INTEGRATED DEMAND MANAGEMENT AND VEHICLE ROUTING PROBLEM IN ATTENDED HOME DELIVERY

Mohammad Abdollahi

A thesis submitted for the degree of *Doctor of Philosophy* in

Operational Research

School of Mathematics, Statistics and Actuarial Science

University of Essex

June 2024

بسم الله الرحمن الرحيم

Abstract

This thesis presents a comprehensive analysis of Attended Home Delivery (AHD) systems within the rapidly expanding e-commerce sector. The research addresses the need for optimising delivery routes, refining pricing strategies, and enhancing customer satisfaction while improving profitability margins. It introduces a new approach to last-mile logistics by incorporating a partially time-windowed dynamic routing method with forecast orders. This method utilises forecast orders without time constraints, enabling more efficient time slot allocation for real orders. This approach reduces delivery costs and increases order acceptance rates, thereby boosting overall profitability.

Building on this concept, the thesis explores the synergy between demand management and vehicle routing with time windows. It proposes an enhanced methodology for estimating opportunity costs and introduces a dynamic slot-combination strategy. These techniques improve displacement cost assessments, informing better decisions regarding delivery charges. The practicality and efficacy of these methodologies are validated through extensive experiments using real-world data, demonstrating that the proposed solutions outperform current state-of-the-art approaches in profitability and delivery efficiency.

Furthermore, the research integrates the impact of delivery price changes, resulting from reduced flexibility in available time slots for future customers, into the opportunity cost estimation framework. This integration is crucial for revenue management and dynamic pricing within AHD, addressing the delivery price as a standalone factor in total revenue. The study refines the approximation of opportunity costs and integrates with existing solution approaches in the AHD problem domain, aligning these approaches with current demand management strategies. This thesis offers practical solutions that address the evolving needs of the AHD industry in an increasingly digital marketplace.

In summary, this thesis contributes to last-mile logistics by providing innovative, cost-effective solutions for the AHD sector. These solutions address trends in online shopping, offering insights and strategies to navigate the complexities of modern e-commerce logistics. The research combines theoretical innovation with practical application, aiming to optimise delivery systems in the digital age.

Dedication

I hereby declare that this thesis entitled "Integrated Demand Management and Vehicle Routing Problem in Attended Home Delivery" is my original work, carried out under the supervision of Prof. Xinan Yang and Dr. Michael Fairbank, at University of Essex, and submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy (PhD) in Operational Research.

I, Mohammad Abdollahi, affirm that:

1. The content of this thesis is entirely my own work unless otherwise indicated through proper citation and acknowledgement.

2. All the sources and references used in this thesis have been duly cited and acknowledged.

3. The work of others incorporated in this thesis, whether published or unpublished, has been appropriately referenced and attributed.

4. This thesis has not been submitted for any other degree or qualification to any other institution.

5. The data and results presented in this thesis are accurate to the best of my knowledge and are based on valid experiments and research.

I understand that any misrepresentation or omission in this declaration may lead to the disqualification of my thesis. I am aware of the university's policies on academic integrity and plagiarism, and I am committed to upholding the highest standards of academic honesty.

Mohammad Abdollahi

Acknowledgement

I would like to extend my deepest gratitude to all those who steadfastly supported and guided me throughout the journey of completing this thesis.

I am profoundly appreciative of my supervisors, Prof. Xinan Yang and Dr. Michael Fairbank, for their invaluable guidance, unwavering encouragement, and steadfast support. Their wealth of expertise, profound insights, and constructive feedback have been pivotal in not only shaping the trajectory of my research but also elevating the overall quality of this work.

My heartfelt appreciation extends to the esteemed faculty members of the School of Mathematics, Statistics and Actuarial Science (SMSAS) at the University of Essex. Their provision of an intellectually invigorating environment for exploration and learning has indelibly enriched my understanding and broadened my perspective.

To the University of Essex, I am profoundly grateful for the financial support that enabled me to channel my focus into research without the burden of distraction. Their benevolence has played an indispensable role in facilitating the successful culmination of this thesis.

I am compelled to acknowledge and extend my gratitude to my dear wife, Fatemeh Mousavi, and my cherished children, Maryam Abdollahi and Mohammad Javad Abdollahi. Their unwavering assistance, valuable insights, boundless encouragement, and unwavering understanding have significantly contributed to every stage of this journey.

Lastly, my heartfelt thanks goes to my family, whose enduring love, limitless patience, and unswerving belief in me have been the cornerstone of my perseverance. Their unwavering encouragement and unwavering support have been my constant motivation throughout this ambitious endeavour.

This thesis stands as a testament to the collective endeavours and contributions of all those aforementioned. Your roles have been pivotal in shaping this academic odyssey, and I extend my deepest appreciation for being an integral part of this journey.

List of Variables

Variable	Definition
α_{is}	Binary variable, 1 if order i is allocated to time slot s, 0 otherwise
$\mathbf{1}_{as}$	One more order will be delivered to area $a \in A$ in time slot $s \in S$
A	Set of areas
$\mathcal{A}_t$	Set of actual orders at time step t
$\mathcal{F}_t$	Set of forecast orders at time step t
b_i	Service time required at order i location
β_0	Base utility
β_s	Utility parameter associated with slot s
β_d	Utility parameter associated with delivery price
$C(\vec{x}_t)$	Minimum cost for solving the CVRPTW for feasible delivery schedules
C_{ij}	Cost of travel from customer i to customer j
CAP_s	Maximum order limit for each time slot s
$\overrightarrow{CAP}$	Set of CAP_s for all the time slots
D	Depot
DC_t	Total delivery cost obtained from the dynamic routing system at time t
DPR_s	Delivery price reduction for time slot s
$F_a(\vec{x}_t)$	Set of feasible time slots for area a given accepted orders $\vec{x}_t$
Q	Vehicle capacity
q_i	Demand of order i
r_i	Revenue of the order i under consideration

Variable	Definition
S	Set of time slots
$\theta_{t,s'}$	Monetary value of the consumed time by the new order 1_{as} in time slot s'
T	Cut-off time beyond which customers cannot place further orders
$\vec{d}$	Collection of $\vec{d_a}$ for all areas
$\vec{d_a}$	Vector of delivery prices displayed to customers in area a
d_{as}	Delivery price to area a at time slot s
U_s^B	Start time of the designated time window s
U_s^E	End time of the designated time window s
l_{ij}	Distance between customers i and j
$\mathcal{M}$	Large constant
μ_a	Probability of an arrival coming from area a
λ_t	Arrival rate of customers entering the booking system at time t
n_i	Number of standard time windows in the augmented time window s_i'
OC_t	Estimated opportunity cost at time t
$P_{s,F_a(\vec{x}_t)}(\vec{d_a})$	Probability that a customer chooses slot s when offered delivery prices $\vec{d_a}$ to feasible slots
Φ_{ik}	Calculated delivery time for order i using vehicle k
q_i	Specific demand of order i
$\vec{E}$	Earliest arrival times of all orders
$\vec{E_b}$	Aligned earliest arrival times of all orders
$\vec{f_t}$	Set of remaining forecast orders at time t
$\vec{f}_{rad}$	Set of forecast orders within a radius of rad miles from the new order
$G(x_t)$	Route plan for the system state x_t
$\vec{L}$	Latest arrival times of all orders

Variable	Definition
$\vec{x}_t$	Reshaped matrix $\mathbb{X}(t)$ in column-major order representing the system state at time t
$V_t(\vec{x}_t)$	Value function representing expected maximum potential profit from time step t until final time step T
v	Vehicle speed
$\mathbb{X}(t)$	Matrix at time step t with dimensions $\|A\| \times \|S\|$, indicating accepted order counts for deliveries in areas $a \in A$ and time slots $s \in S$.
X	Set of all states over the entire time steps leading to feasible delivery plans
Z	Monetary value representing the utility reduction equivalent to increasing a time slot's price by £1

Contents

List of Figures

List of Tables

List of Algorithms

1

Introduction

1.1 Background

In recent years, the rapid expansion of Internet and mobile communication networks has led to significant growth in online retailing. This shift has become an important part of grocers' revenue streams, enabling customers to shop for goods online and have them delivered to their homes. As the e-commerce sector grows, including areas such as e-retailers, couriers, and parcel carriers, researchers have focused on developing strategies to address the complex issue of delivery in online businesses, particularly the challenges associated with last-mile delivery.

