detailed cost breakdown across delivery modes and locker types. To address this gap, this thesis constructs a detailed and structured cost overview based on multiple internal sources and operational cost categories. These cost parameters were developed and subsequently validated with a DHL domain specialist, ensuring that they are both realistic and practically applicable. This financial modelling enables the optimisation framework to evaluate not only where capacity is required, but also which type of OOH infrastructure is cost-effective under varying demand or spatial conditions. The PL capacity values used in this study (e.g., 60 to 220 parcels) are based on the current distribution of PL sizes in DHL's existing network, as previously analysed in Chapter 2. These values reflect realistic capacity levels that DHL can operate and were validated with a DHL domain specialist for use as parameter inputs in the model.

PLs are divided into five capacity categories, each with a fixed daily cost reflecting both operational and installation/ investment-related expenses and a uniform per-parcel handling fee. SPs have a single fixed daily usage cost and a higher per-parcel handling fee. Home delivery is modelled using only a per-parcel cost. Tables 11 and 12 present the full cost and capacity values used throughout the experiments.

5.4 Graphical user interface

To support decision-making for employees at DHL, an interactive GUI application was developed as a front-end of the optimisation model. The GUI provides a user-friendly interface for configuring model inputs and parameters, and visualises the resulting solutions in real time. It is implemented in Python using the packages `PySide6`, and `folium` as a base. Figure 20(a) visualises the input GUI.

Users can configure key model inputs, parameters, and settings through intuitive widgets, including:

- selection of DHL's CityHub or RegioHub regions as model scope;
- selection of the granularity level for demand modelling, if the user wants to model demand on `PC4`, `PC5` or `PC6` level;
- specification of yearly parcel volume and share of OOH deliveries;
- choice between fixed or region-adaptive pickup radius;
- configuration of grid spacing for candidate locker locations;
- constraints on the number of new lockers to place;
- fixation of existing infrastructure or a greenfield scenario;

- the maximum time the solver is allowed to run (after which the best solution found is returned, even if optimality has not been proven);
- the folder where the solution map and solution result text file should be saved.

The input GUI provides real-time feedback during the optimisation process by displaying the complete console output of the Python script. This allows the user to monitor the progress of the model as it runs. During execution, users can follow intuitive print statements that provide insight into the progress of the model, including data processing steps, and important information such as the status of the solver, the gaps, and the solving times. This feature keeps the user informed about key aspects of the optimisation process, helping them understand where the model is in its execution.

Upon completion, the solution is visualised in an interactive HTML map that automatically opens in the user's default browser. This map visualises the optimal placement of PLs and the usage of SPs and home delivery, giving users a straightforward way to evaluate the solution. The interactive solution allows the selection of different layers on the map (such as demand points, used/unused lockers, and SPs) for selective visibility, facilitating a deeper analysis of the model's output via the control panel in the upper right corner. For example, users can choose to display only the new lockers that need to be placed in the optimal solution, making the analysis easier to understand. Figure 20(b) shows the interactive solution map with all layers enabled.

Additional functionality allows users to inspect allocation lines, showing which demand point is served by which locker, and to click on individual lockers, SPs, or demand points to access detailed information. For each element, users can view relevant attributes such as the assigned locker size and used capacity, a breakdown of how many parcels are assigned from each originating postcode, the delivery mode split at each demand location (e.g., locker, service point, home), and the pickup radius applied at that location, as shown in Figure 20(c). These interactive features enable DHL decision-makers to drill down into specific results.

A legend panel in the bottom right corner not only explains the map symbols but also provides a summary of the solution, including the total number of new lockers placed, the number of SPs used, and the overall objective value. Together, these interactive elements support both evaluation and deeper analysis, helping planners to understand, communicate, and compare optimisation scenarios effectively, without requiring technical or mathematical expertise.

5.5 Experimental design

This section outlines the experimental design used to analyse the model's behaviour under a variety of scenarios. The objectives are to: (1) validate the model's performance; (2) analyse its sensitivity to key parameters and settings; and (3) extract managerial insights that can support DHL in real-world decision-making.

Before executing the experimental phases, both the model structure, and its solution outcomes were reviewed with domain experts at DHL. Their feedback confirmed that the logic, assumptions, and results of the model align with operational expectations, thereby providing validation of the model outcomes from their perspective.

Section 5.5.1 introduces the baseline test settings that serve as a reference point for most experiments. The remainder of the chapter is structured into seven targeted experimental phases, each testing a distinct aspect of the model:

- Phase 1 – Model performance benchmark (Section 5.5.2)
- Phase 2 – Impact of maximum locker constraint (Section 5.5.3)
- Phase 3 – Impact of radius strategy (Section 5.5.4)
- Phase 4 – Demand sensitivity (Section 5.5.5)
- Phase 5 – Cost sensitivity (Section 5.5.6)
- Phase 6 – Impact of robust optimisation (Section 5.5.7)
- Phase 7 – Model value potential and scalability (Section 5.5.8)

Each phase follows a consistent structure that outlines the objective of the experiment, the configuration used, and the specific parameters or settings tested.

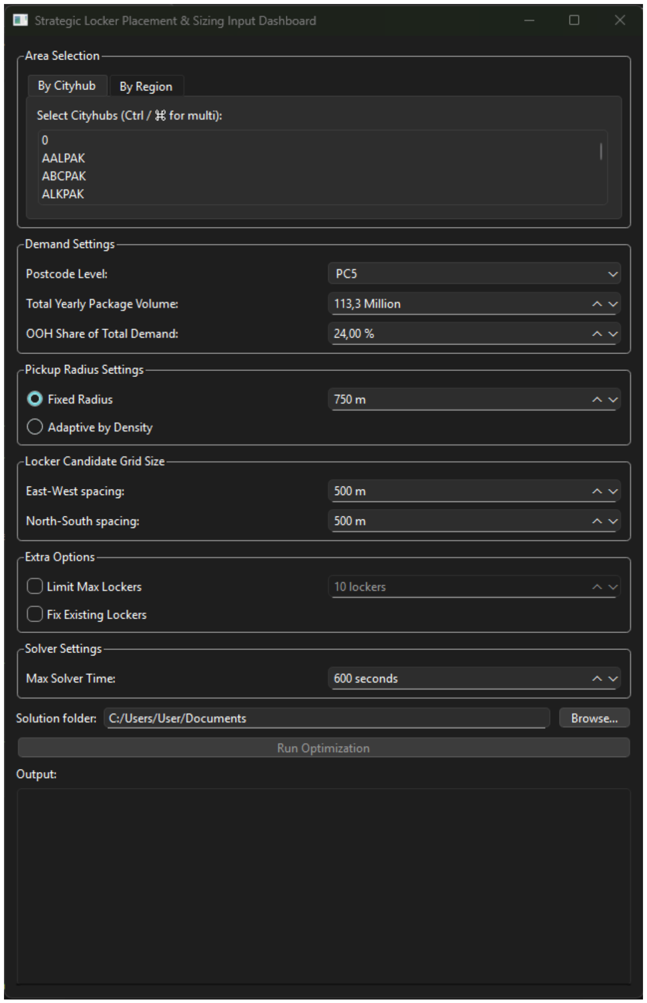

(a) Input dashboard

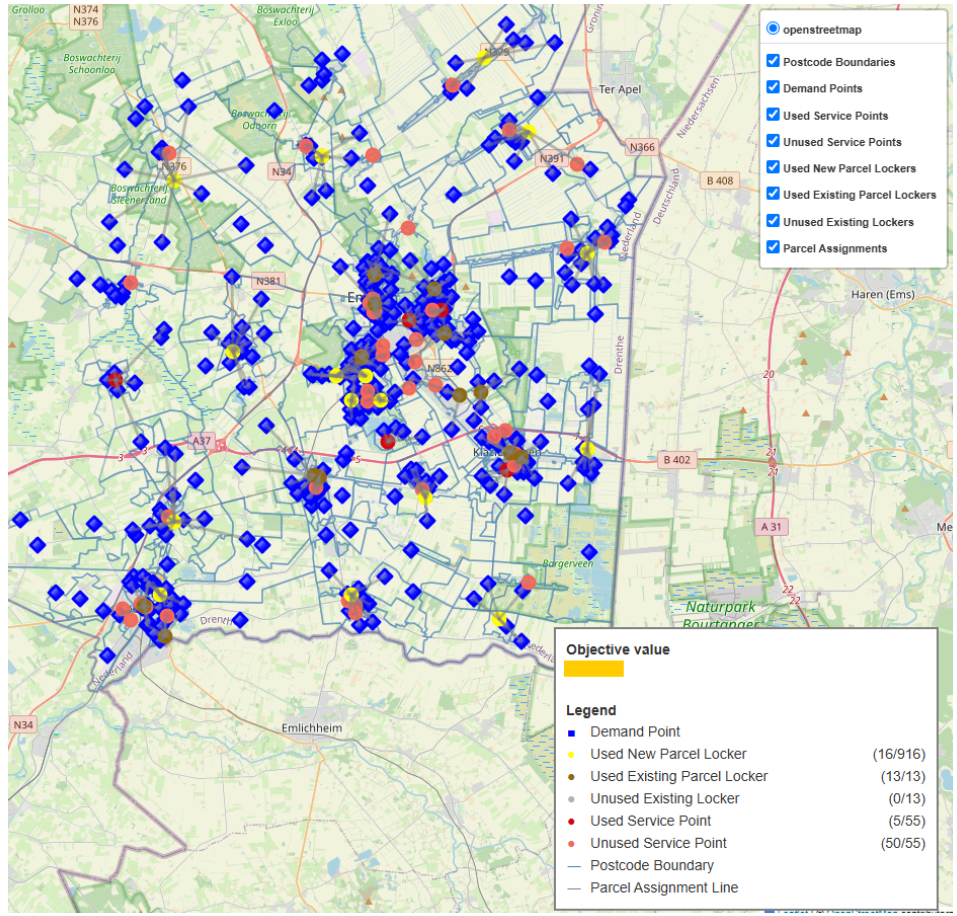

(b) Interactive solution map

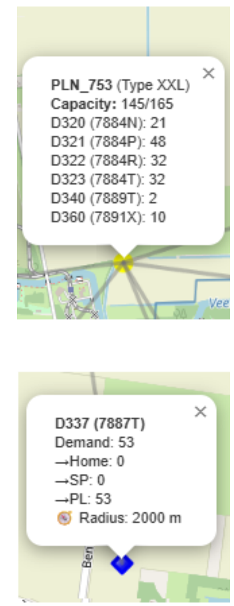

(c) Clicked element information

Figure 20: Overview of the dashboard input interface, solution map, and clicked element information.

5.5.1 Baseline test settings

Unless stated otherwise, all experiments in this study are conducted within a consistent geographical setting. The region comprising Enschede, Hengelo, and Almelo is selected as the baseline test area. As visualised in Figure 21, this area represents a varied spatial environment, combining large cities, mid-sized towns, small villages, and extensive rural zones such as agricultural land and open fields. This choice allows the model to be tested in a heterogeneous setting that captures a broad spectrum of population densities, spatial structures, and accessibility patterns. As such, it can be considered a representative or 'average' environment to evaluate LM delivery configurations within the Netherlands. At the same time, the variation within the region enables direct comparisons across different geographic characteristics, providing valuable insights into how the model behaves under contrasting spatial conditions.

Moreover, the region is structured in such a way that it naturally enables scalable experimentation. It can be incrementally expanded from a single area setting (Enschede) to a dual or triple region setting (Enschede–Hengelo–Almelo). This structure facilitates systematic testing of computational performance and solution behaviour across increasingly large and complex problem instances.

Table 13 presents the baseline configuration. These default parameter values ensure consistency and comparability throughout the experiments. Unless explicitly stated otherwise, individual experiments are conducted under these settings; any deviations are noted within the relevant sections. If only a subset of parameters is changed, we refer to this as the 'new baseline', listing only the modified values from this baseline configuration.

Table 13: Baseline configuration

Parameter	Baseline value
Demand Modelling Level	PC5
Grid Spacing (km)	0.01
Adaptive Radius	Enabled
Max Lockers to Place	Disabled
Fixed Existing Lockers	Disabled
Robust Optimisation	Disabled

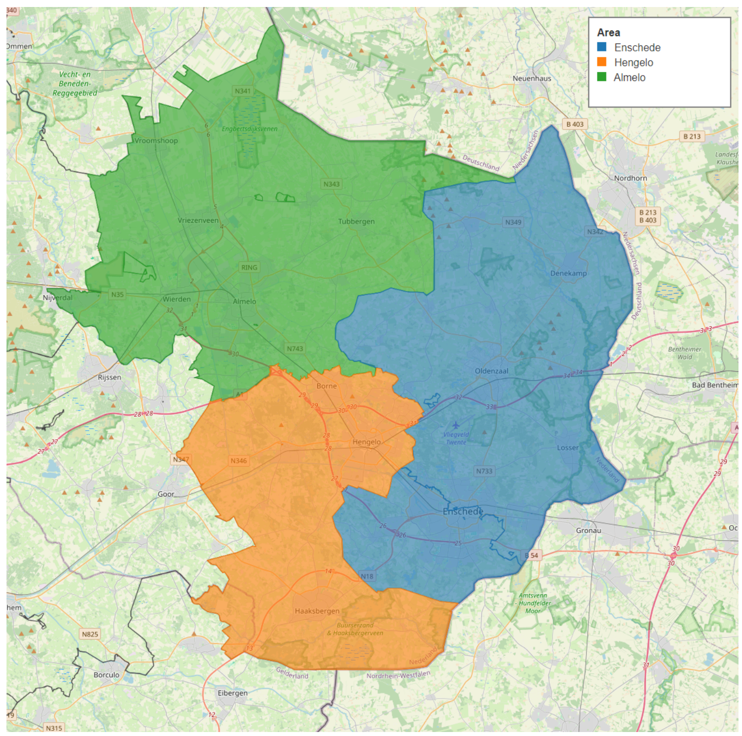

Figure 21: Baseline test area used throughout the experiments

5.5.2 Phase 1: Model performance benchmark

The first experiment benchmarks the computational performance of the model under different configuration settings. The aim is to assess how key modelling parameters affect runtime, convergence, and scalability. Before applying the model to complex real-world scenarios, it is important to understand how sensitive its computational performance is to changes in input settings. Table 14 summarises the baseline configuration used in this experiment, as well as the variations tested for each parameter, where only one parameter is varied at a time while all others are kept at their baseline value. However, later experimental phases do evaluate combinations of these features to assess their joint behavioural and performance effects. This experiment is conducted across three regions of increasing spatial and demographic complexity: (1) Enschede, (2) Enschede-Hengelo, and (3) Enschede-Hengelo-Almelo.

Table 14: Baseline configuration and parameter variations - Model performance benchmark

Parameter	New baseline	Tested variations
Parcel Class Level	PC5	PC4, PC6
Grid Spacing (km)	0.01	0.005, 0.015
Coverage Radius (m)	1000	750, 1250, 2000
Adaptive Radius	Disabled	Enabled
Max Lockers to Place	Disabled	Enabled
Fixed Existing Lockers	Disabled	Enabled
Robust Optimisation	Disabled	Enabled

5.5.3 Phase 2: Impact of maximum locker constraint (U)

This experiment investigates how limiting the number of new PLs to be placed via the parameter U affects the structure of the network and the behaviour of the model. In practice, U reflects real-world constraints such as budgets or rollout strategies. We aim to analyse (1) how the model prioritises locker locations when only a few can be placed, (2) whether these placements follow intuitive logic and remain consistent between scenarios and runs, thereby offering an implicit validation of the model's behaviour, and (3) to assess diminishing returns by evaluating the marginal cost savings of each additional locker and identifying when further expansion becomes economically inefficient. All tests are run only for the Enschede region, to allow a clear visual comparison of placement patterns under different values of U. Table 15 presents the variations tested.

Table 15: Experimental settings – Max lockers experiment

Parameter	New Baseline	Tested Variations
Use Max Lockers	Enabled	—
Fix Existing Lockers	Enabled	—
Max Lockers Allowed (U)	—	1, 2, 5, 10, 20, 50, 75, 100

Table 16: Experimental settings – Radius experiment

Parameter	Tested variations
Pickup Radius (m)	750, 1000, 1250, 2000
Adaptive Radius	Enabled, Disabled

5.5.4 Phase 3: Impact of radius strategy

This experiment investigates how different pickup radius strategies affect the network configuration and solver performance. The pickup radius represents how far customers are willing to travel, and may therefore strongly influence the spatial layout of the resulting network. In addition to testing fixed values, it evaluates the adaptive radius approach introduced in Chapter 4 and is further explained in Section 5.2.7.

The analysis focusses on three key objectives: (1) evaluate how both the fixed and adaptive radius strategies shape the spatial distribution of lockers and assess whether these outcomes are intuitive and explainable, thereby offering implicit validation of the model's behaviour; (2) compare the fixed and adaptive strategies in terms of their ability to simulate real-world realism and assess whether the adaptive strategy, for which it was designed, can offer practical value for DHL's real-world planning context; (3) analyse solver performance across strategies, evaluating how different radius configurations affect computational efficiency. To support a clear visual assessment of the OOH network in different settings, all tests are conducted on the Enschede region only. The complete set of radius strategies tested is summarised in Table 16.

5.5.5 Phase 4: Demand sensitivity analysis

This experiment investigates how the model responds to systematic changes in total demand levels. As introduced in Chapter 4, the model operates using average daily demand estimates per location. This provides a stable and practical basis for long-term planning and already absorbs short-term fluctuations with DHL's dynamic capacity control system. However, it also abstracts away from longer-term demand shifts, such as structural growth or longer seasonal peaks. Given DHL's long-term strategic ambition to significantly scale up its locker network by 2030, it is essential to understand whether the model produces consistent, efficient, and scalable outcomes as demand increases.

The experiment serves three main purposes. (1) Evaluate how demand growth affects the core outcomes of the model, including the number of lockers, sizing patterns, and delivery mode shares. (2) Examine the spatial consistency of the model decisions, whether the same locations remain optimal as demand increases, or whether structural changes emerge in the network layout. These insights are crucial for DHL's long-term infrastructure planning. (3) Provide insight into the scalability of the proposed infrastructure under current demand levels: can growth be absorbed through resizing existing lockers, or will additional placements be required? This helps DHL anticipate when and where proactive expansion may be needed and identify how best to reinforce its network to remain robust during long-term growth or long seasonal peaks. The tested demand multipliers are shown in Table 17. Experiments are performed separately on Enschede, Hengelo, and Almelo to test the consistency of the model in regions with varying urban structure and demand characteristics, and to evaluate the generalisability of the outcomes.

Table 17: Experimental settings – Demand experiment

Parameter	**Tested Variations**
Demand Factor Multiplier	0.25, 0.50, 0.75, 1.00, 1.25, 1.50, 1.75, 2.00, 2.50, 3.00

5.5.6 Phase 5: Cost sensitivity analysis

This experiment explores how the model responds to changes in delivery cost parameters, which directly shape the mode split and overall design of the OOH network. Each delivery mode, lockers, SPs, and home delivery, has a distinct cost profile based on current operational data. However, with DHL planning a significant expansion of its locker network, these cost structures are likely to change. Testing these cost changes not only reveals how sensitive the model is to key cost parameters but also provides insights that may be essential for making robust and future-proof strategic decisions.

This phase serves three goals. (1) Assess how sensitive the model is to changes in cost parameters and then examine how delivery mode shares shift in response, highlighting tipping points where lockers overtake other modes and revealing the competitive range in which delivery mode remains feasible. (2) Examine whether the model consistently adapts to changing input values in a logical and explainable way, offering implicit validation of the behaviour of the model. (3) Provide managerial insight into which cost changes have the largest effect on total system cost, helping DHL prioritise focus areas in cost optimisation or risk mitigation while also exploring the combined effect of plausible simultaneous cost changes and their potential impact.

The model is tested with cost factors ranging from 0.1 to 3.0 (step 0.1), simulating both under- and overpricing scenarios, as shown in Table 18. All experiments are carried out in the combined Enschede–Hengelo region. This setting offers a broader spatial base than a single city, increasing the robustness of results, while remaining computationally manageable for a large set of variations. In the PL fixed cost analysis, only the locker-size-dependent component is varied, representing infrastructure and installation costs. Other fixed costs, such as administrative costs, which also apply to SPs, remain unchanged to ensure a fair comparison between modes and only test the effect of locker-specific investment and operational costs.

Table 18: Experimental settings – Cost sensitivity analysis

Parameter	**Tested Variations**
Variable PL Cost	0.1–3.0 (step 0.1)
Fixed PL Cost	0.1–3.0 (step 0.1)
Variable SP Cost	0.1–3.0 (step 0.1)
Variable home Cost	0.1–3.0 (step 0.1)
PL Fixed ↓, Home ↑	PL fixed: 0.5–1.0, Home: 1.0–1.5 (step 0.1)

Table 19: Experimental settings – Robust optimisation

Δ \ Γ	0	0.2	0.4	0.6	0.8	1
0	✓	✓	✓	✓	✓	✓
0.2	✓	✓	✓	✓	✓	✓
0.4	✓	✓	✓	✓	✓	✓
0.6	✓	✓	✓	✓	✓	✓
0.8	✓	✓	✓	✓	✓	✓
1	✓	✓	✓	✓	✓	✓

5.5.7 Phase 6: Impact of robust optimisation

This phase evaluates how the model performs when explicitly designed to withstand multiple local worst-case demand scenarios. As explained in Chapter 4, there remains a residual operational risk that multiple nearby demand points may simultaneously experience sharp short-term volume peaks, due to, for example, disruptions or local events, that exceed the flexibility of DHL's dynamic capacity control system, which is otherwise capable of eliminating most day-to-day variation. In such cases, it may no longer be possible to redistribute parcels effectively across nearby lockers due to multiple simultaneous spikes, making robustness a possible design consideration. This robustness experiment differs fundamentally from the demand growth scenarios in Section 5.5.5, which simulate longer-term structural increases in volume at all demand points, as explained in Section 4.1.

This phase serves three goals: (1) Assess how solver efficiency responds to different combinations of robustness parameters, indicating whether robust optimisation remains computationally manageable. (2) Evaluate how total system cost changes as robustness increases, highlighting the trade-off between cost-efficiency and robustness. (3) Analyse how robustness requirements influence locker placement and sizing decisions, informing DHL's design choices under uncertainty. All experiments were conducted in the Almelo region to ensure a consistent comparison and keep computation times within a manageable range. The complete set of combinations of the parameters tested is shown in Table 19.

5.5.8 Phase 7: Model value potential and scalability

While other experimental phases also already provide managerial insights, this experiment places an even stronger emphasis on the practical value and usability of the model for DHL by applying it in realistic operational DHL settings, assessing its runtime performance, potential in cost reductions, and testing its scalability across increasingly large DHL areas. As DHL considers deploying the model in nationwide locker planning, it is critical to understand how well the model performs when scaled up and how its output can support real-world decisions.

The experiment serves three main purposes: (1) Test the operational practicality of the model by running it under DHL's preferred planning settings across all RegioHub areas. The goal is to confirm that runtimes remain acceptable and solutions are interpretable for the locker placement teams. (2) Explore the scalability limits of the model by solving the model on a full national scale of the Netherlands. (3) Quantify the value potential of the model by comparing the current locker network with an optimised model solution across Enschede, Hengelo, and Almelo. This aims to assess whether and how the model can reduce operational costs.

The tested configurations are summarised in Table 20. To ensure that this experiment produces relevant and usable outcomes for real-world application, a set of preferred model settings was defined together with DHL. Unless stated otherwise, these preferred settings are used throughout these experiments to generate results that best reflect how DHL intends to apply the model in practice.

Table 20: Experimental settings – Model value potential and scalability

Subexperiment	Scope	Tested Variations
Value Potential	Ens., Heng., Alm.	Cur. situation vs Model Sol.
Oper. Practicality	All RegioHub areas	`PC5`; $U = 1, 5, 10$
Scalability Limits	Entire Netherlands	`PC4`/ `PC5`; $U = 1, 5, 10, 50, 100$

Table 21: Preferred model settings DHL

Parameter	Pref. Setting
Demand Modelling Level	`PC5`
Grid Spacing (km)	0.01
Adaptive Radius	Enabled
Max Lockers to Place	Enabled
Fixed Existing Lockers	Enabled

5.6 Summary & conclusion

This chapter answered the research question: *"How should the developed model be implemented for real-world application, and which experimental setup should be used to evaluate it?"* It described how the optimisation model was operationalised using real-world data and implemented in a modular Python environment, resulting in a functional decision-support tool that is specifically tailored to DHL's operational context. It described how key data sources were preprocessed, how model parameters were defined, and how an interactive GUI was developed to support strategic decision-making. In doing so, the chapter also contributes to the transparency and reproducibility of the research.

Building on this implementation, seven targeted experimental phases were designed: (1) to validate the model's performance; (2) to analyse its sensitivity to key parameters and assumptions; and (3) to extract managerial insights that can inform DHL's strategic OOH network decisions.

With the model setup and experimental design in place, the foundation is set to evaluate how the model performs under a range of scenarios. To ensure that the model is not only theoretically sound, but also practically relevant, its outcomes were reviewed with DHL experts. Their feedback confirmed that the model's logic and results align with operational expectations, validating the model's outcomes from their perspective. The next chapter presents the results, combining both model validation and managerial insights for DHL.

6 Experimental Results

This chapter presents the results of the experimental phases introduced in Chapter 5, thereby answering the research question: *"What are the outcomes of the experimental design?"* Each section corresponds to one of the seven experimental phases and evaluates a specific aspect of the model's behaviour. The results are analysed according to the objectives defined in the experimental design, with a focus on model validation, sensitivity, robustness, and the extraction of managerial insights relevant to the strategic decision-making of DHL.

6.1 Phase 1: Model performance benchmark

The results presented below follow the experimental goals and setup described in Section 5.5.2. Section 6.1.1 evaluates the model's sensitivity to spatial configuration parameters, while Section 6.1.2 examines the impact of enabling individual model features.

6.1.1 Effect of spatial configuration

Table 22 presents the solver runtimes for variations in configurations or parameters: parcel classification level (PC level), grid spacing, and coverage radius. Each experiment was conducted under a fixed time limit of 1200 seconds (20 minutes), and results are reported per region. For incomplete runs, the solver gap at termination is shown.

Table 22: Solver runtimes (s) for PC level, grid spacing, and coverage radius variations per region

Region	**PC Level**			**Grid Spacing (km)**			**Coverage Radius (m)**			
	`PC4`	`PC5`	`PC6`	0.015	0.01	0.005	750	1000	1200	2000
Enschede	0.4	6.2	56.2	2.2	6.2	346.8	2.1	6.2	29.2	317.6
Enschede–Hengelo	0.6	30.5	326.7	2.6	30.5	gap(0.08%)	2.9	30.5	327.7	gap(0.5%)
Enschede–Hengelo–Almelo	1.4	36.0	624.9	5.3	36.0	gap(0.14%)	4.8	36.0	568.9	gap(0.9%)

The results show that solver performance is sensitive to spatial granularity. A higher level of demand modelling (e.g. `pc6`) increases the number of demand points and decision variables, substantially increasing the runtime. In contrast, a coarser level (for example, `pc4`) reduces complexity and speeds up the computation. The impact of grid spacing is even greater: reducing spacing from 0.01 km to 0.005 km (in both x- and y-directions) quadruples the number of potential locker sites, but causes exponential growth in solver time. For example, in the Enschede region, the runtime increases from 2.2 s to 6.2 s (+182%), and from 6.2 s up to 346.8 s (+5493%) at the smallest spacing. In the larger regions, the solver reaches the 20-minute time limit, yet maintains optimality gaps below 0.14%, indicating that the model remains computationally effective and delivers practically usable solutions even under more demanding configurations. The coverage radius shows a similar, also non-linear, pattern. Smaller radii lead to fewer assignable facility, demand pairs $(N(i), A(j))$, resulting in faster runtimes. Larger radii expand the decision space, introducing more options and thus increases complexity, particularly increasing in larger areas.

6.1.2 Effect of model features

Table 23 shows how the individual model features influence the solver runtimes. Each row represents a different region, starting from the baseline configuration (i.e., all features off), followed by the runtime change when enabling one feature at a time. All runs were capped at 300 seconds (5 minutes), with 8 out of 9 configurations solved within 90 seconds (1,5 minutes). For the one run that did not finish in time, the solver gap is reported.

Table 23: Solver runtimes (s) for baseline and impact of enabling individual model features per region

Region	**Baseline**	Adaptive Radius	Max Lockers	Fixed Existing
Enschede	6.2	+0.4	+12.3	-3.9
Enschede–Hengelo	30.5	+69.2	-15.1	-27.1
Enschede–Hengelo–Almelo	36.0	gap(0.06%)	+3.3	-30.4

The results show that the effect of enabling specific model features can vary significantly by region size. The adaptive radius option introduces the strongest performance impact. Although negligible in small regions (e.g., +0.4 s in

Enschede), it leads to a significant increase in larger areas, reaching the time limit in the entire Enschede–Hengelo–Almelo region. This is expected as the adaptive logic increases the radius in rural areas, expanding the assignment sets $(N(i), A(j))$. While the baseline model applies a fixed 1000 m radius throughout, the adaptive version allows this to grow in low-density areas, resulting in a larger solution space and more decisions for the solver to evaluate.

Enabling the constraint that limits the number of PLs shows a mixed impact on the solver time. In the smallest region (Enschede), runtime increases slightly (+12.3 s). This is likely because the constraint adds extra logic without meaningfully reducing the solution space, since the number of potential locker sites is already limited in such a small area. In contrast, in the medium-sized Enschede–Hengelo region, runtime decreases noticeably (-15.1 s), and in the largest region, it remains almost equal (+3.3 s). This suggests that in larger settings, the constraint helps the solver by reducing the number of feasible locker combinations. Although it introduces additional complexity due to the extra constraint, this can be offset by the smaller solution space, potentially allowing the solver to reach a solution faster, or it may not make a significant difference.

Finally, enabling the fixed existing lockers feature leads to a significant improvement in solver performance across all regions. In the smallest region (Enschede), runtime drops from 6.2 s to 2.3 s, a reduction of 62%. In the largest region, the improvement is even more substantial: from 36.0 s to 5.6 s, a speed-up of 84%. This confirms that fixing part of the locker infrastructure reduces the number of decisions the solver must evaluate. This time-saving effect becomes stronger as the region size, and thus the number of candidate locations, increases. This shows that leveraging existing infrastructure not only reflects realistic planning scenarios, but also improves computational efficiency.

6.2 Phase 2: Impact of maximum locker constraint (U)

The following analysis is based on the experimental design and objectives outlined in Section 5.5.3, examining the model's response under different values of U. All experiments were solved to optimality within 10 seconds.

Figures 23, 25a and 25b present some of the effects of varying the maximum number of new lockers to be placed in a zoomed-in snippet of the Enschede area, as constrained by parameter U. Looking at the results of the experiment, the first locker is placed in Enschede, close to the city centre, in an area with high demand but limited coverage due to the relatively small locker radius within high-density areas. This aligns with expectations, as the model prioritises areas with high demand that are currently underserved. This pattern is even more clearly illustrated in Figure 24a, which shows demand allocation lines and highlights how many customers in the city centre would otherwise rely on home delivery in the absence of a nearby locker.

(a) $U = 1$ (b) $U = 2$ (c) $U = 5$

Figure 23: Model solution for $U = 1$, 2, and 5.