1.2 Online Grocery Sales Surge

The demand for e-grocery services surged during the pandemic and has remained high. Despite initial predictions of a decline due to COVID-19 control measures, online grocery sales in 2021 increased by 17.1% Mintel (2022). This increase reflects a broader trend. For example, Mintel (2023) reports that the United Kingdom's online grocery market grew by 88% between 2019 and 2021. This trend is also observed in other major European markets, including France, Germany, Spain, and Italy, highlighting the growing importance of e-commerce in retail.

Within this changing landscape, Attended Home Delivery (AHD) has become a significant area of focus for e-grocers Wang et al. (2014). The importance of AHD is due to the nature of grocery orders, which often include perishable and frozen goods that require prompt delivery Agatz et al. (2013).

1.3 The intricacy of Attended Home Delivery (AHD)

The landscape of AHD presents a complex challenge, combining demand management with a form of the vehicle-routing problem. This combination forms a framework for the efficient provision of AHD services. Comprehensive surveys by Fleckenstein et al. (2022) and Waßmuth et al. (2022) highlight how AHD intertwines demand management and vehicle routing. Essentially, it involves dispatching delivery vehicles to customers within predefined time slots, requiring effective geographical and time-window management.

Customer satisfaction stands as a paramount concern in the AHD domain, with delivery wait times directly impacting this critical metric. Consequently, the provision of narrower time slots has become a significant consideration, echoing the findings of

Agatz et al. (2011). However, the pursuit of narrower time slots introduces a noteworthy challenge in optimising delivery routes. Shorter time windows inherently limit the flexibility available for scheduling deliveries, thereby amplifying the intricacies of the route optimisation process.

1.4 Complex Decision-Making

The decision-making process in AHD is inherently multi-faceted, involving two fundamental dimensions: the Capacitated Vehicle Routing Problem with Time Windows (CVRPTW) and Demand Management (DM). These two domains are intricately intertwined, creating a complex challenge that businesses must navigate.

At its core, this challenge revolves around striking a delicate balance between minimising delivery costs for the company and maximising customer convenience. This dynamic lies at the heart of AHD operations, profoundly influencing the strategies employed by businesses to thrive in a dynamic marketplace.

1.4.1 Capacitated Vehicle Routing Problem with Time Windows

The CVRPTW dimension focuses on optimising delivery routes. Companies are tasked with the intricate process of creating routes that are not only cost-efficient in terms of fuel, time, and labour but also adhere to specific time windows. The strict adherence to these time windows is of paramount importance as it ensures that customers receive their orders within the committed time slots. However, this optimisation process is anything but straightforward, particularly when factoring in customers' strong preference for shorter time windows to enhance their convenience. The core challenge here revolves around the efficient allocation of orders to delivery routes while rigorously adhering to these stringent time constraints.

1.4.2 Demand Management (DM)

Simultaneously, companies face the intricate task of demand management. This aspect centres on order acceptance and pricing strategies. To encourage order placement, companies often use competitive delivery fees and various incentives to attract customers. These incentives can include financial discounts, promotions, or appeals that emphasise convenience and reliability.

1.4.3 Integration of Demand Management and Vehicle Routing Problem with Time Windows

The true complexity emerges from the synergy between these two domains. Decisions made regarding delivery routes directly impact the feasibility of accepting new orders. For example, accepting an order in a popular time slot without considering potential routing inefficiencies may lead to a shortage of available slots in the future, affecting overall order acceptance and profitability. On the other hand, rejecting orders too frequently can result in customer dissatisfaction and lost revenue. Additionally, the availability of time slots plays a pivotal role in this dynamic. A greater number of open time slots not only increases the likelihood of customers finding suitable delivery times but also reduces pressure on the company to offer steep discounts to incentivise orders during less popular slots. Managing this slot availability is, therefore, a strategic element in optimising profits within the AHD framework.

In essence, decision-making in the AHD landscape requires a delicate dance between route optimisation, demand management, and slot availability. Achieving the right balance between these facets is the ultimate challenge, necessitating thoughtful consideration of cost-efficiency, customer satisfaction, competitive pricing, and environmental responsibility. The intricate interplay between these elements shapes the strategies employed by companies to thrive in the complex world of AHD.

1.5 Research Questions

This thesis aims to optimise real-time, intelligent delivery plans for attended home delivery (AHD) systems, with a focus on the integration of dynamic pricing, forecasting, and real-time routing adjustments. The research questions presented here are designed to address the challenges and opportunities identified in the comprehensive literature review and the innovative contributions that follow in later chapters. These questions guide the investigation into enhancing the effectiveness and efficiency of AHD systems:

1. **How does the incorporation of forecast orders without time windows enhance the performance of routing and the accuracy of opportunity cost calculations in AHD systems?** This question explores the benefits of integrating forecast orders into routing systems, examining how these considerations improve routing efficiency and opportunity cost assessments, contributing to more effective AHD operations.

2. **How can the distribution of forecast orders among time windows improve the displacement cost calculations for each slot, thereby refining the overall opportunity cost estimations?** This question investigates the impact of strategically distributing forecast orders across different time slots, aiming to enhance the precision of displacement cost calculations and improve the fidelity of opportunity cost estimations.

3. **In what ways can slot availability during the booking period be enhanced by offering discounts to customers in return for accepting extended delivery time window flexibility?** This question delves into the potential of augmented time windows to increase slot availability and provide customers with greater flexibility, which could lead to enhanced customer satisfaction and improved resource utilisation.

4. **How can the inclusion of dynamic pricing mechanisms that alter slot prices enhance opportunity cost calculations by considering the effects of slot price reductions and customer slot selection?** This question focuses on the role of dynamic pricing strategies in refining opportunity cost calculations, especially how these pricing adjustments affect customer behaviour and slot selection, thereby optimising AHD system performance.

5. **How can the estimation of opportunity cost be enhanced by considering the effects of slot price alterations/reductions on remaining slots due to current customer slot selection when using dynamic pricing?** This question examines the influence of dynamic pricing strategies on opportunity cost estimations in AHD systems. It specifically investigates how adjustments in slot prices—alterations or reductions—applied to remaining slots after a customer has made a selection, can improve the accuracy of opportunity cost calculations. This analysis seeks to determine how such pricing adjustments can optimise the allocation of delivery slots, thereby influencing future customer behaviour and enhancing the operational efficiency and attractiveness of delivery options.

These research questions are integral to developing a cohesive and comprehensive understanding of the theoretical and practical enhancements possible in the field of last-mile delivery, particularly for attended home deliveries. They lay a foundational framework for the subsequent chapters, which provide detailed analyses and contributions to these areas.

1.6 Thesis Overview

This thesis presents a series of innovative approaches aimed at improving the integrated demand management and vehicle routing problem inherent in AHD from various perspectives. The structure of the thesis comprises the following chapters:

- Chapter 2 provides a comprehensive review of the existing literature on AHD. The review covers the historical and contemporary landscape of AHD research, as well as the challenges and complexities that continue to shape this field. The chapter also identifies the gaps and limitations that persist in current AHD solutions.
- Chapter 3 introduces the concept of Dynamic Routing with Forecast Orders. The chapter starts by constructing and optimising tentative routes for forecast orders devoid of time constraints. As actual orders arrive, the forecast orders are substituted with actual orders to gauge the opportunity cost for each time slot. Dynamic pricing then takes the reins to influence customer slot selection. Depending on the customer's choice, the actual order steps in, and the route map undergoes re-optimisation. This methodology results in more efficient scheduling and increased slot availability for incoming customers.
- Chapter 4 focuses on the estimation of Opportunity Cost and Slot Management. This chapter presents two crucial contributions. Firstly, it proposes an enhanced method for estimating opportunity cost by leveraging a dynamic routing and distribution approach integrating forecast orders. This refined approach facilitates more precise displacement cost assessments, thereby informing decisions on deliv-

ery charges more effectively. Secondly, it introduces a dynamic slot-combination strategy aimed at fully harnessing customer flexibility in receiving their deliveries. This, in turn, elevates overall route efficiency and customer satisfaction.

- Chapter 5 discusses the Integration of Opportunity Cost and Delivery Price Reduction. Enhancing the estimation of opportunity cost involves adjusting delivery prices for future time slots based on the reduction in delivery prices resulting from slot allocation at the present time. This adjustment aims to optimise pricing strategies to better reflect customer preferences and route efficiency.