When two lockers are allowed ($U = 2$), the model chooses to serve an uncovered area in Oldenzaal, where a new locker not only fills a local gap in coverage but also absorbs demand from a nearby SP. This is particularly efficient, as consolidating demand into a single locker is generally cheaper than relying on a combination of home delivery and a potential less efficient SP. This effect is clearly visible in the allocation visuals (Figure 24b and 24c), where demand that was previously split between home delivery and a SP is now fully captured by the new locker.

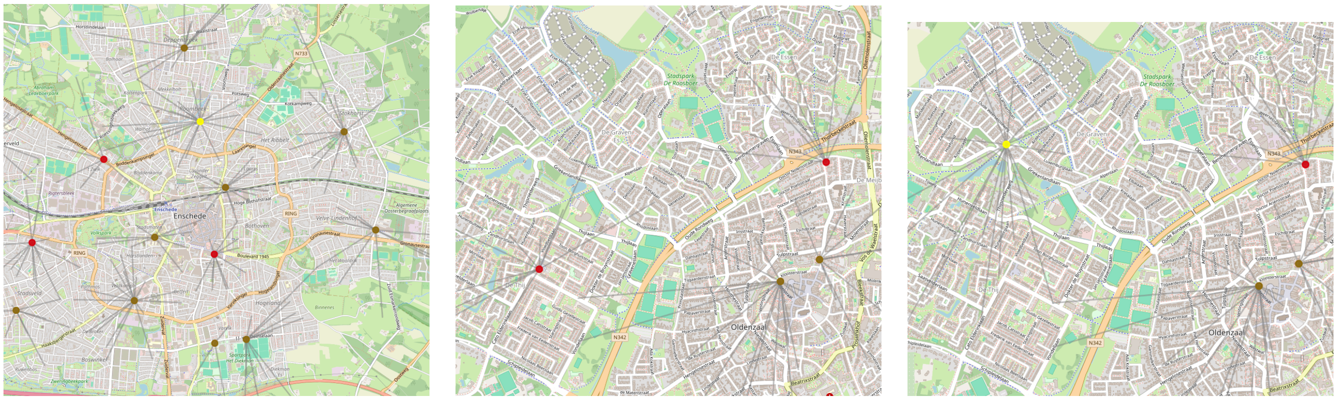

(a) Enschede, U=1 (b) Oldenzaal, U=1 (c) Oldenzaal, U=2

Figure 24: Geographical distribution of lockers and service points with demand allocation lines in Enschede and Oldenzaal

As we move to $U = 5$, we see that these patterns become more noticeable. An additional locker is placed in the south of Enschede, another underserved area with significant demand close to the city centre that was previously served with home delivery. Furthermore, the model increasingly takes over demand of SPs, especially those nearing their capacity limits, due to their relatively high CpP. A clear example of this is in Losser: at $U = 5$, this demand is served by a single SP operating at full capacity (150 parcels), while at $U = 10$, this is demand is taken by a large locker (the change in red to yellow dot in Losser) that can handle the same volume more cost-effectively.

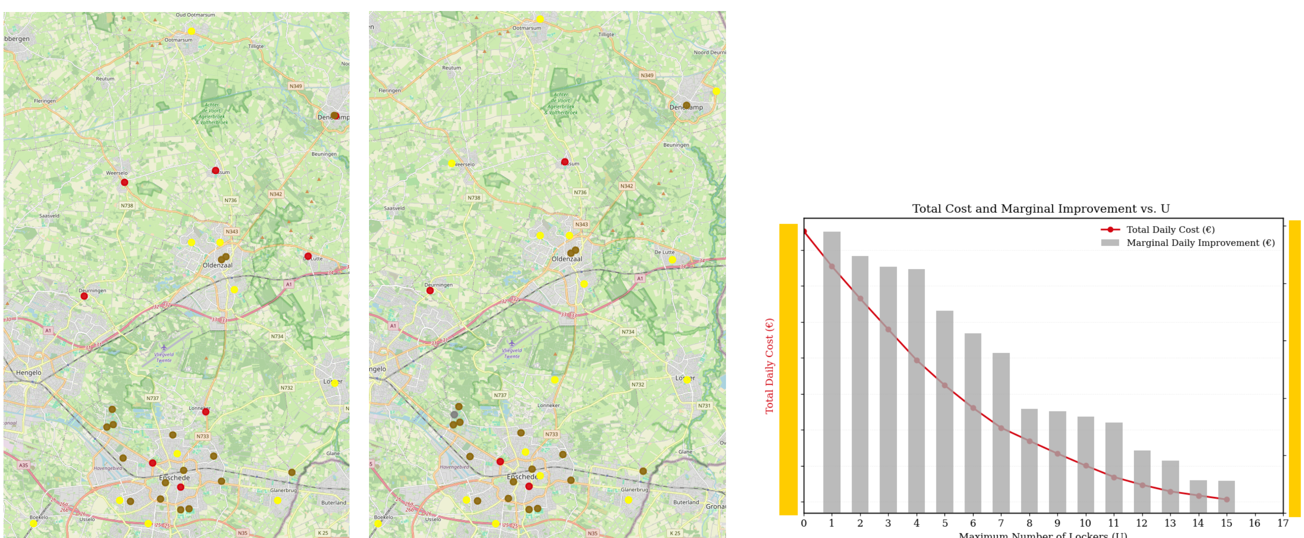

(a) $U = 10$ (b) $U = 20$ (c) Diminishing Returns

Figure 25: Model solutions for $U = 10$, $U = 20$, and visualisation of diminishing returns.

As U increases, additional locations begin to emerge, not just in urban areas, but also in more rural areas, such as the region between Enschede and Oldenzaal and in Boekelo. In these cases, demand is more spread out, but people are generally more willing to travel slightly further, and lockers prove to be cost-efficient alternatives to home delivery.

This progression follows a clear and intuitive pattern. The model initially prioritises locations with high unmet demand (currently with home delivery), where the placement of a locker generates the greatest cost-saving potential. Hereafter, it targets SPs with high parcel volumes, as these are relatively expensive due to their high CpP, making them cost-inefficient compared to lockers that benefit from economies of scale. Beyond this point, lockers also begin to appear in more rural or peripheral areas with lower demand densities. In these cases, smaller lockers can still be more cost-effective than maintaining home delivery. SPs with only limited volumes tend to remain in the solution, as their relatively high CpP still outweigh the cost of installing and operating an entirely new locker. At some point, the model no longer places additional lockers because the investment is no longer justified. For the remaining, more dispersed demand, existing SPs or home delivery offer a cheaper alternative.

As visually confirmed, the model consistently selects the location with the highest overall cost-saving potential at each incremental value of U. As a result, locker placements chosen at $U = 5$ reappear in the solutions for higher values such as $U = 10$. Figure 25c illustrates the diminishing returns of additional PLs. The total cost drops sharply after the first 4 lockers, with marginal benefit per locker decreasing from approximately €[redacted] to €[redacted], €[redacted], and €[redacted] per day. From approximately the eighth locker onwards, the savings fall below €[redacted] per locker and continue to decline. Around $U = 15$, the curve begins to flatten, indicating that placing additional lockers beyond this point does not generate sufficient cost savings to justify the investment. Consequently, no new lockers are placed after $U = 15$. This extra analysis could support DHL in making strategic placement decisions by quantifying the marginal savings of each additional locker.

6.3 Phase 3: Impact of radius strategy

Table 24: Solver runtimes and delivery mode shares for different radius strategies in Enschede

Strategy	OOH Total	% Locker	% SP	% Home	S	M	L	XL	XXL	SP Total	Time (s)
750	33	71.6	5.7	22.7	22	8	3	–	–	6	1.9
1000	28	82.9	3.1	13.9	11	9	5	3	–	4	6.2
1250	21	86.2	1.8	12.0	5	3	4	7	2	2	27.5
2000	18	94.7	0.9	4.4	2	3	3	2	8	1	322.5
Adaptive Radius	29	88.9	3.5	7.5	11	9	6	3	–	4	7.6

The following analysis is based on the experimental design and objectives outlined in Section 5.5.4. As shown in Table 24, the radius configuration has a clear and substantial impact on both the solver performance and the resulting OOH network structure. Solver runtimes to optimality increase non-linearly with a larger radius: from just 1.9 seconds at 750 m to over 322 seconds at 2000 m. The increase is expected because a larger pickup radius expands the assignment sets, increasing the number of decision variables the solver must evaluate.

At the same time, a larger radius allows for greater demand coverage through lockers, increasing from 71.6% at 750 m to 94.7% at 2000 m. This comes with a notable merging effect: fewer lockers are needed and their size increases. While this is attractive in rural zones, it comes at the cost of granularity in dense urban settings such as the city centre of Enschede, where proximity is often more important than reach. This is also visible in the 2000 m solution in Figure 26a, where very few lockers are placed in the city centre. A smaller radius, such as 750 m or 1000 m, leads to a dense locker placement in the city centre, often clustering multiple small lockers in close proximity. However, these configurations struggle to reach the more rural zones.

The adaptive radius option offers a strong middle ground. It solves in just 7.6 seconds, while still achieving a locker delivery share of almost 89%. The adaptive approach automatically tailors the radius per location. It keeps the coverage tight in the city (where demand is dense) and stretches further out where needed, without requiring any manual adjustments or separate model runs. Looking at Figure 26, visual inspection confirms this, with the grey dots representing the OOH points placed or used in the solution. The adaptive layout places lockers densely in urban cores, achieving a balance between the 750 m strategy, which places a very high number of lockers in Enschede and surrounding areas, and the 1000 m strategy, which is somewhat less dense in the very centre of the city. At the same time, it still manages to reach more rural areas that remain uncovered in the 750 m and/or 1000 m configurations, such as De Lutte, Deurningen, Weerselo, Rossum and Vliegveld Twente, as with the 2000 m radius.

Overall, the results show intuitive and consistent placement patterns, reinforcing the logic embedded in the model. They also highlight that no single fixed radius performs optimally across a heterogeneous region. Small radii offer (too) high precision in urban areas but lack sufficient rural reach, while large radii improve overall coverage at the cost of increased computation time and reduced locker granularity in dense zones such as city centres. The adaptive radius strategy avoids this trade-off by flexibly adjusting to the local context, offering an alternative that is both realistic and computationally efficient. This demonstrates its potential as a valuable extension for DHL, capable of running across varied regions without additional configuration effort.

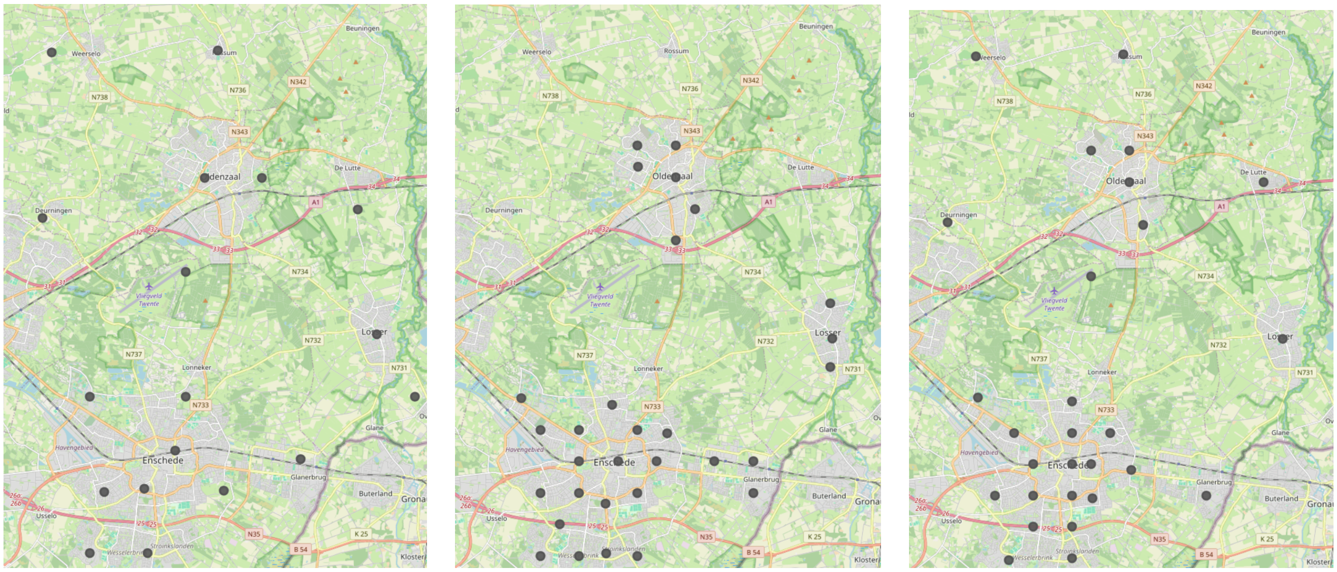

(a) Radius 2000 m (b) Radius 750 m (c) Adaptive Radius

Figure 26: Model solutions for 750 m, 2000 m, and adaptive radius strategies

6.4 Phase 4: Demand sensitivity analysis

The results presented below correspond to the experimental setup and goals described in Section 5.5.5. Figures 27 and 28 together illustrate how the model responds to increasing demand across the regions. As the total volume increases, the number of lockers placed increases in all regions, but growth is not linear. The dotted line in Figure 27 shows how the number of lockers would increase if the scaling were strictly proportional. In contrast, the empirical curves rise more gradually, particularly at higher demand levels, indicating that the model increasingly absorbs volume by placing large lockers instead of deploying additional locations.

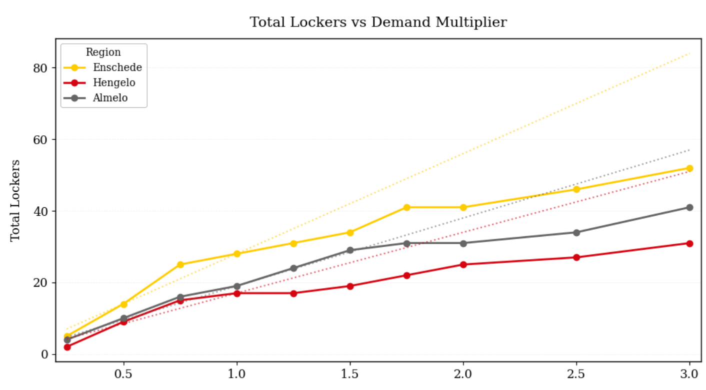

Figure 27: Total number of lockers placed versus demand multiplier, with dotted linear baseline for comparison.

This observation is further supported by the locker-type distributions in Figure 28. While demand increases by 50% between multipliers 1.0 and 1.5, the number of lockers placed increases by only about 21%. However, this may give a somewhat misleading picture of what is truly required from a capacity perspective. As shown in Section 6.2, the model starts replacing SPs with nearby lockers as volumes increase, not because these SPs lack capacity, but because lockers become more cost-efficient at scale. SPs can handle around 150 parcels per day, even though the model often assigns them far fewer (e.g., 20 per day) due to their narrow optimal operating window. In other words, there often remains considerable unused buffer capacity available at SPs. This makes SPs a valuable asset during extended seasonal peaks or gradual volume growth, as discussed in Chapter 2. Leveraging their existing flexibility allows DHL to absorb additional demand without immediately investing in new lockers. While lockers may offer long-term cost advantages, SPs already outperform home delivery and can delay infrastructure expansion until higher, more stable demand levels are reached. For long-term planning, it remains relevant to explore where lockers might eventually replace SPs. However, when evaluating the network's robustness under temporary or incremental demand increases, the more important question is: how much new capacity is truly needed? In other words, how many new OOH locations are added because the system would otherwise hit its operational limits?