- Chapter 6 concludes the thesis and explores avenues for future research. This chapter concludes the thesis by summarising the key findings and contributions. It also discusses the implications of the research for businesses and practitioners in the field of AHD. Additionally, it identifies and explores potential avenues for future research, providing a foundation for further exploration and development in this evolving domain.

This thesis makes a number of notable contributions to the field of AHD. The proposed approaches have been evaluated using a comprehensive set of computational experiments. The results of these experiments demonstrate the effectiveness of the proposed approaches in terms of profitability, delivery efficiency, and customer satisfaction. The research presented in this thesis has the potential to considerably impact the field of AHD. The proposed approaches can be used by businesses to improve the efficiency, profitability, and customer satisfaction of their AHD operations. The research also provides valuable insights for researchers and practitioners in the field of last-mile logistics.

2

Background

2.1 Introduction

Last-mile delivery involves transporting goods from a distribution centre to an end user's location. It is the final step in the delivery of goods, which is usually the most expensive portion of the process. The goal of last-mile delivery logistics is to ensure that demands are delivered in a timely, accurate, and cost-effective fashion. There are two major classes of problems associated with last-mile delivery, namely demand management and vehicle routing. Therefore, a satisfactory solution to last-mile delivery will be achieved if the objectives of both problems are taken into account. The following sections review the existing e-fulfilment literature, particularly attended home deliveries (AHD), which fits the direction of our research. For a broader overview of revenue management, the reader is advised to refer to studies by Agatz et al. (2013), Agatz et al. (2008) and Klein et al. (2020). Recently, Snoeck et al. (2020), Waßmuth et al. (2022) review the researches in last-mile delivery.

The structure of this chapter is outlined as follows: Section 2.2 reviews the literature pertinent to the demand management problem in Automated Home Delivery (AHD). Section 2.3 is dedicated to examining the development of customer choice models as presented in various studies. Following this, Section 2.4 discusses different methodologies for approximating opportunity costs, including insights into marginal delivery cost estimates, displacement costs, and delivery price reductions. Section 2.5 explores the management of different time window systems in AHD, as documented in the literature. Finally, Section 2.6 provides a concise summary of the topics covered in this chapter.

2.2 Demand Management in AHD

Following the standard categorisation of demand management in time-slotted deliveries by Agatz et al. (2008), literature can be grouped into four main groups: static slotting (e.g. Agatz et al. (2011)), dynamic slotting (e.g. Campbell and Savelsbergh (2005)), static pricing (e.g. Klein et al. (2017)) and dynamic pricing (e.g. Campbell and Savelsbergh (2006)). Slotting focuses on time-slot allocation to customer regions. In contrast, pricing assigns delivery prices to balance demands across time slots and/or steer customer choices towards the best delivery-time options. Amongst static and dynamic approaches, the latter attracts more attention as it reflects the nature of online booking. Readers are referred to the surveys by Klein et al. (2020) and Snoeck et al. (2020) on the industrial application of revenue management and advances in choice-based models for more information. This thesis concentrates on dynamic approaches to demand management with pricing.

In attended home delivery, requests appear progressively over time during the booking horizon until a cut-off time when no further requests can be accepted, as in the works of Yang et al. (2016) and Yang and Strauss (2017). Each request chooses a certain time of the delivery day if the desired delivery time is available. Various requests arise from different locations and times, and the goal is to fulfil as many of them as possible. In other words, the delivery sequence should be operationally feasible while maintaining a low total delivery cost. Dynamic routing systems play an important role in such a framework to accommodate stochastic arrival of customers in the chosen time windows while minimising the total delivery cost. From a demand management perspective, the number of accepted requests should be maximised by means of pricing policies on the delivery fees to increase the yield of delivery services. To ensure that the final route plan is feasible, the dynamic routing system must continuously verify the feasibility of the route plan for each potential customer.

Another crucial aspect of demand management in AHD is the specification of delivery fees for the available time slots. Pricing involves assigning delivery prices to balance demands across time slots and guide customer choices towards optimal delivery time options. Under a dynamic pricing strategy, the delivery prices offered can exhibit flexibility, spanning a continuous range, as demonstrated by Yang et al. (2016), or they can be selected from a predetermined set of discrete price points, as showcased in Koch and Klein (2020). This adaptability in pricing caters to various operational and market circumstances, further solidifying dynamic pricing as the favoured approach for optimising attended home delivery processes.

In contrast to static pricing models, which generally assign uniform prices to all time slots, dynamic pricing distinguishes itself by offering varied prices for different slots, aiming to maximise company profits. This method is particularly well-suited to the nature of online booking systems, where customer delivery requests are placed and adjusted dynamically throughout the booking horizon, reflecting the fluctuating nature of demand and availability.

2.3 Customer Choice Model

From the grocer's point of view, delivering an order in a time slot or another typically yields different costs, motivating the company to steer customers' selection of time slots to increase its profit. One can apply many promoting strategies to achieve this aim, such as offering discounted delivery prices/discount vouchers for some slots, highlighting the slot in different colours to reflect their pollution/environmental impacts, etc. In this context, the work of Campbell and Savelsbergh (2006) stands out as an early pioneer in employing a relatively straightforward customer behaviour model. Their objective was to grasp how incentives, such as delivery charges, influenced the probability of customers choosing specific time slots. Their emphasis centred on steering slot preferences to lower delivery costs, rather than solely optimising anticipated profits, which resonates with the central theme of our study. To gain deeper insights into customer behaviour, Asdemir et al. (2009) employed a sophisticated approach known as the multinomial logit (MNL) model, designed to address the complexities of demand management in attended home deliveries. This model assigns varying levels of utility to each delivery option, reflecting the assumption that customers make choices to maximise their perceived value. The framework developed by Asdemir et al. (2009) incorporates a dynamic programming approach with fixed delivery costs, however, it does not include dynamic pricing considerations essential for addressing fluctuating time slot demands. Furthermore, the dynamic programming model they propose has a state space that expands exponentially with the number of delivery slots, rendering it impractical for large-scale applications. Specifically, the model aims to maximise expected profits without accounting for delivery routing, presenting challenges in real-world scenarios where delivery fees significantly impact demand for time slots and thus affect routing schedules and costs. Additionally, the model assumes geographic independence, which simplifies its application to smaller areas. For greater practicality, it is necessary to integrate routing costs and account for demand variations across neighbouring areas in the routing plans.

In their research, Asdemir et al. (2009) combined the MNL model with a dynamic pricing framework to refine the incentive scheme for delivery fees. Employing dynamic programming with predefined delivery costs, their method aimed to solve the dynamic time slot pricing problem. However, this approach is limited by its assumptions that the delivery capacity for each time slot within a region is known in advance, which restricts its utility to smaller, less complex scenarios. As a result, while the MNL model provides a robust framework for analysing customer choice behaviour, its application in larger or more dynamic settings is constrained by the rapid growth of the state space in the dynamic programming model, necessitating simplified operational environments for its effective use. Mackert et al. (2019) proposed a new choice model based on a finite-mixture MNL choice model that can simplify the nonlinear optimisation model to a linear one for solving. A model-based, profit-oriented slotting approach is developed to accurately approximate customer choice behaviour.

One noteworthy advantage of dynamic pricing is its potential to steer customer slot choices toward those that benefit the routing system, enabling accommodation of more orders and ultimately increasing profits by the end of the booking period. This advantage is contingent on simulating customer responses to the varying prices of delivery slots throughout the delivery day. Consequently, the incorporation of a customer choice model becomes pivotal in simulating customer selection behaviour and adjusting prices accordingly.

The research conducted by Yang et al. (2016) delves into the realm of dynamic pricing within the context of an MNL choice model. This model was trained using a real data-set sourced from previous customer booking data in the UK. The study aimed to understand customer willingness to pay for various displayed delivery time slots and to analyse the probability of selecting specific time slots from the available options. The model, built on simulations, successfully replicated customer choice behaviour. The parameters used were then employed to shape and guide the dynamic pricing strategy. They employ insertion heuristics to update a pool of feasible routes as orders come in over the booking horizon and deploy the marginal delivery cost as estimates of the opportunity cost of accepting an order into a particular time slot. The proposed "foresight" approach, which uses previously planned routes to bring in the effects of forecast orders, is justified as superior to the "hindsight" approach, which only considers accepted orders. While using forecast orders based on previously planned routes can help to build better routes, their method is restricted compared to our approach for the reason that such forecasts are not updated with the acceptance of new orders, and the routing for actual orders is independent of that for forecast orders.