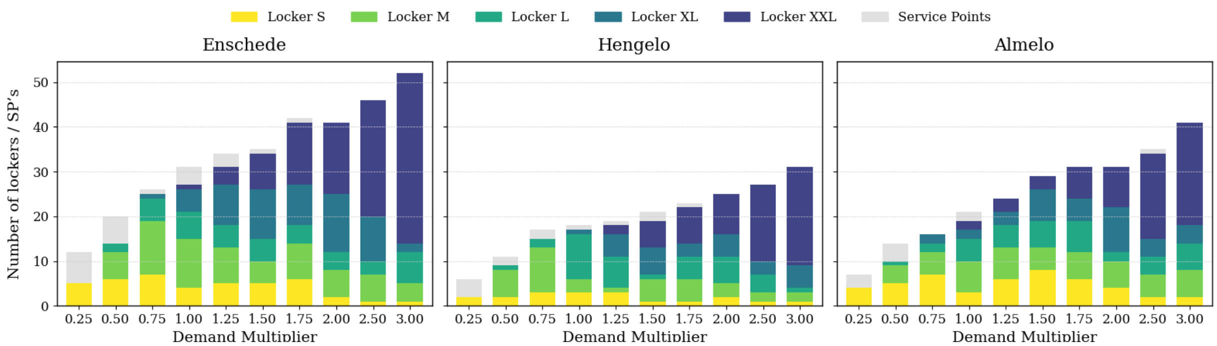

Figure 28: Distribution of facility types across demand multipliers per region.

To answer this, we distinguish between lockers added purely for cost-efficiency (i.e., replacing a SP) and those added to expand total capacity. We do so by including SPs in the overview and comparing total OOH locations across demand levels in Figure 28, focusing only on truly new placements, net increases in OOH points, rather than substitutions. This distinction is clearly visible in the Enschede region. While demand rises by 50%, the total number of lockers increases by just 9%, three additional units, compared to a 21% increase when replacements are also counted. This shows that over 82% of this extra volume is absorbed by upgrading locker sizes, instead of increasing the number of locker locations.

Figure 30 visualises the optimal model solution for locker placements at demand factors 1.0 and 1.5. An interesting pattern emerges: despite a substantial increase in total parcel volume, the spatial distribution of locker locations remains largely unchanged. The majority of lockers placed under baseline demand (demand factor = 1.0) continue to appear in the optimal solution at higher demand levels, often at the exact same coordinates. Rather than reallocating lockers or placing new, larger lockers elsewhere, the model primarily upgrades existing ones, replacing smaller formats with larger ones. Lighter colours (S, M, L) are substituted by darker shades (XL, XXL), reflecting a systematic upscaling of locker capacity at the same location.

The few additional lockers (at new locations) that appear under increased demand are interestingly mostly small in size. They are placed (i) in dense urban areas, where a short pickup radius and concentrated demand require supplementary lockers between existing sites, and (ii) in rural regions, such as north of Lonneker, where rising demand now makes it viable to install lockers. Furthermore, in some cases, these new lockers replace existing SPs as the lockers become more cost-effective once volumes increase.

These observations confirm that the model's selected locations under baseline demand already reflect structurally advantageous positions, capable of efficiently absorbing future growth without requiring major spatial reconfiguration. Despite a few small additions, the overall network layout remains remarkably stable, making the resulting infrastructure inherently more robust to future growth than a system that would require a complete spatial reconfiguration as volumes increase.

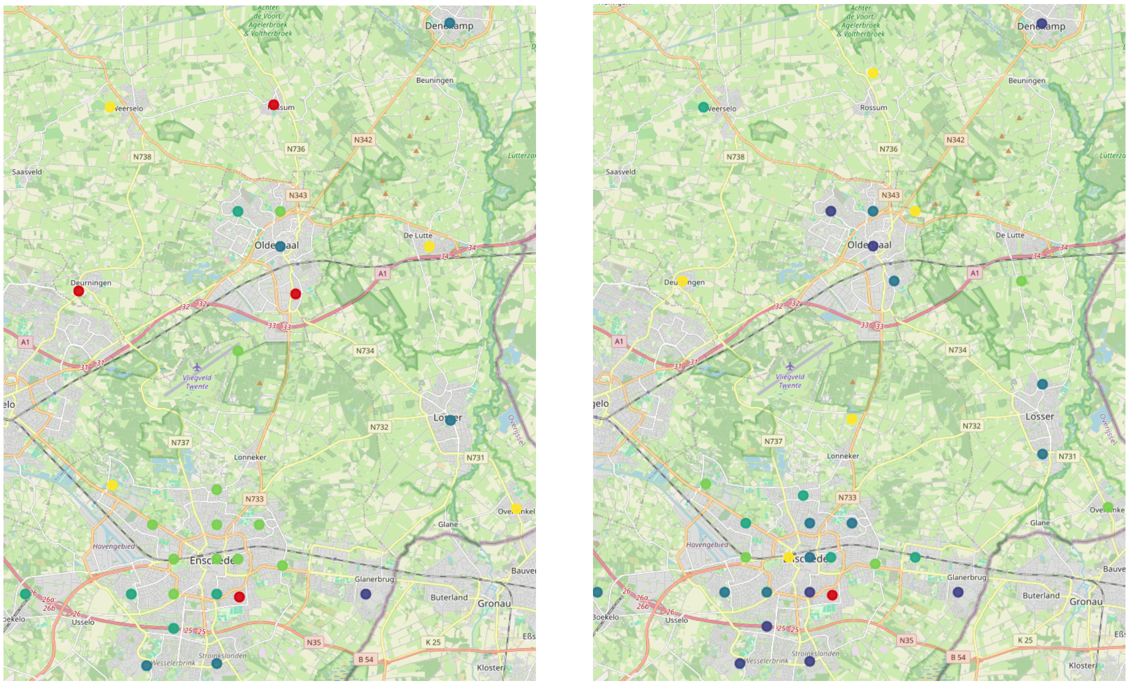

(a) Demand factor = 1 (b) Demand factor = 1.5

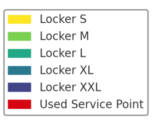

Figure 30: Model solution for demand factor 1 and 1.5.

Figure 31 presents the resulting delivery mode shares across regions as demand increases. In all three cities, the share of lockers rises sharply with demand, quickly overtaking both home delivery and SPs. In the current situation, PLs already account for around 86% of mode share in Enschede and Hengelo, and 80% in Almelo, levels that were already nearly reached at a demand multiplier of just 0.75. While locker share continues to increase, the growth slows considerably, eventually levelling off near full adoption at the highest demand levels. SPs consistently account for only a small share of total demand, both under baseline conditions and further decline as volumes rise. This highlights their limited role: they operate effectively only in a narrow volume window. Below that, home delivery is cheaper; above that, a small locker becomes more cost-efficient. The share of home delivery drops sharply between demand factors 1.0 and 2.0, falling from around 10–20% to just 3–5% across all regions. The remaining share is concentrated mainly in rural areas, where demand levels are still too low to justify the installation of lockers. In these cases, home delivery remains the most practical and cost-effective solution. As demand increases, these areas may also accumulate sufficient volume for (small) locker placement to become cost-efficient, gradually phasing out home delivery almost entirely.

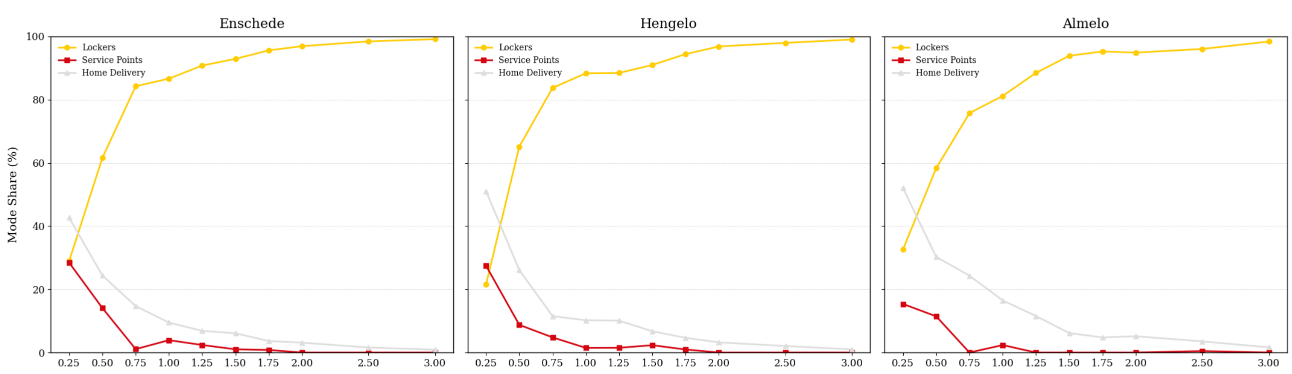

Figure 31: Mode share distribution across delivery methods (parcel lockers, service points, home) per region.

6.5 Phase 5: Cost sensitivity analysis

The experiments in this phase follow the experimental setup and objectives described in Section 5.5.6. Section 6.5.1 presents the effect of varying cost factors on delivery mode selection, while Section 6.5.2 examines the corresponding impact on the total daily cost in the objective function. Finally, Section 6.5.3 introduces a combined cost shift scenario, exploring a future-oriented trade-off between rising home delivery costs and declining locker infrastructure costs.

6.5.1 Delivery mode share

Figure 32a shows the effect of increasing the cost factor for home delivery. PLs quickly overtake home delivery between cost factors 0.4 and 0.6, becoming the dominant mode. From cost factor 1 onwards, SPs begins to also absorb the remaining demand for home delivery. This suggests that once home delivery becomes too expensive, the model prefers SPs in lower volume areas, likely due to their lower fixed costs. At cost factors between 2 and 3, home delivery is almost entirely phased out.

Figure 32b shows the effect of changing the variable costs for SPs. At baseline (factor 1.0), they account for just 2% of the total demand. A 20% cost increase reduces this to almost zero, further illustrating the narrow range and situations in which SPs remain a preferred option. Halving the cost allows SPs to briefly overtake lockers as the dominant delivery mode. This highlights their high sensitivity to cost reductions within a narrow window (factors 0.4-1.0). Below this, SP share stabilises around 70%, and further cost changes to SPs no longer affect the mode split, likely because the fixed SP infrastructure is too limited to absorb additional demand. Home delivery remains largely unchanged throughout.

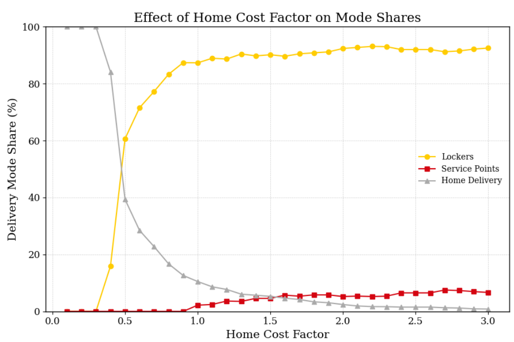

(a) Effect of home cost factor

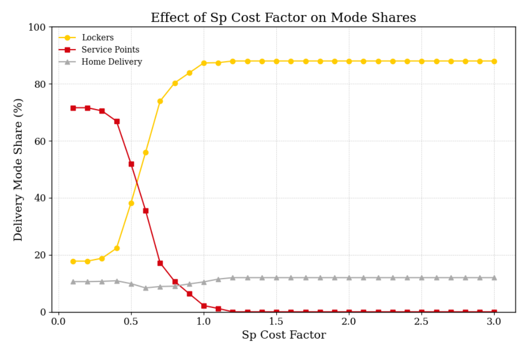

(b) Effect of service point cost factor

Figure 32: Impact of cost factors for home and service point delivery on mode shares.

Initially, lockers hold a dominant mode share of more than 85%. As shown in Figure 33a, this share remains relatively stable as variable locker costs increase, up to a factor of about 1.5, suggesting limited sensitivity within that range. Beyond this point, locker share drops sharply, falling below 80% at factor 1.7. Around factor 2.2, SPs overtake lockers as the most-used mode. Notably, home delivery remains fairly stable across most of the range, increasing only beyond factor 2.0.

Figure 33b, by contrast, shows both a more gradual increase and decrease in locker share in response to changes in fixed locker costs. As costs decrease, locker share rises gradually, reaching nearly 100% when fixed costs are extremely low. Conversely, when fixed costs exceed 2.8 times the baseline, lockers begin to lose their dominant position. Unlike in the variable cost scenario where home delivery share remains stable, this also leads to a clearer reduction in home deliveries, suggesting that lowering fixed costs enables lockers to become viable even in low-demand or rural areas. However, this effect emerges only after substantial reductions (typically 20%). Reducing variable costs alone does not show the same effect, likely because it fails to overcome the initial locker infrastructure investment threshold.

Figure 33c visualises this difference: a 20% reduction in fixed costs yields about a 3 percentage point increase in locker share, compared to just 1 point for an equivalent drop in variable costs. As prices deviate further from the baseline (factor 1), this gap widens. Locker share also declines more sharply under rising fixed costs than under rising variable costs, indicating that the model is more sensitive to fixed locker costs changes near the baseline.

While variable costs have little effect initially, crossing a critical threshold triggers a sharp drop in locker share, ultimately shifting demand back to home delivery.

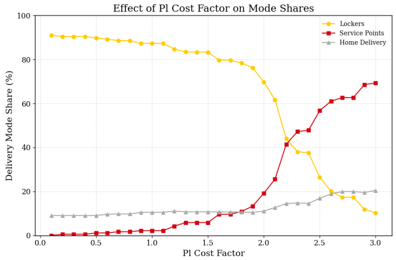

(a) Effect of parcel locker variable cost factor

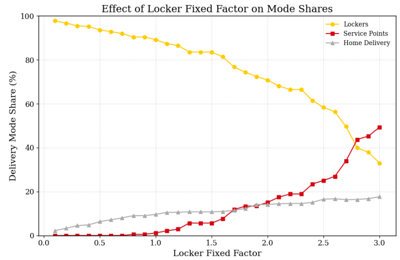

(b) Effect of locker fixed cost factor

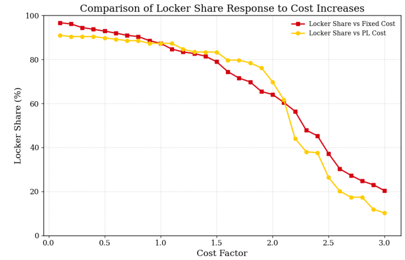

(c) Comparison of Fixed vs. Variable PL Cost

Figure 33: Impact of locker-related cost factors on the share of parcel lockers.

6.5.2 Impact on objective value

Figure 34a shows how changes in the cost factors for home delivery and SPs affect the total objective value. The impact of home delivery is clearly larger: increasing its cost immediately drives up the total cost, while reductions yield substantial savings. This aligns with previous findings, home delivery remains essential in low-density rural areas, meaning its cost is directly felt in the model's outcome. As shown in Table 25, a 20% cost reduction yields 4.6% savings in total system cost, while a 20% increase results in a 3.2% increase. In contrast, changes in SP costs have limited impact. Cost increases barely affect the objective value, consistent with their minimal role in the mode split as shown in Section 6.5.1. Only substantial cost reductions allow them to contribute meaningfully to OOH network savings.