In this thesis, the benefits of the MNL choice model established by Yang et al. (2016) are leveraged. Building on this framework, several studies, including Yang and Strauss (2017) and Strauss et al. (2020), have incorporated the MNL choice model in conjunction with dynamic pricing. This approach depends on high-quality inputs that forecast future customer behaviours, enabling price adjustments aimed at increasing the probability of time slot selection as predicted by the MNL choice model. Opportunity cost, driven by the potential revenue from future customers, plays a pivotal role in this dynamic pricing strategy. The next section will discuss various methods for estimating the opportunity cost.

2.4 Opportunity Cost Approximation

Opportunity cost, a fundamental concept in economics and decision-making, assumes a pivotal role when resources are limited. In the context of Attended Home Deliveries (AHD), the opportunity cost of a delivery slot encompasses both the insertion cost of a new actual order into the route plan and the potential future order displacement cost. This displacement cost arises when the current order occupies the limited capacity for a particular time slot, thereby preventing the accommodation of potential future orders that may have higher value or priority. Given the large size and stochastic nature of industrial applications in e-fulfilment, the opportunity cost cannot be calculated precisely. Approaches such as Approximate Dynamic Programming (ADP), Linear Approximations, and look-ahead heuristics have been deployed to tackle the computational difficulties.

In detail, Figliozzi et al. (2007) model the carrier-pricing problem in the dynamic vehicle routing environment as a stochastic dynamic program, which is solved through a one-step look-ahead heuristic. Since the context is in-freight transportation, the approaches proposed by Figliozzi et al. (2007) are not readily applicable to AHD problems. In the AHD context, Klein et al. (2018) present an approximation approach based on a Mixed-Integer Linear Programming (MILP) reformulation to approximate opportunity costs. However, the MILP suffers from computational challenges; even with further simplifications and parallel computing, their approach has not proven suitable for scenarios with more than 15 vans.

Simulation-based approaches and predictive models, as exemplified by studies such as Yang and Strauss (2017) and Ulmer (2020) are other ways to find the opportunity cost. To solve the computational difficulty, Yang and Strauss (2017) exploit a continuous approximation of the delivery costs and propose an ADP method that estimates the opportunity cost in real-time. Their approach is justified as efficient in industrial-size implementations; however, the routing approach used omits many practical restrictions to achieve this aim such as the use of opportunity-cost estimates that depend neither on time of booking within the booking horizon nor on the level of orders accepted. In contrast, this thesis proposes an approach that can directly deploy the routing package a van is currently using, which ensures that all practical restrictions of routing are accommodated in the online pricing decisions.

The study conducted by Koch and Klein (2020) introduces an innovative approach that incorporates a tentative route plan to identify feasible time slots and estimate the two pivotal components of opportunity cost. Their method draws inspiration from studies like Campbell and Savelsbergh (2006) and Yang et al. (2016), where a route plan is meticulously constructed based on two distinct route maps. The initial route commences as an empty slate and evolves incrementally with the arrival of new requests, employing a myopic approach focused on minimising insertion costs. Additionally, they employ a tentative route map, often referred to as a skeletal route, populated with artificial requests that simulate future customers. This forward-looking tentative route plan is crafted through the simulation of multiple booking periods, with each subsequent simulation employing the previously created route map as a foundation. Furthermore, they introduce the concept of a time window time budget to quantify the idle time within each time window, subsequently utilising this metric to calculate the value of each time window as an estimate of the displacement cost. A fundamental distinction between their problem setting and ours lies in their assumption of a known probability distribution of requests.

Unlike all previous works employing forecast orders, including Koch and Klein (2020) and Yang et al. (2016), the research presented in Chapter 3 does not employ the "previous route" for forecasts. This is because the "previous route" was constructed based on orders received under a fixed-demand scheduling policy, such as fixed pricing or fixed order acceptance. These works consider previously allocated time windows when constructing the forecast route, implicitly assuming that the previously allocated time window was optimal (or sufficient) and that repeating it (or guiding the system to reinforce it) will lead to preferable solutions. However, this research avoids using previous routes and their time-window restrictions imposed on forecast orders. Instead, it starts with the *optimal* route (or the best route one can find with a heuristic) *without* imposing any time windows to mimic the best possible route. This is achieved by adopting an "excellent" pricing policy in an ideal scenario, where all customers select what turns out to be the optimal time slot (from the route-optimisation software's perspective) to receive their orders. Although this strategy may seem overly optimistic initially, it is adapted over time, gradually incorporating more actual information as orders arrive, and re-optimising the routing upon every committed order with a fixed, known time window. This approach does not require forecasting time windows but only the number of orders and their locations. As it updates according to actual order arrivals, it remains robust concerning forecast errors. For detailed experiments and results on shifted forecast levels, please refer to Section 3.4.4.

With regard to dynamic routing with forecast orders, but without pricing to influence time-slot choices, our approach in Chapter 3 shares certain similarities with Bent and Van Hentenryck (2004), Ichoua et al. (2006), and Voccia et al. (2019). However, these works assume known probability distributions of future demands when generating future delivery requests, which, in our case, is unknown because the distribution is influenced by the dynamically-changing pricing policy our system generates. Recently, Soeffker et al. (2022) comprehensively discuss scenario-based approaches in which information models are integrated to tackle stochastic dynamic vehicle routing. Also, Klein and Steinhardt (2022) extend scenario-based approaches to address same-day deliveries and incorporate value function approximation approaches as support to include the dynamism over time to accomplish anticipatory decision-making. Likewise, Ulmer (2020) uses a value function approximation approach to find the best pricing policy in the same-day delivery problem, which proves effective. However, if these approaches rely on a lookup table, such as the variant approach in Ulmer (2020), they suffer from the curse of dimensionality as the size of the vehicle fleet increases.

Additionally, other articles discussing AHD resolutions under different problem settings exist, such as Dayarian and Savelsbergh (2020) exploring the employment of in-store customers to deliver online orders while they return home, Strauss et al. (2020); Köhler et al. (2020a) working on the use of flexible time slots and Agatz et al. (2021) focusing on the effect of green labels. Further, Ojeda Rios et al. (2021) have recently published a survey on Dynamic Vehicle Routing Problems (DVRP) by providing taxonomies of the problems and solution methods. They reported that heuristics and meta-heuristics provided 66% of solutions to DVRP. In this thesis DVRP will be used to estimate opportunity costs and guide choices of time slots.

2.4.1 Marginal Delivery Cost Estimation

Estimating the marginal delivery cost through insertion heuristics applied to tentative routes is an established approach. Pioneers like Campbell and Savelsbergh (2005) integrated dynamic vehicle routing with the scheduling problem for home delivery services. They conducted feasibility checks on tentative routes and approximated opportunity costs based on marginal delivery costs using insertion heuristics. When feasible, they evaluated whether to accept or reject incoming requests based on the associated profit or cost. Further exploration into diverse applications of insertion heuristics in vehicle routing problems is provided by Liu et al. (2023).

The estimated marginal delivery cost did not contribute to the decision of e-Retailers to incentivise customers by Campbell and Savelsbergh (2005). Customer behaviour models are used to determine how likely a time slot is to be selected by a customer. As a result of incorporating this information into the firm's decision-making process on delivery prices, they can examine an adaptive incentive scheme to display the price of slots to the customers. Campbell and Savelsbergh (2006) propose the first simple model of customer behaviour in a linear format to reduce delivery costs by driving customers to time slots with lower travel costs. To determine the feasibility of time slots in these studies, insertion heuristics have been used to estimate marginal delivery costs by identifying the slot with the lowest insertion costs (e.g., the works by Campbell and Savelsbergh (2006), Klein et al. (2018), Yang et al. (2016), and Köhler et al. (2020b)). In addition, Bühler et al. (2016) present another approach for calculating the cost of delivery for a fixed pool of feasible routes using linear mixed-integer algorithms. In recent studies, tentative route plans have proven to be even more relevant because they can effectively contribute to estimating the displacement cost of future demands. In the literature, different mechanisms for generating and maintaining tentative route plans have been proposed.