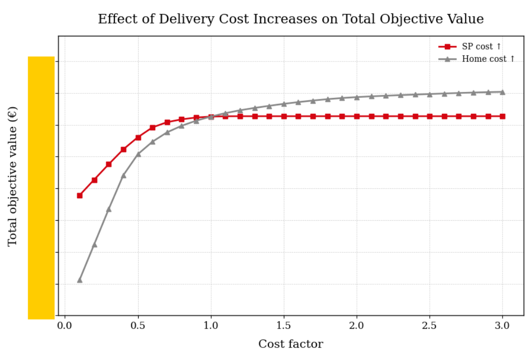

(a) Effect of service point and home delivery costs on total objective value

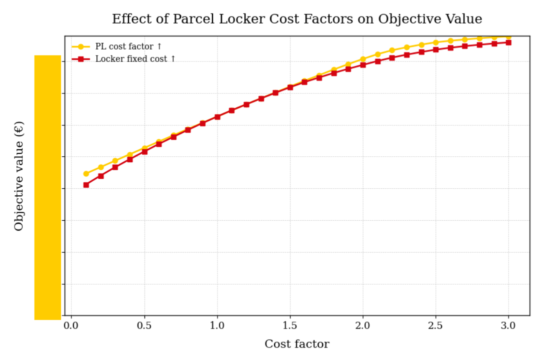

(b) Effect of fixed versus variable locker costs on objective value

Figure 34: Comparison of total objective value under different cost assumptions

Table 25: Percentage change in total cost in response to cost factor variations (relative to base = 1)

Factor change	SP cost factor	Home cost factor	PL cost factor (variable)	Locker Fixed Cost Factor
-20%	−1.4%	−4.6%	−6.2%	−6.7%
+20%	+0.2%	+3.2%	+6.2%	+6.2%

Figure 34b compares the effect of increasing PL costs, both fixed infrastructure and variable per-parcel, on the total objective value. In both cases, cost increases lead to steady but substantial rises in total cost. As Table 25 confirms, a 20% increase or decrease in either variable or fixed locker costs results in a cost change of approximately 6–7%, highlighting the objective function's sensitivity to both components. This is expected, given the dominant share of lockers in the delivery mix. However, beyond a 50% deviation from the baseline, the effects begin to diverge: large cost increases make variable costs the primary cost driver, whereas substantial cost reductions make fixed costs relatively more influential. This reinforces earlier insights; while both cost components are important, the fixed cost structure offers greater potential for long-term savings if cost reductions can be achieved.

6.5.3 Combined cost trade-off scenario

A supplementary scenario is included to assess how the model responds when locker fixed costs decrease while home delivery costs increase, reflecting a plausible future shift due to DHL's strategic goals. This setting simulates a context in which lockers become cheaper due to, e.g., scaling, while home delivery becomes more expensive due to shrinking volumes and reduced economies of scale as more parcels shift to OOH delivery.

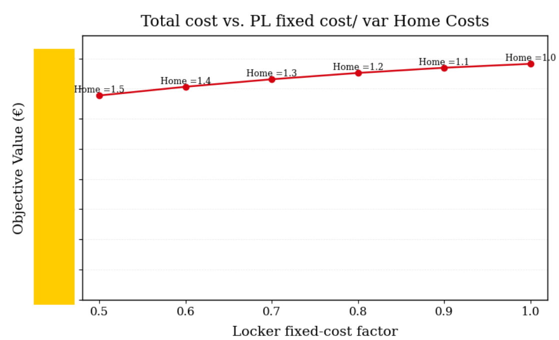

(a) Locker share vs. PL fixed cost

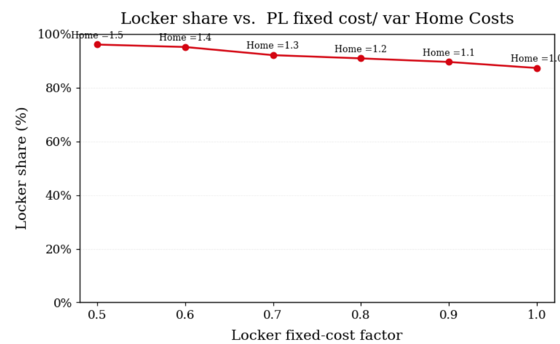

(b) Total cost vs. PL fixed cost

Figure 35: Impact of increasing home cost and decreasing PL fixed cost by the same percentage on the objective value and locker share

As shown in Figure 35, even though home delivery becomes more costly, the reduction in fixed locker costs still results in a net decrease in the total objective value. This suggests that the savings from cheaper infrastructure outweighs the rising cost of home delivery, making such a transition financially favourable for DHL. Locker share also increases under this combined change. This highlights that reducing fixed locker costs has a stronger impact on both total system cost and locker adoption than increasing home delivery costs, assuming that both cost changes occur at similar percentage magnitudes.

6.6 Phase 6: Impact of robust optimisation

This phase follows the experimental setup and objectives described in Section 5.5.7. Section 6.6.1 analyses solver performance across different levels of robustness and Section 6.6.2 examines how the total objective value, as well as the spatial allocation and sizing of lockers, change with robustness.

6.6.1 Impact on model efficiency

Figure 36 visualises the solver time required to solve the robust formulation under varying levels of demand uncertainty. Both the maximum deviation magnitude (Δ) and the uncertainty budget (Γ) significantly impact solver performance. Although low to moderate uncertainty levels generally yield fast solutions, combinations with high values for both parameters result in a sharp, often exponential, increase in computation time. Protecting against both larger maximum deviations and many simultaneous disruptions substantially increases model complexity, and thus solver time. This is particularly evident in the top-right region of the heat-map. In contrast, almost all runs with either lower Δ, lower Γ, or a combination of low and high values solve well within 300 seconds. Only a few higher combinations show a steep rise in runtime.

Despite this, the model remains solvable in nearly all configurations. As Table 26 shows, more than 90% of the runs finish within 20 minutes, and 66% complete in less than 5 minutes. Only three runs resulted in relatively extreme runtimes between 20 and 60 minutes. While these are high compared to other cases, they remain within acceptable bounds for strategic decision-making purposes.

Table 26: Model runs completed within time thresholds

Threshold (min)	Runs ≤ Threshold	Cumulative %
1	1	2.8%
5	24	66.7%
10	28	77.8%
20	33	91.7%
45	35	97.2%
60	36	100.0%

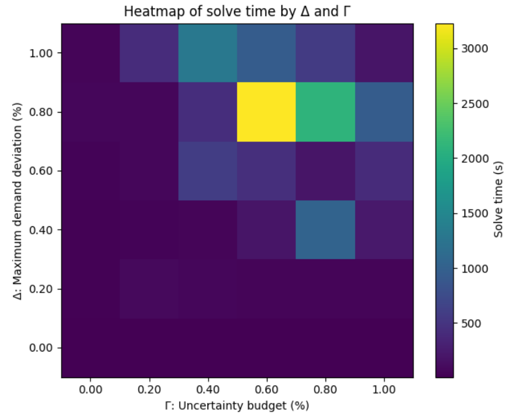

Figure 36: Solve time (s) across combinations of Δ (max demand deviation) and Γ (uncertainty budget).

6.6.2 Impact on model outcome

Objective value

The 3D surface plot in Figure 37 shows how the total objective value responds to increasing levels of robustness. As both the maximum-deviation parameter (Δ) and the uncertainty budget (Γ) increase, the model allocates additional capacity within the OOH network to safeguard against worst-case demand scenarios, which results in higher overall costs.

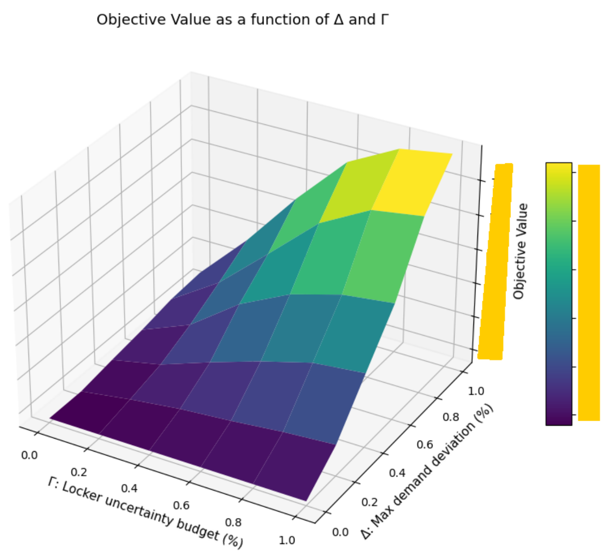

Figure 37: Total daily objective value (in €) as a function of Δ and Γ for the Almelo region.

Figure 37 illustrates that increasing robustness has a clear and escalating cost impact. While modest robustness levels lead to limited increases in the objective value, higher parameter settings cause costs to rise sharply. For instance, as shown in Table 33 in Appendix A.0.9 showing all the results of the experiment, increasing from a baseline with no robustness to $\Delta = 0.4$ and $\Gamma = 0.4$ results in an increase of just €12.29. However, increasing both parameters further to $\Delta = 0.6$ and $\Gamma = 0.6$ leads to a cost of €939.90, €33.01 above baseline and around three times the cost increase seen at $\Delta = 0.4$ and $\Gamma = 0.4$, even though the parameter increase was only half as large (0.2 instead of 0.4). This demonstrates a distinctly non-linear relationship between robustness and cost. The effects of Δ and Γ interact in a compounding way and show a convex cost structure: the marginal cost of increasing one parameter grows as the other increases, meaning that designing the network to withstand both larger maximum individual demand location deviations and broad simultaneous disruption becomes disproportionately expensive.

At first glance, the effects of Δ and Γ appear symmetric. However, a closer examination of Table 33 and the detailed colour gradients in Figure 37 reveals that increases in Δ tend to raise costs more sharply than equivalent increases in Γ. The results indicate that large, concentrated demand shocks are more costly than more distributed uncertainty, suggesting that from a cost perspective, mitigating extreme local single deviations are more expensive than guarding against multiple moderate fluctuations across the network. This also highlights the value of the dynamic capacity control project, which aims to smooth parcel flows across lockers and prevent extreme local spikes in demand.

Locker placement and sizing

Figures 38 and 39 together illustrate how robustness settings impact both the size and spatial distribution of PLs in Almelo. First, the stacked bar chart in Figure 38 shows that the total number of lockers remains almost constant, growing only slightly from 19 at $\Delta/\Gamma = 0.0$ to 22 at $\Delta/\Gamma = 1.0$. Rather than opening new sites, the model upscales locker sizes: S, M, and L lockers are gradually replaced by large XL, and XXL units as uncertainty increases. This suggests that, under local worst-case demand shocks, having higher capacity-lockers is more cost-effective than introducing extra locker sites, and protects the network from overloads. Second, the maps in Figure 39 show that the locker locations remain nearly identical. Comparing (a) (without robustness) and (c) (with relatively high robustness), the lockers are still clustered in the same Almelo neighbourhoods, often at the exact same coordinates, with only minor local location shifts in some cases. Only at relatively high robustness settings a few new sites emerge, as already reflected in Figure 38. These findings mirror the results from the demand-sensitivity experiment in Phase 4: the strategic sites selected under current (baseline) conditions remain effective under structural demand growth and longer seasonal peaks, and, as shown here, a similar pattern emerges under local worst-case demand spikes. By scaling capacity at the locations that were identified as optimal under non-robust conditions, the network absorbs uncertainty without requiring any major spatial reconfiguration, while simultaneously offering the most cost-effective response.

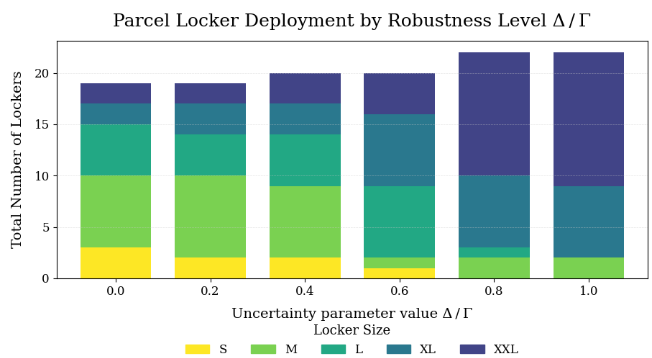

Figure 38: Parcel locker deployment by robustness level (Δ/Γ), differentiated by locker size in the Almelo region.

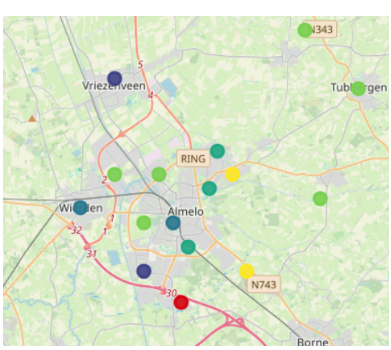

(a) $\Gamma/\Delta = 0.0$

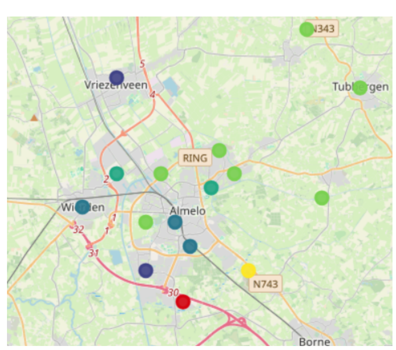

(b) $\Gamma/\Delta = 0.4$

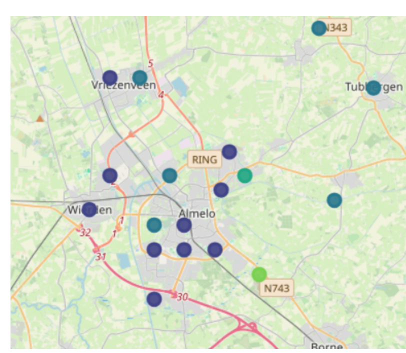

(c) $\Gamma/\Delta = 0.8$

Figure 39: Locker placement under increasing robustness levels (Γ/Δ) in a subsection of the Almelo region.

6.7 Phase 7: Model value potential and scalability

Phase 7 evaluates the value potential and scalability of the model for DHL, based on the objectives and experimental setup described in Section 5.5.8. While earlier experimental phases already provided meaningful managerial insights, this phase places a stronger emphasis on the practical applicability and value of the model in real-world settings. Section 6.7.1 assesses computational performance across realistic DHL decision regions. Section 6.7.2 explores its scalability across larger geographic datasets. Finally, Section 6.7.3 demonstrates the model's potential to reduce costs when applied to actual DHL use cases.

6.7.1 Operational practicality

To evaluate whether the model is efficient enough for real-world deployment, we tested it under DHL's preferred settings (as described in Section 5.5.8), using a placement limit of $U = 1$, 5, and 10 lockers. These tests were conducted for each RegioHub area individually (Figure 40), being geographic clusters that DHL considers coherent planning regions for locker deployment. As shown in Table 27, all runs were completed in 230 seconds, and the vast majority solved in less than 100 seconds. This confirms that even in larger geographic regions the model remains computationally efficient under realistic parameter settings.

Table 27: Solve Time (in seconds) per region for different U values

Region	$U = 1$	$U = 5$	$U = 10$
Alkmaar	73	79	106
Amersfoort	77	82	109
Amsterdam	89	87	92
Arnhem	76	77	86
Beek	71	74	84
Den Bosch	74	94	92
Den Haag	86	89	92
Drachten	120	148	230
Eindhoven	96	97	108
Hengelo	78	119	226
Roosendaal	97	106	125
Rotterdam	82	82	110
Utrecht	82	87	93
Zwolle	108	132	197

This figure has been removed due to confidentiality.

Figure 40: RegioHub regions DHL

6.7.2 Scalability and model limits

The model was also tested on the entire Netherlands region to evaluate how far it can be pushed for practical use by DHL. The model was solved at two different postcode resolutions: postcode level 4 (`PC4`), which provides a more aggregated representation of demand, and postcode level 5 (`PC5`), which is part of DHL's preferred planning configuration. All other parameters were kept consistent with the preferred DHL model settings. All experiments were run with a maximum solver time of 3600 seconds (1 hour), which is still a feasible duration for strategic decision-making of the largest decision area possible.

Table 28: Solve time (in seconds) across U values for different postcode levels, based on a model run on the entire Netherlands

Postcode Level	$U = 1$	$U = 5$	$U = 10$	$U = 50$	$U = 100$
`PC4`	169	64	74	120	353
`PC5`	2590	2733	3568	gap(0.04%)	gap(0.07%)

Table 28 summarises the solver times across different U values. The results for `PC4` show that even for the entire country, the model runs remain computationally efficient, with solve times well under 6 minutes, and typically below 2 minutes, demonstrating that national-scale decision-making is feasible. For `PC5`, the solver times increase significantly due to the higher spatial detail, with times ranging from 2590 to 3568 seconds for different U values. For higher values of U (50 and 100), the solver did not complete within the one hour time limit, resulting in gaps of approximately 0.04% and 0.07%, respectively. These gaps are relatively small and acceptable, indicating that the model is very close to providing the optimal solution. The results suggest that the model also remains usable for national-scale planning at `PC5` level, especially for smaller U values (< 50). These tests provide valuable insights into the model's scalability and inform future adaptations if DHL chooses to support decision making on a more detailed or larger scale.

6.7.3 Value potential: from current to optimised

Finally, to analyse the model's potential value, we compare the current network configuration with an optimised solution generated by the model. In this comparison, the existing infrastructure is held fixed, allowing the model to determine optimal expansions and delivery mode choices under current network conditions. This setup enables a fair assessment of whether the model can improve upon the existing situation. The fixed costs of existing lockers are excluded from both cases as they are equal and unaffected by model decisions. The detailed cost and configuration comparison is shown in Table 29 (Enschede), Table 30 (Hengelo), and Table 31 (Almelo). The tables clearly illustrate where changes occur: which additional lockers are placed (including their sizes), how delivery options are utilised, and how these affect the overall cost structure. The final rows of each table show the resulting difference in the daily objective function and total reduction in daily operating costs.

In the optimal solution, the model places additional lockers. This is explained by the lower CpP and greater cost-efficiency of PLs at higher demand levels compared to SPs and home delivery, as previously discussed. SPs and home delivery remain in use where demand is too sparse to justify a locker. However, the model strategically deploys new lockers to absorb the majority of demand wherever feasible. This aligns with earlier findings in this chapter, which showed that lockers consistently emerge as the most cost-effective delivery mode across various settings, making targeted placement in high-demand regions an effective way to reduce costs.

Across the three regions, the optimised configuration yields substantial savings, ranging from €20 to €100 per day, which corresponds to a cost reduction of 18% to 22%. This provides strong evidence that the model can significantly contribute to cost reduction. As discussed in Chapter 1, reducing operational costs is one of the action problems of DHL, and this result highlights the potential of the model to directly address this problem.

Table 29: Comparison – Enschede

Feature	Current	Model
Lockers (exist.)	21	21
New lockers	–	16
Sizes S/M/L/XL	–	5/5/4/2
SP usage	16	4
Var. locker cost	€47.18	€92.83
Fix. locker cost	–	€78.79
Var. SP cost	€121.70	€19.06
Fix. SP cost	€24.65	€6.16
Home cost	€104.59	€43.40
Total cost	**€298.12**	**€240.24**
Diff. per day	–	**€-57.88**
% diff	–	**-19.41%**

Table 30: Comparison – Hengelo

Feature	Current	Model
Lockers (exist.)	11	11
New lockers	–	11
Sizes S/M/L/XL	–	4/5/2/–
SP usage	9	2
Var. locker cost	€16.70	€34.50
Fix. locker cost	–	€32.58
Var. SP cost	€36.44	€5.07
Fix. SP cost	€8.50	€1.89
Home cost	€56.43	€18.61
Total cost	**€118.07**	**€92.65**
Diff. per day	–	**€-25.42**
% diff	–	**-21.53%**

Table 31: Comparison – Almelo

Feature	Current	Model
Lockers (exist.)	13	13
New lockers	–	12
Sizes S/M/L/XL	–	1/8/2/1
SP usage	11	2
Var. locker cost	€63.36	€146.55
Fix. locker cost	–	€139.47
Var. SP cost	€212.82	€12.87
Fix. SP cost	€39.91	€7.26
Home cost	€226.37	€135.93
Total cost	**€542.46**	**€442.08**
Diff. per day	–	**€-100.38**
% diff	–	**-18.51%**

6.8 Summary & conclusion

This chapter evaluated the model's behaviour across seven experimental phases, addressing the research question: *"What are the outcomes of the experimental design?"*

The results demonstrate that both the granularity of spatial input (e.g., grid resolution and postcode level) and the activation of model features (e.g., adaptive radius, locker constraints) significantly influence solver performance. While halving the grid spacing leads to exponential increases in runtime due to the explosion in candidate locations, enabling the fixed existing lockers feature substantially reduces computation time. Overall, the model remains computationally tractable across a broad range of configurations, validating its practical applicability for different use cases.

When experimenting with the maximum number of new lockers to be placed, the findings reveal a clear and intuitive pattern: the model initially prioritises locations with high unmet demand (typically served by home delivery), followed by areas currently relying on expensive or high-volume SPs, where a locker can offer a more cost-efficient alternative. As U, the maximum number of new lockers to place, increases, additional lockers begin to appear in rural or peripheral areas where cost savings are still achievable, often by deploying smaller lockers. In areas where demand is too low or where existing alternatives like SPs or home delivery are still cheaper, the model simply does not place any lockers. A diminishing returns curve, which can also support DHL's strategic decision making, illustrates that beyond a certain threshold, placing additional lockers no longer yields sufficient savings. At that point, the model naturally stops allocating new lockers.

Smaller radii result in dense locker placement within urban areas, but leave rural areas underserved. In contrast, larger radii improve rural coverage and reduce the total number of lockers needed, though at the cost of exponential longer runtimes and a loss of placement precision in densely populated areas. The adaptive radius offers a strong alternative: it adjusts coverage to local density, reaches rural demand efficiently, and solves relatively quickly. This makes it a practical and scalable extension, enabling DHL to run a single model across diverse regions without requiring multiple separate configurations while also achieving faster solution times.

As parcel demand increases, the model responds primarily by upscaling existing lockers rather than expanding the total number of locations. The experiment shows that over 82% of the additional volume is absorbed by upgrading locker sizes rather than increasing the number of locker locations, when applying a 50% volume increase. Moreover, locker locations selected under baseline demand consistently reappear in the solutions generated for higher demand scenarios, confirming their strategic positioning and the robustness of the network. While some small new lockers are added in dense or newly viable rural areas, the overall spatial layout remains stable. SPs gradually decline as lockers become more cost-efficient, and home delivery is mostly phased out, except in remote areas where demand remains too low to justify fixed infrastructure. These results demonstrate that the model supports scalable growth without requiring major spatial reconfiguration.

Currently, across the studied regions, lockers hold a dominant position with an average mode share of approximately 85%. Cost sensitivity analysis shows that the delivery mix responds intuitively, but differently, to various cost drivers. PLs remain the dominant mode in most scenarios, and among their cost components, fixed infrastructure costs have the strongest impact on their share. Home delivery keeps a crucial role in sparsely populated areas and continues to strongly influence both total cost and mode shares. SPs, by contrast, operate effectively only within a narrow cost-efficient range and show limited impact outside of it. A combined scenario shows that if home delivery costs rise while locker costs drop by the same percentage, total system costs still decrease, and locker use increases.

The experiments show that incorporating robustness is computationally feasible and allows DHL to evaluate trade-offs between protection levels and cost. Most configurations solve within a reasonable time, though high values for both Δ and Γ lead to steep increases in runtime. Robustness increases total cost due to added buffer capacity within the OOH network, and this cost rises non-linearly: costs escalate rapidly when both parameters increase, revealing a convex pattern. The model is more sensitive to the maximum deviation parameter (Δ) than to the deviation budget (Γ), suggesting that, from a cost perspective, addressing extreme deviations at a few locations is more expensive than protecting against more smaller, distributed fluctuations across the network. This reinforces the value of DHL's dynamic capacity control project, which aims to balance parcel flows and reduce the likelihood of single high spikes in demand. Moreover, the model's response to robustness aligns with the demand-sensitivity findings from Phase 4: uncertainty is primarily absorbed by scaling up capacities at the locations also selected under non-robust conditions, with minimal spatial reconfiguration required, as the robust and non-robust site selections remain nearly identical.

This final phase demonstrates that the model is both practically valuable and technically scalable under DHL's preferred (parameter) settings. Applied to real-world DHL regions, it produces optimised configurations that significantly reduce operating costs by improving locker placement and sizing, and delivery mode allocation, resulting in a cost reduction of 18% to 22%, providing strong evidence of the model's potential value in cost reduction. Runtime evaluations confirm that the model solves efficiently (within 0–4 minutes) across all DHL RegioHub areas, supporting its practical deployability in large-scale planning contexts. The interactive visual outputs further enhance decision-making by offering clear and interpretable insights into the model's solution. Finally, national-scale runs demonstrate that the model can handle even larger or more complex areas, suggesting that it could be applied to more detailed or expansive planning scenarios if needed.

The results confirm that the model is capable of generating realistic, consistent, and operationally valuable placement and sizing decisions for PLs across a variety of configurations and input conditions. The experiments demonstrate its sensitivity to key cost and uncertainty parameters and provide clear strategic managerial insights, such as when and where lockers are cost-effective, which can directly support DHL's OOH network planning. Furthermore, the intuitive and explainable outcomes across all experimental phases validate the model's internal logic.

7 Conclusion and Recommendations

This final chapter brings together the key findings of the research. It answers the main research question and summarises the core conclusions derived from the research. Based on these findings, it presents concrete managerial recommendations for DHL. Finally, it reflects on the limitations of the study and proposes directions for future research.

7.1 Conclusions

This research addressed the main research question: "*How can DHL eCommerce Benelux make data-driven strategic decisions regarding the placement and sizing of parcel lockers within its OOH network, to reduce operational costs and improve customer satisfaction?*" Despite the critical role of PL placement in DHL's long-term strategy, placement and sizing decisions for PLs are currently made based on intuition and driven by growth targets, rather than data-driven analysis. As a result, the current network experiences inefficiencies and increased operational costs, primarily caused by geographical mismatches between locker capacity, location, and actual demand.

To answer the main research question, a Mixed-Integer Linear Programming (MILP) model was developed, providing a formulation of the Last Mile Capacitated Parcel Locker Location Problem (LMCPLLP), to support the strategic placement and sizing of PLs within DHL's OOH network. The model allocates demand across delivery modes in a cost-efficient way, while determining optimal placement and sizing of PLs. Its framework, implementation, and data preprocessing account for the real-world constraints and operational realities of DHL, incorporating cost structures developed in this research that guide all placement and sizing decisions, and uses historical pickup behaviour to reflect realistic capacity usage. The model handles heterogeneous locker types with size-dependent capacity constraints, supports partial demand allocation, and includes fallback delivery modes, such as SPs and home delivery, to ensure full demand coverage where lockers are not viable. A robustness extension enables the assessment of worst-case local demand fluctuations, allowing DHL to evaluate trade-offs between cost efficiency and network resilience.

To better align with DHL's planning requirements and real-world use cases, additional features were introduced to the model. These include the integration of existing lockers (fixed infrastructure), placement limits for new lockers, and an adaptive pickup radius that scales according to local density. These features enhance both scalability and adaptability, enabling design of new networks (greenfield design), incremental expansion of the current network, and deployment across heterogeneous areas without the need for separate configurations. Furthermore, the study provides DHL with an interactive GUI for both input and output. This interface, tailored to DHL's strategic needs, allows a user-friendly configuration of model parameters and offers interactive visualisation of the results, with customisable layers for detailed analysis.

Seven experimental phases were conducted to evaluate the model's behaviour under various settings. These experiments validated the model's internal logic, confirmed by both DHL domain expert reviews and intuitive, explainable solution patterns, while also assessing the model's sensitivity to key parameters and providing managerial insights. The model results show that:

- PLs emerge as the preferred delivery mode, capturing an average share of 85% at current cost levels across all experiments, and remain the dominant choice even under moderate cost variations. Only under extreme cost increases do lockers lose their dominance, highlighting their strategic robustness. This supports DHL's long-term locker expansion strategy, concluding that lockers are not only operationally scalable but also the most cost-efficient mode in most real-world conditions.
- The model demonstrates a cost reduction of 18-22% across three real-world test regions, providing strong evidence of the value of the model in optimising locker placement and sizing to decrease DHL's operational costs.
- The model solves each DHL RegioHub area within 0-4 minutes in DHL's preferred settings, demonstrating its scalability for large-scale planning. It also solves full-scale Netherlands scenarios in under one hour, even at detailed postcode levels, confirming national-level scalability.
- As demand grows, locker locations remain relatively stable, with primarily locker sizes increasing, indicating the model's long-term effectiveness in initial placements.

- The adaptive radii effectively adjust to smaller radii in high-density areas and larger radii in low-density areas, without the need for separate runs, while also solving relatively quickly.
- Both the granularity of spatial input (e.g., grid resolution and postcode level) and the activation of model features (e.g., adaptive radius, locker constraints) significantly influence solver performance.
- The model is sensitive to cost shifts, particularly in locker and home delivery costs, affecting the delivery mode selection and objective function.
- Robust optimisation scenarios showed that the model remains stable under demand uncertainty. Experiments show that addressing extreme, concentrated demand spikes tends to be more costly than mitigating multiple moderate, distributed fluctuations, further highlighting the importance of DHL's dynamic capacity control project. The local worst-case uncertainty is primarily absorbed by increasing capacities, while locker locations remain relatively stable, further indicating the model's effectiveness in initial (non-robust) placements.

7.2 Contribution to theory

This thesis contributes to the academic literature on PL facility location problems and OOHD networks by proposing a novel MILP-based optimisation framework for the strategic placement and sizing of PLs within an OOH network. The model addresses several key gaps identified in previous research. A complete contribution statement and a detailed explanation of the theoretical gaps addressed can be found in **Section 3.4 Contribution statement**.

To briefly summarise, this study contributes a novel MILP-based framework that:

- integrates PLs and SPs in a capacitated FLP for PLs;
- models overflow via fallback/delivery options to SPs and home delivery;
- allows full fractional (non-binary) demand allocation within a radius;
- supports region-adaptive pickup radii for geographic pickup realism;
- integrates robust optimisation for fractional demand and mixed facility types, while maintaining linear solvability.

First, the inclusion of novel features introduces elements not present in existing models. In addition to these individual features, the integrated combination of both novel and established modelling features, as can be seen in Table 9 in Chapter 3, into a single framework results in a unique approach and framework not explored in previous studies. Together, this model contributes to key literature gaps and supports DHL's ambition of data-driven decision support in locker placement and sizing within its existing OOH network.

7.3 Contribution to practice

This section demonstrates the practical impact of this research by highlighting its contributions to decision-making quality, operational usability, and broader applicability to strategic planning within the logistics sector and beyond.

First, the developed MILP model demonstrated cost reductions of 18–22% in tested DHL regions using real-world data, confirming its strategic value and design quality. In addition, the study offers directly applicable concrete recommendations that DHL can implement to improve strategic decision-making, enhance operational efficiency, increase customer satisfaction for both senders and receivers, and strengthen long-term network resilience.

Second, the model has been developed with a strong focus on usability and practical deployment. It runs efficiently for the large-scale RegioHub areas of DHL (0 to 4 minutes per region) on standard hardware and is implemented as a standalone Python tool with an intuitive GUI. This allows planners without programming or mathematical expertise to configure input parameters and model features via a dashboard, and explore the model results through interactive maps with layered toggles and clickable elements. These features support detailed spatial analysis and enable direct interpretation. The model also translates strategic optimisation outcomes into intuitive spatial outputs, allowing decision-makers across departments to interpret and communicate network scenarios more effectively.

Third, the model is designed as a modular framework, offering high flexibility across a wide range of planning contexts. Users can activate or deactivate key features or extensions depending on strategic needs. This configurability allows the model to address diverse use cases, ranging from incremental network expansion to fully greenfield network designs, or specific settings for experimentation and the evaluation of different future strategic plans, each producing substantially different and context-specific outcomes. Moreover, the model accepts any combination of input data, such as potential and existing locker locations, cost structures, SPs, and demand nodes. Although it was applied to DHL's RegioHub and CityHub areas in this study, its flexible design allows for easy adaptation to different organisational infrastructures or extension to new geographical regions, making it broadly applicable across diverse planning contexts.