Instead of using marginal delivery costs as estimates of the opportunity costs and focusing on tentative routes and insertion heuristics, there is a different strategy in other research. They use alternative modelling approaches to simplify the underlying VRPTW solutions to emphasise the potential revenue loss by occupying the slot capacity. Most of these studies investigate the "acceptance scheme" of customer requests, such as Ehmke and Campbell (2014) and Cleophas and Ehmke (2014), which aims to maximise the number of requests accepted for delivery.

2.4.2 Displacement Cost (Revenue Loss)

Displacement cost quantifies the monetary value associated with integrating a new request into an existing delivery route. In the literature as proposed by Klein and Steinhardt (2022), two primary methods are identified for calculating this cost:

- **Learning-Based Approaches:** These methods approximate opportunity cost through simulations learning from different system states, incorporating factors like marginal delivery cost. Value Function Approximations (VFAs) are commonly used, as shown in studies by Ulmer (2020) and Yang and Strauss (2017). VFAs tackle the dynamic programming complexities in AHD, addressing the exponential growth of the state space in vehicle routing problems with uncertain requests. Yang and Strauss (2017) notably employ an offline methodology with approximate dynamic programming to optimise request distribution within delivery areas. These methods can operate either offline or in real-time during booking but often face the challenge of the "curse of dimensionality," particularly when using look-up tables for state value determination.

- **Non-Learning-Based Approaches:** This category includes methods that calculate opportunity cost components using current or tentative route plans, derived from scenarios as by Bent and Van Hentenryck (2004) or historical data, as in Yang et al. (2016). Recent research like Klein et al. (2018), Mackert et al. (2019), and Strauss et al. (2020) incorporates future customer behaviour and final delivery charges into opportunity cost calculations. These methods are particularly useful in stochastic dynamic VRPs with uncertain requests and are notable for integrating dynamic pricing and choice models like the MNL, as evidenced in works by Koch and Klein (2020) and Abdollahi et al. (2023). They focus on the interplay between demand management and vehicle routing, using forecast route plans for estimating opportunity costs, especially in scenarios with unknown future demand distributions.

In this thesis, displacement cost is addressed using the forecast route plan as a proxy for the final route. Chapter 3 introduces a heuristic-based method for creating an optimal route plan for forecast orders without time windows, which includes assigning time windows for feasibility and identifying cost-effective insertion points for new requests. This approach also involves allocating time windows to forecast orders to reflect the count of orders per time window when calculating displacement costs.

Moreover, as delineated in Chapter 4, forecast orders are allocated based on a metric that sets the upper limit of orders per time window. This metric, derived from simulations to determine the maximum feasible orders per window, allows for a more accurate estimation of displacement costs. This new methodology, which assigns time windows based on earliest availability and regulates order distribution within each window, offers a nuanced approach that aligns more closely with customer preferences and is adaptable to varying pricing policies. Further specifics of this method are discussed in Section 4.3.3.

The researches conducted by [Yang and Strauss (2017), Klein et al. (2018), Mackert et al. (2019), and Strauss et al. (2020)] similarly factor in future customers and final delivery charges when estimating opportunity costs. Yang and Strauss (2017) employ an offline methodology through approximate dynamic programming to estimate the opportunity cost. This involves dividing the delivery area served by a single van into segments and determining the highest count of requests within each sub-area.

2.4.3 Delivery Price Reduction

The framework for estimating opportunity cost, as outlined by Klein and Steinhardt (2022), guides the approach to delivery price reduction in Attended Home Deliveries (AHD). Chapter 5 introduces an novel method for estimating opportunity costs related to price adjustments for future time slots. This method, grounded in non-learning-based strategies, accurately quantifies the reduction in slot prices necessitated by the integration of new orders. It leverages the principles of dynamic pricing and the MNL model, as discussed in the works of Yang et al. (2016), Koch and Klein (2020), and Abdollahi et al. (2023), enhancing the precision of opportunity cost calculations in dynamic pricing scenarios.

The methodology employed in this thesis estimates potential revenue loss from pricing adjustments, factoring in customer utility levels and preferences. This approach is crucial for dynamic pricing strategies that optimise time slot allocation based on booking behaviours. By quantifying real-time price reductions for each time slot, the calculation of marginal insertion and displacement costs is enhanced.

2.5 Time Window Design and Management

In standard Attended Home Deliveries (AHD), customers are typically offered one-hour time slots, each tailored to individual factors such as location, order size, and existing bookings. This customisation allows e-retailers to adapt their slotting process efficiently. However, as Lin and Mahmassani (2002) and Punakivi and Saranen (2001) highlight, focusing too heavily on reducing waiting times can lead to higher delivery charges and potential order rejections. This is because it necessitates the use of more vehicles and staff, mandates narrower delivery windows, and consequently offers less flexibility to customers. Conversely, Gevaers et al. (2014) demonstrate that longer time windows can enhance e-retailer profitability, citing significant cost differences between one-hour and four-hour delivery windows.

The concept of introducing flexibility into time window management has garnered attention in the existing literature, from the perspectives of both the firm and the customer. Strauss et al. (2020) propose a method that offers customers increased flexibility in selecting their preferred time windows, allowing them to choose from a range of standard time windows that align with their preferences. Once the delivery vehicle is dispatched, customers are promptly informed of the specific delivery window. This flexibility could potentially be coupled with reduced delivery charges and a focus on environmental factors as a form of compensation. In Chapter 4 of this thesis, a novel approach is introduced to enhance flexibility in time window management by incorporating standard and extended time windows with longer durations, while offering discounts to improve slot availability and motivate customers to place their orders.

The study in Chapter 4 aligns with the strategies wherein e-retailers introduce flexibility into time window management by offering both extended and standard time windows. Köhler et al. (2020b) propose a strategy where new customers are presented with a selection of short and/or long time windows based on factors such as travel time, current booking horizon, insertion time, and the time span associated with order insertion. While their approach assumes knowledge of future demands, the context in this thesis lacks detailed information about customer distribution. Additionally, their model offers fixed delivery fees for both short and long windows, which may not align with potential customer preferences for lower charges in exchange for increased flexibility. In contrast, this research incorporates both types of time windows to cater to diverse customer needs. The dynamic pricing strategy adopted here, similar to the one discussed by Strauss et al. (2020), involves tailored delivery prices derived from current and anticipated customer preferences, incentivising or emphasising the environmental benefits of selecting extended time windows. Moreover, the inclusion of varying time window lengths is intended to enhance the likelihood of bookings, ultimately aiming to increase acceptance rates by the conclusion of the process.

2.6 Conclusion

In conclusion, the literature review has highlighted the key factors and approaches in the field of last-mile delivery, particularly in attended home deliveries (AHD). It has discussed the importance of demand management and dynamic pricing in optimising AHD operations. Moreover, it has explored various models, including customer choice models and marginal delivery cost estimation, that are essential in implementing dynamic pricing strategies. Additionally, the review has emphasised the significance of accurate opportunity cost estimation and effective time window design and management in enhancing the efficiency and profitability of last-mile delivery. The insights gained from this literature review will provide a solid foundation for the research conducted in this thesis.

Despite these comprehensive insights, several research gaps remain apparent. Firstly, there is a need for enhanced models that integrate dynamic pricing with real-time routing adjustments to better reflect the stochastic nature of customer bookings and demand fluctuations. Additionally, the existing models often simplify geographical and temporal dependencies, which can limit their applicability in dense urban environments or complex logistical networks. Another important gap lies in the opportunity cost estimation methods which, while advanced, still struggle with computational efficiency and scalability in larger operational settings. Further research should also explore the inter-dependencies between dynamic pricing, customer choice behaviour, and their cumulative effects on routing efficiency and overall system profitability. Addressing these gaps will not only refine the theoretical models but also enhance their practical applicability in diverse last-mile delivery contexts.

3

Demand Management in Time-slotted Last-mile Delivery via Dynamic Routing with Forecast Orders

This chapter is reproduced with some changes from Abdollahi, M., Yang, X., Nasri, M. I., and Fairbank, M. (2023). Demand management in time-slotted last-mile delivery via dynamic routing with forecast orders. European Journal of Operational Research.

3.1 Introduction

Nowadays, most online grocers use fixed, one-hour time slots. Naturally, time slots receive different customer preferences. Suppose nothing is done through the booking horizon. In that case, the company will likely have to face unbalanced demand for time slots, which leads to inefficient delivery routes and a potential waste of fleet capacity. To deal with this issue, researchers developed demand-management strategies in the past and recent literature aiming to control/incentivise customers to select specific time slots to ease the routing difficulty, including Agatz et al. (2013); Asdemir et al. (2009); Campbell and Savelsbergh (2006); Cleophas and Ehmke (2014); Yang et al. (2016); Koch and Klein (2020). Which slot price vector is the best is the key question to answer in this context, which has not yet been fully addressed by the existing literature. Difficulties in answering this question lie in two facts:

1. As shown in Figure 3.1, the best slot for satisfying an order depends on the insertion cost for placing that order into the *final* delivery route and the potential revenue loss of displacing another order. However, the full order list is not known during the booking horizon when decisions on which slot to promote have to be made. Therefore, a forecast is needed for the final delivery route so that when the booking is made, the extra cost caused by the likely route deviation to satisfy that new order can be estimated reasonably.

2. The likelihood of a customer accepting an order and its time slot depends on the promotional decisions made during the booking horizon, so the final delivery route is not entirely predictable from historical information where, essentially, a different pricing/order-acceptance approach was used.

There is a mutual dependency between the forecast delivery route and the incentive decision. This work proposes a methodology that employs dynamic routing to disrupt this cycle, enabling accurate demand forecasting while preserving the interdependence between variables to the greatest extent possible.

Past literature in AHD concentrates primarily on how to deal with time windows, both in the demand-management step (via pricing or other incentive means) and in solving the so-called Capacitated Vehicle Routing Problems with Time Windows (CVRPTW) affirmed by Kumar and Panneerselvam (2012). While Yang et al. (2016) and Koch and Klein (2020) suggest route-based approaches by which a virtual route map with time windows is created to predict the final route map to help the CVRPTW find a more suitable time window for servicing an order, an obvious fact is that a more appropriate time window can be found by solving the VRP *without* time windows. Inspired by this idea, in this study, we propose a novel dynamic routing and pricing approach based on solving and updating a partially time-windowed CVRP with a combination of real and forecast orders. Specifically, at any time during the booking horizon, we maintain a set of already-committed orders, with fixed known time windows, locations and order sizes and a set of forecast orders, with given order sizes and locations but without any time windows. We make this distinction because the locations and sizes of forecast orders can be more reliably concluded from historical data

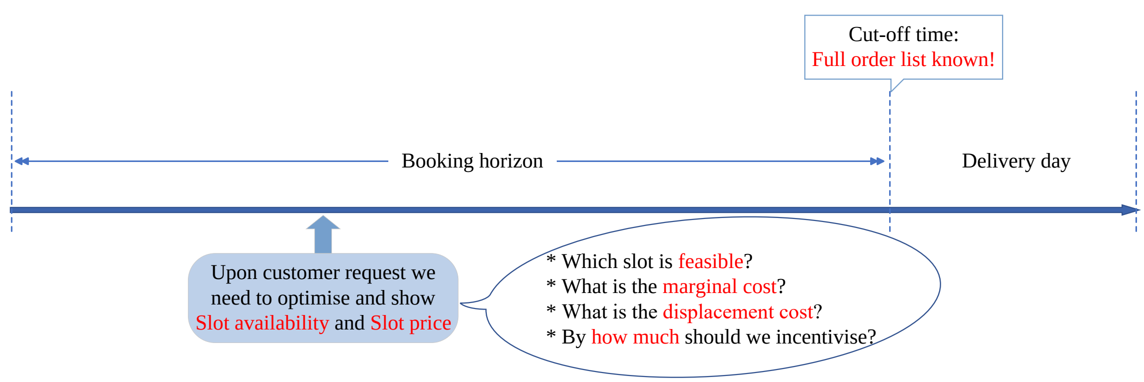

Figure 3.1: Time frame

(than delivery slots), as they are not meant to be influenced by the incentive policy. As for the delivery-slot choices of forecast orders, however, we do not impose any time windows. Instead, we allow the dynamic routing approach to choose its preferred time slot for each order, which is then offered to customers following the incentive policy. Essentially, this approach assumes that future customers will generally choose the incentivised time slots to receive their deliveries, which aligns with our ultimate aim of developing and deploying an incentive policy to steer customer choices of time slots. In detail, at any point in time, we solve a partial CVRPTW (p-CVRPTW) made up of accepted orders with time windows and forecast orders without time-windows. We feed back more suitable time window for satisfying the forthcoming customer based on the solution of the p-CVRPTW. This p-CVRPTW is solved dynamically online and updated whenever new orders are committed.

The major contributions of this chapter are:

- For the first time, incorporating forecast orders *without time windows* into the vehicle-routing system to allow the p-CVRPTW to suggest a more appropriate time slot to accommodate every forecast order and guide the choice of incoming orders accordingly;

- Proposing a simple-to-implement dynamic opportunity-cost approximation for marginal delivery cost and potential revenue loss, based on the dynamically managed routing system with both actually accepted orders and forecast orders without time windows;

- Presenting an order-replacement and routing re-optimisation framework to capture the influence of new order commitments and facilitate opportunity-cost approximation, which evolves as more information becomes available;

- Presenting an approach that is capable of incorporating the firm's specific routing method, which may include considerations such as clustering, shifting, traffic prediction, etc., to the maximum extent;

- Demonstrating the superiority of the developed approach over four benchmark approaches on real data-sets taken from four typical geographical and demographic settings;
- Investigating the trade-off between responding time and the online decision process accuracy.

The chapter is organised as follows: Section 3.2 explains different aspects of the AHD problem and its dynamic programming model. Section 3.3, presents our methodology and how to incorporate forecast orders, pricing optimisation and the customer-behaviour model. The experiment settings and results obtained are reported in Section 3.4. Finally, we conclude in Section 3.5.

3.2 Problem Specification

For the problem under consideration, the company manages an online booking system that allows customers to book their delivery a couple of days in advance, which we refer to as the booking horizon. The orders committed during the booking horizon have to be delivered to the customer's front door during the agreed time slot by the company, using its fleet. Time slots are predefined by the company which may overlap. A scheme of the slot-booking process is shown in Figure 3.2. To purchase goods and book for delivery, a customer has to log in to their account with the grocer, which allows the system to identify their address. We refer to this as a "customer arrival". Next, assuming the customer decides to place an order, the customer chooses their delivery day. This step may happen before or after filling their shopping basket. Note that once the truck has been loaded and dispatched, it does not need to come back to the depot to collect more orders until its planned route has finished. After the customer's selection of a delivery day, the company has to identify in real-time all the feasible time slots which could be used to service this order, together with their incentive scales and additional cost of delivering the order. Based on this information, the customer chooses a time slot, finishing the order commitment. Decisions in this problem must be made in a stochastic dynamic environment, with randomness coming from both customer arrival and customer selection of delivery slots. The requirement for a fast response time, between a customer's click of delivery day and the display of available slots and prices, adds another layer of difficulty to the problem.