Beyond the parcel domain, the model's optimisation logic and modular structure are broadly applicable to other sectors facing spatial planning and capacity allocation challenges. Any context that requires the strategic placement of limited-capacity facilities based on distributed demand patterns can benefit from this framework. For example, shared mobility providers, such as scooter, bike, or car-sharing platforms, can use our modelling logic to determine the optimal locations for vehicle deployment or charging infrastructure. Even outside logistics, applications exist in sectors like healthcare, where the framework could guide the placement of mobile testing units, pop-up vaccination sites, or care facilities in response to local needs.

Finally, this research provides a practical roadmap for companies aiming to adopt advanced (facility) optimisation solutions. It covers the entire process, from context analysis and problem formulation to data preparation, model development, deployment, and performance evaluation, it offers a step-by-step methodology to translate complex operational challenges into actionable, data-driven decision frameworks. As such, the study not only delivers a solution to a specific case but also serves as a replicable guide for organisations seeking to enhance their logistics networks through structured optimisation practices.

7.4 Recommendations

- **Embed the optimisation model into DHL's strategic network planning**

To replace the current intuition-based approach to PL placement and sizing, DHL is advised to structurally adopt the developed optimisation model as a core decision-support tool. The model directly addresses the core research problem by enabling data-driven decisions for the OOH network that better aligns locker locations and capacity with local demand. The accompanying GUI allows planners to easily configure input parameters and interactively explore model outputs, without requiring complex training. This supports transparent decision-making, in-depth scenario analysis, and improved visibility of the OOH network's structure and performance.

- **Promote the use of parcel lockers over alternative delivery modes to increase cost-efficiency and scalability**

The experimental results consistently show that PLs are DHL's most cost-efficient delivery mode, capturing on average over 85% of OOH volume at current cost levels. Their dominance persists even under moderate cost fluctuations and growing demand, confirming that lockers remain the most cost-effective option in most scenarios. This also demonstrates their scalability and robustness as a sustainable backbone for DHL's future OOH network. Therefore, DHL is advised to actively promote the usage of PLs over alternatives such as SPs or home delivery, for instance, by setting PLs as the default delivery option, introducing incentives (price) or prioritising lockers in marketing communication. This would help increase cost-effectiveness and reduce operational pressure on less scalable delivery modes. Further research or internal experimentation may help determine which measures are most effective in encouraging locker use, while maintaining customer satisfaction.

- **Integrate real-time locker availability data into the sender interface**

As diagnosed in the context analysis, real-time data on locker availability is already accessible within DHL's internal systems, but it is not exposed to senders via the public interface. As a result, senders cannot check the occupancy status of the locker before selecting a drop-off point at, for example, the DHL website or app. This lack of transparency can lead to failed drop-off attempts, requiring senders to travel to a second location, ultimately harming the user experience and reducing customer satisfaction. DHL is advised to integrate this real-time information into the sender interface to improve customer satisfaction and transparency.

- **Develop dynamic reallocation mechanisms for overloaded OOH points**

When a PL is full, drivers currently select an alternative drop-off point manually. This process lacks dynamic logic based on proximity, local capacity, or overall network load. As identified in the context analysis, this often leads to repeated diversions to a few 'driver-preferred' fallback locations, typically large and easily accessible SPs, causing local congestion while underusing other sites. Although DHL's planned dynamic capacity control system is expected to mostly eliminate such issues, some overflow will likely remain. DHL is therefore advised to expand this system with a real-time diversion mechanism that applies the same logic and also incorporates short-term OOH point capacity usage forecasts. This would enable more intelligent reallocation, improve load balance, reduce delivery inefficiencies, improve locker availability, and strengthen overall network resilience.

- **Develop data-driven insight into the reasons behind parcel diversions**

While parcel diversions are currently tracked, the underlying causes remain unclear. As diagnosed in the context analysis in Chapter 2, it is not recorded whether reroutes stem from capacity constraints, courier preferences, technical issues, or local deviations from standard procedures. As a result, DHL lacks the information needed to effectively target and reduce diversions. Therefore, it is recommended that DHL expands its data collection to explicitly capture the reasons behind each diversion.

- **Proactively scale locker sizes at key locations to accommodate future demand and improve robustness, if financially possible**

If DHL intends to prepare early for higher future volumes, it may be effective to install larger lockers than currently necessary at strategically selected locations based on current demand, even if this involves a higher upfront investment. This recommendation is supported by both the demand growth and robustness-extension experiments.

The results indicate that most future demand can be absorbed by increasing the size of the lockers at existing sites, without the need for a widespread expansion of new OOH locations. In fact, the experiments show that under a 50% increase in total demand, approximately 82% of the additional volume is absorbed by upsizing existing locker locations, highlighting the effectiveness of scaling current locations as a proactive strategy to accommodate expected growth. Robustness experiments support this same conclusion: uncertainty is almost entirely absorbed by increasing capacities at the same locations selected under non-robust conditions. This reinforces the recommendation to proactively scale locker sizes at key sites, as these locations also remain effective even under local worst-case demand spikes.

This recommendation also aligns well with DHL's dynamic capacity control system, which routes parcels to nearby available lockers. That system performs best in a network with broadly distributed moderate excess capacity, allowing greater operational flexibility during short-term volume spikes. By proactively scaling locker sizes, DHL can strengthen both its long-term growth preparedness and day-to-day resilience, provided that the additional investment costs are acceptable within its strategic budget.

- **Retain service points to delay large-scale locker expansion and increase network resilience**

The limited need for additional locker locations, observed in the demand growth experiments, is partly enabled by the residual capacity of existing SPs. While the model tends to favour PLs at higher demand levels due to their superior cost-efficiency, such placements are often not financially sustainable when demand increases are only temporary. In these cases, replacing SPs is not strictly necessary, as their substantial residual capacity allows them to absorb the volume anyway. Their cost structure, with no fixed costs and only per-parcel fees, makes SPs particularly suited to act as a structural buffer during longer seasonal peaks, demand fluctuations, or gradual volume growth. They become essential for absorbing additional demand without triggering immediate infrastructure investments. Such investments are more cost-effective when postponed until growth stabilises or reaches a sufficiently high constant level to justify permanent capacity expansion. In the meantime, service points remain a more financially efficient solution than home delivery for handling these volumes.

Moreover, if home delivery becomes structurally more expensive, due to, for example, reduced volumes or increasing environmental pressure, SPs, if available, can absorb low-density demand that cannot be efficiently served by lockers. This makes SPs strategically important, not only as a transitional and buffering mode, but also as a fallback option in areas where lockers remain unviable. Additionally, the SPs selected and used within the model's optimal solution are not only useful as transitional or fallback options, but also represent the most cost-effective choice at their specific locations. This makes them already operationally valuable, as they are the best solution to that local demand.

DHL is therefore advised to actively maintain SPs during the present period, in which recruiting new SP partners has become increasingly difficult (as explained in Chapter 2), due to the strategic value of SPs for both the current and future OOH network. Preserving these locations for as long as operationally viable allows DHL to operate efficiently until a complete transition to a PL network becomes feasible.

- **Prioritise fixed locker cost reduction to enable wider PL network viability**

The cost sensitivity experiments reveal that both fixed and variable costs of PLs significantly influence the model's objective value. This is expected, given that PLs capture the largest share of OOH demand in the optimised network, making cost shifts in this delivery mode highly impactful. However, reducing fixed costs shows a distinct additional advantage over variable cost reductions. Specifically, lowering fixed costs leads to a visible increase in PL adoption and enables lockers to take over volumes that would otherwise be served by home delivery, particularly in low-density or rural areas. This effect does not occur when only variable costs are reduced, as the initial investment threshold remains too high to justify locker placement in these areas. DHL is therefore advised to focus cost optimisation efforts on fixed cost reduction, through e.g., supplier agreements, modular locker designs, or long-term depreciation strategies to increase the reach and viability of PLs across the full network.

- **Monitor and manage cost levels of home delivery relative to other delivery modes**

In addition to the strong sensitivity to PL cost shifts, home delivery also has a smaller, but still notable impact on total daily costs. This is mainly because home delivery remains crucial in low-density rural areas, where the substitution potential of PLs is limited. Maintaining current home delivery cost levels is therefore important for DHL to prevent sudden increases in total system costs. At the same time, combined scenario testing shows that reducing locker costs while home delivery becomes equally more expensive (in percentage terms) results in a net cost reduction and increased locker usage. This suggests that if home delivery becomes structurally more expensive in the future, due to, for example, shrinking volumes or reduced economies of scale, there may be strategic room for DHL to adjust home delivery pricing accordingly. DHL is therefore advised to closely monitor home delivery cost levels, given its continued necessity in certain areas, and to adjust pricing between modes where possible to maintain overall network cost efficiency.

7.5 Limitations and future research

The limitations discussed in this section stem from a combination of modelling choices, simplifying assumptions, and data availability constraints. Data access was limited due to the absence of a direct connection to DHL's Oracle-based systems, restricting the analysis to aggregated and preprocessed datasets. Moreover, detailed OOH data was only consistently available for the period between 1 November 2024 and 27 March 2025, further constraining the temporal scope of the analysis. Based on these underlying choices and constraints, the remainder of this section outlines key limitations and suggests directions for future research.

First, the model relies on static average daily demand per location. This is motivated by the strategic focus of the model and the planned implementation of a dynamic capacity control system to manage intra-day variability. In addition, daily-level demand distributions per locker were not accessible due to data restrictions. To still account for uncertainty and growth, the robustness extension and demand scaling experiments provide directional safeguards against local worst-case demand uncertainties, longer seasonal peaks, and long-term growth. However, despite these measures, this representation remains a simplification: real-world locker usage is inherently dynamic, with temporal and spatial fluctuations not fully captured by average demand inputs. As such, the model's outcomes may deviate from operational realities in some cases. If more precise or short-term capacity estimation is required, future research could explore simulation-based optimisation approaches that explicitly capture temporal fluctuations, both within and across days, and assess their operational impact on locker performance and required capacity. Moreover, demand is spatially aggregated at postcode level using centroid coordinates (e.g., `PC4/5/6` depending on the settings), which may oversimplify local variation in delivery patterns or distort actual distances between customers and OOH points. While often negligible, centroid-based modelling may lead to inaccurate representations of customer-to-facility distances in larger postcode areas, especially when customers live near the edges of such regions. Modelling demand at the individual address level could improve location accuracy and better reflect real-world proximity effects, although this would increase computational complexity. Future work could explore (meta) heuristic approaches to maintain tractability while capturing finer spatial granularity.

Second, the modelling approach accounts for parcel pickup behaviour within the locker capacity using an average delay across all lockers, based on a discrete approximation of Little's Law and the overall pickup-times distribution. Due to data limitations, no location-specific or time-dependent pickup data was available. As a result, the model neglects regional or seasonal variation in collection behaviour, which may lead to periodic over- or underestimation of capacity pressure at specific sites and reduce sizing accuracy. However, if such data were available, incorporating it into the model would be a relatively straightforward extension.

Next to that, the optimisation model does not account for individual compartment sizes or the actual dimensions of parcels. This simplification, driven by the lack of data on parcel and compartment size distributions and the strategic scope of the research, may result in an overestimation of usable capacity in practice. While DHL currently mitigates this by relying on local expert opinion to determine appropriate locker configurations during actual placement, the model may still overestimate the actual service levels that can realistically be achieved for both LM and FM deliveries once the locker is deployed. Future research could explore a model extension that incorporates locker layouts with predefined compartment mixes (e.g., small, medium, large). Instead of modelling capacity as a single aggregate value, each locker type would define the number of parcels it can hold per size category. Parcel demand would be disaggregated accordingly and constraints updated to ensure a feasible allocation. An extra rule could be added, allowing smaller parcels to use larger compartments when needed, but not the reverse. Given our modular MILP structure, this extension is straightforward to implement and could improve allocation realism, provided that reliable data becomes available.

In addition, the cost component of the model assumes a fixed locker lifespan of 12 years based on current expectations of DHL. However, this lifespan is uncertain. If actual lifespans are shorter, the effective annual cost increases, reducing cost-efficiency. If they are longer, lockers become more financially attractive. This uncertainty may affect optimal sizing and placement decisions. Future research could address this by applying a probabilistic depreciation model that captures the likelihood of different lifespan scenarios.

The model also assumes that customers are indifferent between SPs and PLs, aligning with DHL's goal to promote a uniform OOH experience. However, real-world preferences may vary due to factors like accessibility, operating hours, or perceived convenience. In addition, the model applies distance thresholds to define customer reach, while actual willingness to travel may vary between individuals, even within the same urban or rural context, or deviate from the thresholds defined in this research in collaboration with DHL expert insights. This simplification may limit the accuracy of demand allocation. Future research could integrate customer preferences by incorporating mode-specific utility parameters or choice probabilities into the model, to better mimic receiver/sender behaviour. This might involve approaches such as a nested logit model informed by behavioural data.

Furthermore, the model identifies cost-optimal locker locations under idealised assumptions of full placement flexibility. While this provides valuable directional insight, guiding planners toward high-potential areas, real-world legal, spatial, or operational constraints may prevent installation at the exact suggested sites. As a result, realised cost savings or service levels may be lower than projected.

Finally, the model optimises locker placement solely from a cost-minimisation perspective, without accounting for broader strategic objectives such as market expansion or long-term brand positioning. While this aligns with DHL's current focus, it overlooks potential trade-offs that may be relevant under more growth-orientated strategies. For instance, a company might accept short-term inefficiencies or higher costs to pre-empt competitors, secure prime locations, or accelerate adoption of OOH delivery. These non-cost-optimal decisions may yield long-term strategic value. Future research could extend this work by explicitly incorporating such trade-offs and identifying conditions under which deviations from cost-efficiency become strategically justified.

References

Balinski, M. L. (1965). Integer programming: Methods, uses, computations. *Management Science*, *12*(3), 253–313. https://doi.org/10.1287/mnsc.12.3.253

Bertsimas, D., & Sim, M. (2004). The price of robustness. *Operations Research*, *52*(1), 35–53. https://doi.org/10.1287/opre.1030.0065

Centraal Bureau voor de Statistiek. (2023). *Postcode areas and demographic statistics 2023 (pc4, pc5, pc6)* [Includes polygon geometries (GeoPackage files) and tabular population statistics]. Retrieved May 7, 2025, from https://www.cbs.nl/en-gb/dossier/the-netherlands-regionally/geodata/data-per-postcode

de Armas, J., Juan, A. A., Marquès, J. M., & Pedroso, J. P. (2017). Solving the deterministic and stochastic uncapacitated facility location problem: From a heuristic to a simheuristic. *Journal of the Operational Research Society*, *68*(10), 1161–1176. https://doi.org/10.1057/s41274-016-0155-6

Deutsch, Y., & Golany, B. (2018). A parcel locker network as a solution to the logistics last mile problem. *International Journal of Production Research*, *56*(1-2), 251–261. https://doi.org/10.1080/00207543.2017.1395490

DHL eCommerce Netherlands. (2025). *Over ons – dhl ecommerce netherlands*. Retrieved February 18, 2025, from https://www.dhlecommerce.nl/nl/over-ons

DHL Express Belgium. (2025). *Our story – dhl express belgium*. Retrieved February 18, 2025, from https://www.dhlexpress.be/en/our-story/

DHL Group. (2025a). *Dhl ecommerce solutions*. Retrieved February 18, 2025, from https://www.dhl.com/nl-en/home/ecommerce.html

DHL Group. (2025b). *Über uns - dhl group*. Retrieved February 18, 2025, from https://group.dhl.com/de/ueber-uns.html

Faugère, L., & Montreuil, B. (2018). Smart locker bank design optimization for urban omnichannel logistics: Assessing monolithic vs. modular configurations. *Computers & Industrial Engineering*, *130*, 393–406. https://doi.org/10.1016/j.cie.2018.11.054

Heerkens, H., & van Winden, A. (2017). *Solving managerial problems systematically* [Translated into English by Jan-Willem Tjooitink]. Noordhoff Uitgevers.

Janinhoff, L., Klein, R., Sailer, D., & Schoppa, J. M. (2024). Out-of-home delivery in last-mile logistics: A review. *Computers & Operations Research*, *168*, 106686. https://doi.org/10.1016/j.cor.2024.106686

Kahr, M. (2022). Determining locations and layouts for parcel lockers to support supply chain viability at the last mile. *Omega*, *113*, 102721. https://doi.org/10.1016/j.omega.2022.102721

Lee, H., Chen, M., Pham, H. T., & Choo, S. (2019). Development of a decision making system for installing unmanned parcel lockers: Focusing on residential complexes in korea. *KSCE Journal of Civil Engineering*, *23*(6), 2713–2722. https://doi.org/10.1007/s12205-019-1398-y

Lin, Y. H., Wang, Y., He, D., & Lee, L. H. (2020). Last-mile delivery: Optimal locker location under multinomial logit choice model. *Transportation Research Part E*, *142*, 102059. https://doi.org/10.1016/j.tre.2020.102059

Luo, R., Ji, S., & Ji, Y. (2022). An active-learning pareto evolutionary algorithm for parcel locker network design considering accessibility of customers. *Computers & Operations Research*, *141*, 105677. https://doi.org/10.1016/j.cor.2021.105677

Lyu, G., & Teo, C.-P. (2022). Last mile innovation: The case of the locker alliance network. *Manufacturing & Service Operations Management*, *24*(5), 2425–2443. https://doi.org/10.1287/msom.2021.1000

Mancini, S., Gansterer, M., & Triki, C. (2023). Locker box location planning under uncertainty in demand and capacity availability. *Omega*, *120*, 102910. https://doi.org/10.1016/j.omega.2023.102910

Northwestern University. (2022). Facility location problems [Accessed: 11 November 2022]. Retrieved November 11, 2022, from https://optimization.mccormick.northwestern.edu/index.php/Facility_location_problems

Ottaviani, F. M., Zenezini, G., Marco, A. D., & Carlin, A. (2023). Locating automated parcel lockers (apl) with known customers' demand: A mixed approach proposal. *European Journal of Transport and Infrastructure Research*, *23*(2), 24–45. https://doi.org/10.18757/ejtir.2023.23.2.6786

Pitney Bowes. (2023). 2023 global parcel shipping index. Retrieved February 26, 2025, from https://www.pitneybowes.com/content/dam/pitneybowes/us/en/shipping-index/23-mktc-03596-2023_global_parcel_shipping_index_ebook-web.pdf

Rabe, M., Gonzalez-Feliu, J., Chicaiza-Vaca, J., & Tordecilla, R. D. (2021). Simulation-optimization approach for multi-period facility location problems with forecasted and random demands in a last-mile logistics application. *Algorithms*, *14*(2), 41. https://doi.org/10.3390/a14020041

Raviv, T. (2023). The service points' location and capacity problem. *Transportation Research Part E: Logistics and Transportation Review*, *176*, 103216. https://doi.org/10.1016/j.tre.2023.103216

Reiffer, A., Kübler, J., Briem, L., Kagerbauer, M., & Vortisch, P. (2023). Analysing long-term effects of the covid-19 pandemic on last-mile delivery traffic using an agent-based travel demand model [Working paper, Karlsruhe Institute of Technology. Accessed: 2025-02-26]. *n/a*. https://www.semanticscholar.org/paper/Analysing-long-term-effects-of-the-Covid-19-pandemic-Reiff-K%C3%BCbler/d5681c7ebe32656a2835583040f9a84baae6b9aa

Savelsbergh, M., & Woensel, T. V. (2016). City logistics: Challenges and opportunities. *Transportation Science*, *50*(2), 579–590. https://doi.org/10.1287/trsc.2016.0675

Sawik, B., Serrano-Hernandez, A., Muro, A., & Faulin, J. (2022). Multi-criteria simulation-optimization analysis of usage of automated parcel lockers: A practical approach. *Mathematics*, *10*(23), 4423. https://doi.org/10.3390/math10234423

Song, L., Cherrett, T., McLeod, F., & Guan, W. (2009). Addressing the last mile problem: Transport impacts of collection and delivery points. *Transportation Research Record*, *2097*(1), 9–18. https://doi.org/10.3141/2097-02

Sweidan, A., Elomri, A., & Kerbache, L. (2022). Quantifying smart parcel station network usage as a logistical solution for the last-mile problem. *IFAC PapersOnLine*, *55*(10), 127–132. https://doi.org/10.1016/j.ifacol.2022.09.379

Tadic, S., Krstic, M., Stevic, Ž., & Veljovic, M. (2023). Locating collection and delivery points using the p-median location problem. *Logistics*, *7*(1), 10. https://doi.org/10.3390/logistics7010010

Ulukan, H. Z., & Demircioğlu, M. E. (2015). A survey of discrete facility location problem. *Proceedings of the International Conference on Operations Management and Industrial Engineering (ICAMIE)*, 2332–2337.

Wadhwa, V. M., & Garg, D. (2011). Facility location problem using genetic algorithm: A review [Accessed: 2025-02-26]. *Soft Computing for Problem Solving (SocProS 2011)*, 177–189. https://doi.org/10.1007/978-81-322-1602-5_16

Wang, Y., Bi, M., Lai, J., & Chen, Y. (2020). Locating movable parcel lockers under stochastic demands. *Symmetry*, *12*(12), 2033. https://doi.org/10.3390/sym12122033

Wang, Y., Zhang, Y., Bi, M., Lai, J., & Chen, Y. (2022). A robust optimization method for location selection of parcel lockers under uncertain demands. *Mathematics*, *10*(22), 4289. https://doi.org/10.3390/math10224289

Xu, X., Shen, Y., Chen, W., Gong, Y., & Wang, H. (2021). Data-driven decision and analytics of collection and delivery point location problems for online retailers. *Omega*, *100*, 102280. https://doi.org/10.1016/j.omega.2020.102280

A Appendix

A.0.1 Overview CityHubs

This figure has been removed due to confidentiality.

Figure 41: CityHub netwerk

A.0.2 Afternoon trip workflow

This figure has been removed due to confidentiality.

Figure 42: Afternoon trip workflow for OOH point collections

A.0.3 Morning trip workflow

This figure has been removed due to confidentiality.

Figure 43: Morning trip workflow for OOH point deliveries

A.0.4 Geographic distribution capacity parcel lockers

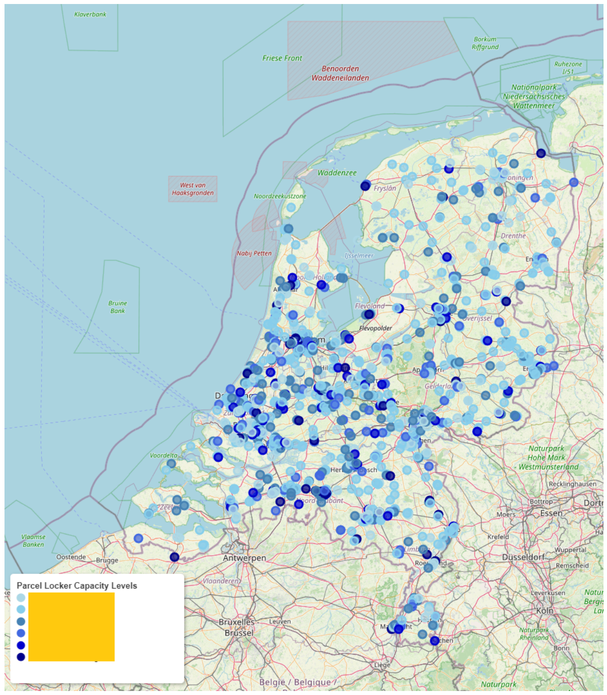

Figure 44: Geographic distribution

A.0.5 Scatter plot turnover against utilisation

This figure has been removed due to confidentiality.

Figure 45: Scatter plot: turnover vs utilisation, over the period 01-11-24 till 27-03-25.

A.0.6 Summary statistics turnover

This table has been removed due to confidentiality.

Table 32: Descriptive statistics of turnover for parcel lockers and service points

A.0.7 Heatmaps route diversion from/to

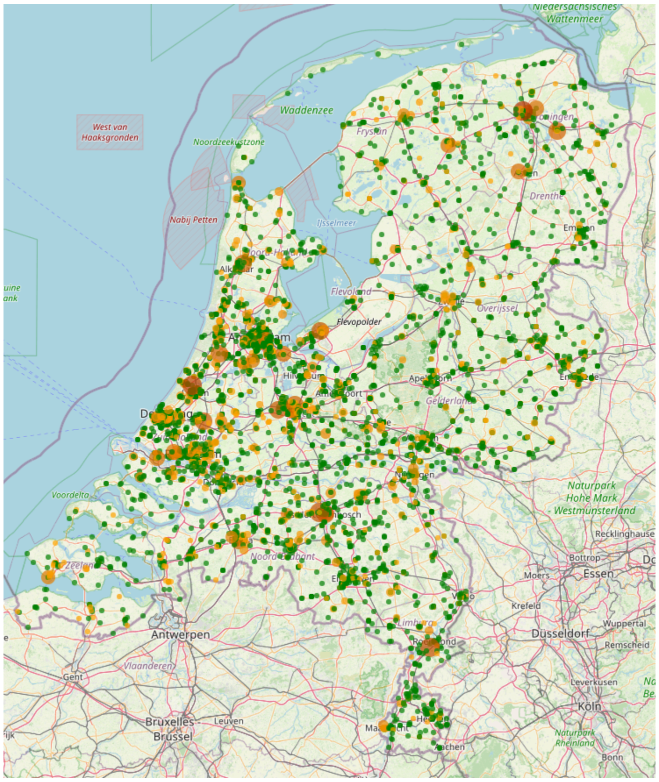

Figure 46: Heatmap route diversion at service points

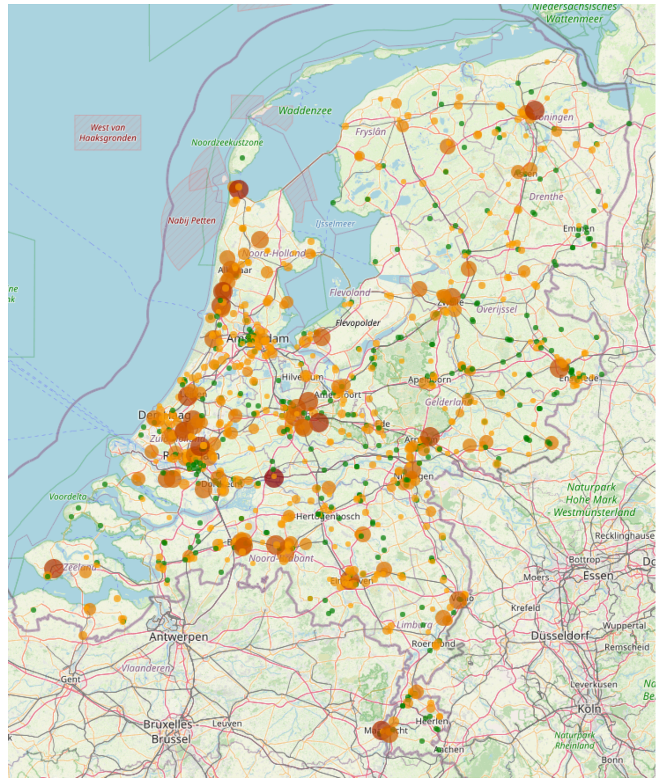

Figure 47: Heatmap route diversion at lockers

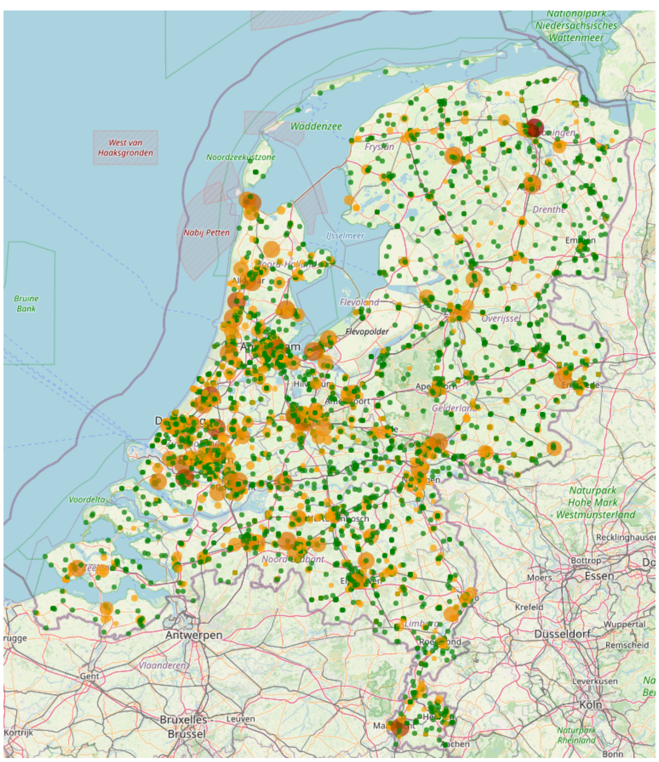

Figure 48: Heatmap diversion to Service Points

Figure 49: Heatmap diversion to Lockers

A.0.8 Geographical distribution of average parcel locker utilisation (hours 11–13)

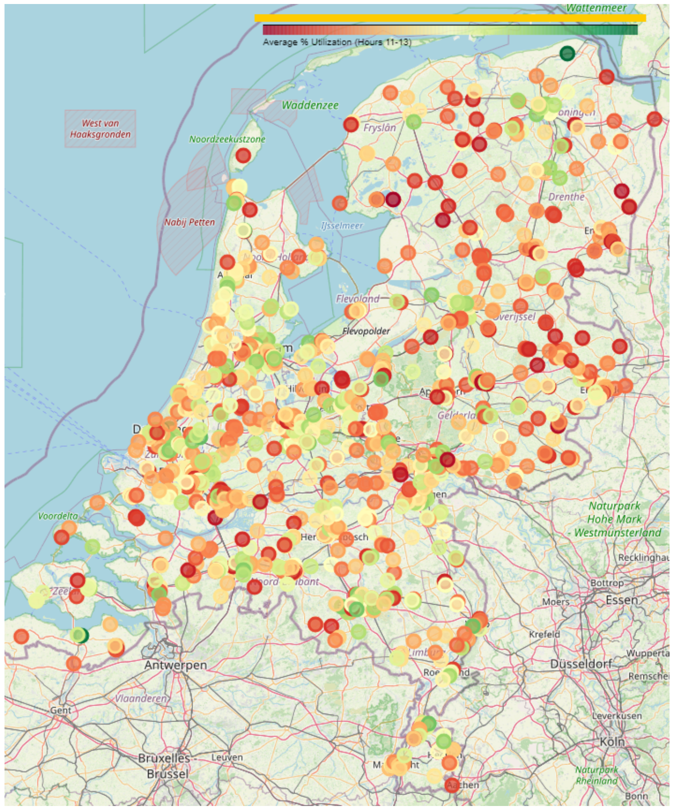

Figure 50: Geographical distribution of average parcel locker utilisation (hours 11–13) over period 01-11-24 till 27-03-25

A.0.9 Outcome robustness experiment

Table 33: Objective values for the almelo region under different robustness parameter settings

Δ (Maximum relative demand deviation)	Γ (Uncertainty budget in %)					
	0.0	0.2	0.4	0.6	0.8	1.0
0.0	906.89	906.89	906.89	906.89	906.89	906.89
0.2	906.89	909.67	911.93	913.42	915.29	916.60
0.4	908.87	912.72	919.18	925.11	930.20	932.08
0.6	910.64	917.99	931.01	939.90	945.12	947.35
0.8	912.94	924.53	940.32	956.48	966.16	968.92
1.0	913.90	931.54	952.86	970.91	980.30	982.28