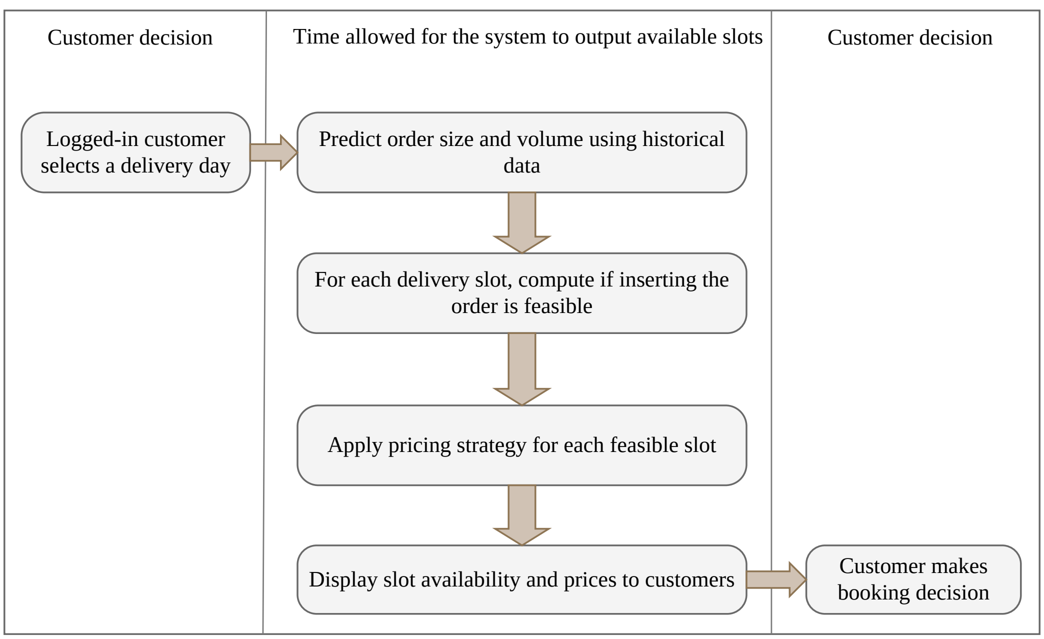

Figure 3.2: Slot booking process

3.2.1 Dynamic Programming Model

In this work, we inherit the Markov Decision Process (MDP) model formulated by Yang et al. (2016). We consider a discretized booking horizon with T periods, by which we mean customer arrivals to the website during the booking horizon shown in Figure3.1. Each booking period is sufficiently small such that the probability of having more than one arrival of a booking request is negligible. The final time period T denotes the cut-off time after which no further bookings are taken. The stages of the dynamic program are the time periods $t \in \{1, 2, \ldots, T\}$. At time step t within the booking horizon, the system's state can be described by a matrix $X(t)$, with $|A|$ rows and $|S|$ columns. The $[a, s]$th component of $X(t)$ represents the number of orders accepted up to time t (in the booking horizon), to be delivered to area $a \in A$ in time slot $s \in S$. In what follows, we use $\vec{x}_t$ to denote the matrix $X(t)$ reshaped in column-major order to be stored in a one-dimensional array.

Let $V_t(\vec{x}_t)$ denote the value function at stage t and state $\vec{x}_t$; it represents the expected maximum profit [1] obtainable from the sales process from time t until the cut-off time T. The dynamic-programming recursion at stage $t \in \{1, 2, \ldots, T\}$ is:

$$\begin{aligned}
V_t(\vec{x}_t) = \max_{\vec{d}} \Bigg[& \Bigg(\sum_a \lambda \mu_a \sum_{s \in F_a(\vec{x}_t)} P_{s, F_a(\vec{x}_t)}(\vec{d}_a) \Big[r_i + d_{as} + V_{t+1}(\vec{x}_t + \mathbf{1}_{as}) \Big] \Bigg) + \\
& \Big[1 - \sum_a \lambda \mu_a \sum_{s \in F_a(\vec{x}_t)} P_{s, F_a(\vec{x}_t)}(\vec{d}_a) \Big] V_{t+1}(\vec{x}_t) \Bigg] \\
= & \Bigg[\max_{\vec{d}} \sum_a \lambda \mu_a \sum_{s \in F_a(\vec{x}_t)} P_{s, F_a(\vec{x}_t)}(\vec{d}_a) \Big[r_i + d_{as} - \big(V_{t+1}(\vec{x}_t) - V_{t+1}(\vec{x}_t + \mathbf{1}_{as}) \big) \Big] \Bigg] \\
& + V_{t+1}(\vec{x}_t), \\
\forall\, \vec{x}_t \in \mathcal{X}, &
\end{aligned} \tag{3.2.1}$$

where λ indicates the arrival rate of customer requests; μ_a denotes the probability that the arrival comes from area a for a given customer arrival; $\vec{d}_a$ is a vector of length $|S|$

[1] In this context, "profit" specifically refers to the total shopping revenue plus delivery charges minus the cost of delivery, which is described here as the "Net Delivery Margin." This distinction is important as the traditional definition of profit generally encompasses a wider range of costs and revenues. The term "Net Delivery Margin" is used to specifically highlight the financial outcomes directly associated with delivery operations, distinguishing it from broader profitability measures.

specific to area a, with components d_{as}, where component d_{as} represents the delivery price to area a at time slot s; $\vec{d}$ is a collection of $\vec{d_a}$ over all areas; $F_a(\vec{x}_t) := \{s : C(\vec{x}_t + \mathbf{1}_{as}) < \infty\}$ denotes all feasible time slots for area a, into which order (a, s) can be feasibly inserted given orders $\vec{x}_t$ that have been accepted, where $\mathbf{1}_{as}$ is the unit vector equal to the flattened single-entry matrix with a 1 in position (a, s); $P_{s,F_a(\vec{x}_t)}(\vec{d_a})$ denotes the probability that a customer chooses slot s when the firm offers the vector of delivery prices $\vec{d_a}$ to feasible slots in $F_a(\vec{x}_t)$; r_i denotes the revenue of the order i that is under consideration. The boundary condition for the MDP model is given by:

$$V_{T+1}(\vec{x}_{T+1}) = -C(\vec{x}_{T+1}) \quad \forall\, \vec{x}_{T+1} \in \mathcal{X}, \tag{3.2.2}$$

where $C(\vec{x}_t)$ represents the minimum cost of servicing all accepted orders during their agreed time slots, which is the optimal solution of a Capacitated Vehicle-Routing Problem with Time Windows (CVRPTW); $\mathcal{X}$ denotes the set of all states that allow a feasible delivery schedule. If there is no feasible solution for a given $\vec{x}_{T+1}$, $C(\vec{x}_{T+1}) := +\infty$.

The dynamic program is intractable due to the large state space and the fact that the optimal solution of large-scale CVRPTW alone is intractable. Nevertheless, the formula (3.2.1) shows that the time-slot pricing decision is a trade-off between the immediate income, $(r_i + d_{as})$, and the expected opportunity cost $(V_{t+1}(\vec{x}_t) - V_{t+1}(\vec{x}_t + \mathbf{1}_{as}))$ arising from reserving the delivery capacity in (a, s) at time t for a future order. Suppose the opportunity cost can be estimated, then the problem can be divided into single-stage decision problems and becomes tractable. There are two major components of the opportunity cost:

1. the marginal delivery cost of servicing one more order in (a, s), and
2. the potential revenue loss from filling fleet capacity at t with an order in (a, s).

Both of these two terms depend on the final delivery routes.

This study aims to estimate the opportunity cost via dynamic routing with forecast orders. Both marginal delivery costs and potential revenue loss will be estimated by solving a CVRPTW dynamically over the booking horizon, with a set of already-accepted orders with fixed time windows and a set of forecast orders with relaxed time windows. More details about the approximation are discussed in Section 3.3.2.

3.3 Methodology

This section presents the solution methodology for the stochastic dynamic-pricing problem (3.2.1), which is intractable via backward induction. In detail, we will address how forecast orders are generated, integrated and updated in the dynamic setting and used to estimate opportunity cost in the following sub-sections.

3.3.1 Forecast orders

As explained, we aim to incorporate forecast orders into the dynamic routing process to enable making better incentive decisions. Full information about an order in AHD includes arrival time, customer address, order size and delivery time window. In this work, however, we only forecast the total number of orders over a day, their addresses and order volumes, but not the delivery time window of each order. The reason is that while we optimise the incentive decision, we aim to steer customers' choices of time windows. Any forecasting model ignoring the impact of the incentive decision will not do a good job of predicting how many orders would select each time slot in the end. On the other hand, the incentive decision is optimised dynamically over the entire booking horizon, changing over time and highly dependent on previously placed orders.

Therefore, we propose simplifying by assuming that all customers will select the time window most beneficial for the route planning to receive their orders. We also assume this is consistent with the company's goal of providing incentives. How to calculate more appropriate time window for an upcoming order will be discussed in more detail in Section 3.3.2. Here we only need an approach to predict the total number of forecast orders and their locations/sizes and assume that all forecast orders are granted a 24-hour time window.

For the total number of orders on a specific delivery day, n, we use a Simple Moving-Average (SMA) model:

$$n = \frac{1}{k}\sum_{i=1}^{k}\hat{n}_i \tag{3.3.1}$$

where $\hat{n}_i$ denotes the number of orders we received i weeks prior on the same weekday; k indicates the number of samples we consider for the prediction. We argue that this number n is not influenced significantly by the slot price/incentives we offer, as the number of customers in a fixed area and their intention to purchase from the e-retailer is mainly concerned with the demography of the area and the loyalty of customers. We also note that the moving average might not be the best possible approach that one can choose to predict the number of orders. More complicated machine-learning methods could be used to forecast the number of orders based on historical data. However, for this study, we only aim to demonstrate that incorporating forecast orders without any time windows into the routing process helps improve delivery efficiency and increases the total profit, even if a simplified model generates the forecast orders. For every single forecast order, its address, order size and order revenue are randomly simulated from historical data. Specifically, to generate one forecast order for a particular day of the week, we randomly choose (with uniform probability distribution) one order from the previous k weeks on that weekday and note its address, order size and revenue.

3.3.2 Opportunity cost approximation

As noted above, model (3.2.1) is not tractable for real implementations. This section presents an efficient approximation of it using dynamic routing. Many firms use dynamic routing to build their fulfilment plan while orders are still collected.

This work calculates approximations of opportunity cost by creatively using the given dynamic-routing package to make the system (3.2.1) solvable in a practical dynamic setting. One key idea here is to distribute the delivery cost in boundary condition (3.2.2) into stages and calculate the incremental delivery cost of accepting one more order in a specific area and time window. Similar ideas are effective in works such as Campbell and Savelsbergh (2006), and Yang et al. (2016); however, the approach used in this chapter is more advanced since it considers the potential revenue loss and incremental delivery cost by incorporating a forecast of future accepted orders.

The number of forecast orders and their locations/order sizes can be generated using the methodology presented in Section 3.3.1. We indicate the list of forecast orders by a vector $\vec{f}$, with $|\vec{f}| = n$, ($|\cdot|$ indicates the cardinality of a set), where n is given by (3.3.1). These virtual orders are put into the problem to help predict the final routes. The time window is the most significant difference between committed and forecast orders. Committed orders have their own time windows, as selected by customers, which are not changeable. However, forecast orders can be placed in whichever time window is most suitable because they have yet to be agreed upon with customers. This procedure allows the optimisation algorithm to choose an optimised delivery slot for the forecast orders based on the location and agreed slot of all actual orders collected so far. The optimised delivery time window for these forecast orders is then used as the time window to promote when an actual arrival is seen in the same area as that of the forecast order. In summary, forecast orders serve as dummy orders without time windows that guide the booking process.

3.3.2.1 Insertion cost

Let $DC_t(\vec{x}_t, \vec{f}_t)$ denote the total delivery cost obtained from the dynamic routing system at booking-horizon time t, of a list of already-accepted orders $\vec{x}_t$, and a set of forecast orders, $\vec{f}_t$ that remains in the system until time t. When it is infeasible to fulfil all orders $(\vec{x}_t, \vec{f}_t)$ with the given capacity and committed time slots for $\vec{x}_t$, we define $DC_t(\vec{x}_t, \vec{f}_t) = \infty$. We build a series of incremental delivery-cost approximations in the interim periods via dynamic routing with forecast orders. In more detail, the insertion cost of having one more delivery in area a and slot s, $IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, is approximated by

$$IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) = \begin{cases} \min\limits_{j_s \in \vec{f}_{rad}} \left[DC_{t+1}(\vec{x}_t + \mathbf{1}_{as}, \vec{f}_t - j_s) - DC_{t+1}(\vec{x}_t, \vec{f}_t) \right] & \text{if } \vec{f}_t \neq \emptyset \\ DC_{t+1}(\vec{x}_t + \mathbf{1}_{as}, \emptyset) - DC_{t+1}(\vec{x}_t, \emptyset) & \text{otherwise} \end{cases} \tag{3.3.2}$$

where $\vec{f}_{rad} \subseteq \vec{f}_t$ denotes all forecast orders that are not exceeding a radius equal to rad miles away from the new order, $j_s \in \vec{f}_{rad}$ indicates the forecast order to be removed while inserting the new order (a, s), and $j_s^* \in \vec{f}_{rad}$ denotes the best forecast order which is identified to be removed. Note that the location of the removed order might be different from that of the new order due to the forecast error and the rule of replacement (explained in detail in Section 3.3.3).

Note further that we did not need to re-run the VRP when computing equations (3.3.2). Instead, to calculate the extra driving time/cost in reaching order (a, s) and omitting the order j_s^*, we just estimated the extra driving distance and time which would be required as a deviation from the existing route found by the VRP, to accommodate these two changes.

3.3.2.2 Displacement cost (Revenue loss)

The insertion cost (3.3.2) forms one part of the opportunity-cost estimation, whereas the other part comes from the expected revenue loss by accepting order (a, s), denoted by $RL_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$. Provided that the current route (consisting of both actual and forecast orders) is always treated as the optimal route at the end of the booking horizon, we can construct our displacement cost/revenue loss estimation. According to the state-of-the-art approach used in attended home delivery literature, such as Koch and Klein (2020), we interpret the potential revenue loss as the "additive monetary value of the time window consumption" due to the acceptance of a new order. Let $w_{s'}(\vec{x}_t, \vec{f}_t)$ denote the idle time of the current route in time slot s'; then, after performing the replacement of forecast order j_s^* by the new order $\mathbf{1}_{as}$, the idle time is represented by $w_{s'}(\vec{x}_t + \mathbf{1}_{as}, \vec{f}_t - j_s^*)$. The revenue loss, therefore, is formulated as:

$$RL_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) = \sum_{s' \in S} \theta_{t,s'}(w_{s'}(\vec{x}_t, \vec{f}_t) - w_{s'}(\vec{x}_t + \mathbf{1}_{as}, \vec{f}_t - j_s^*)) \tag{3.3.3}$$

where $\theta_{t,s'} \in \mathbb{R}$ denotes the expected future revenue income per unit-time in slot s', that is evaluated at booking horizon t. While unlike Koch and Klein (2020) who learn the $\theta_{t,s'}$ values through sample simulation, in this work, we estimate the value of $\theta_{t,s'}$ using the current best route from the dynamic vehicle-routing solutions, as:

$$\theta_{t,s'} = \frac{\sum_i \{r_i | i \in \vec{f}_t,\ u_{s'-1} \leq \tau_i \leq u_{s'}\}}{u_{s'} - u_{s'-1}} \tag{3.3.4}$$

where τ_i indicates the delivery time of order i in the current best route, and $u_{s'}$ denotes the finishing time of slot s'. The numerator of (3.3.4) represents the total revenue (i.e.

$\sum r_i$) for all forecast orders i scheduled to be delivered in slot s'; and the denominator represents the duration of time slot s'. The opportunity-cost estimation is then:

$$OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) = IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) + RL_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) \tag{3.3.5}$$

which can be used to replace the $(V_{t+1}(\vec{x}_t) - V_{t+1}(\vec{x}_t + \mathbf{1}_{as}))$ in equation (3.2.1), so that the DP program can be reformulated as:

$$\tilde{V}_t(\vec{x}_t) \approx \left[\max_{\vec{d}} \sum_a \lambda \mu_a \sum_{s \in F_a(\vec{x}_t, \vec{f}_t)} P_{s, F_a(\vec{x}_t, \vec{f}_t)}(\vec{d}_a) \left[r_i + d_{as} - OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) \right] \right] + V_{t+1}(\vec{x}_t),$$

$$\forall\, \vec{x}_t \in \mathcal{X}, \tag{3.3.6}$$

with all elements known (except for $V_{t+1}(\vec{x}_t)$; but this term is not relevant in pricing optimisation) for every new order arriving at the system. This approximation decomposes the MDP into single-stage decision problems, provided that the opportunity cost $OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$ is evaluated dynamically over time. Note that the difference in (3.3.3) can be negative, indicating that replacing a forecast order with an actual one can help to travel less. We can interpret the negative value for RL as accepting the incoming order will not endure revenue loss and lead to higher slot availability for future customers. Therefore, the opportunity cost regarding RL, in this case, favours accepting the incoming order.

In real-world scenarios, determining the probability of selecting a time slot s at time step t or opting not to select any time slot is typically a challenge. In this context, we build upon the Multinomial Logit (MNL) model, as introduced by McFadden et al. (1973). This model defines the probability of choosing a time slot ($P_s(\vec{d})$) or not choosing any time slot ($P_0(\vec{d})$) at time step t, based on the available time slots and the delivery charge vector $\vec{d}$ at each customer's instance. Since the total probabilities sum to unity, we can calculate the probability of not making a booking, denoted as $P_0(\vec{d}) = 1 - \sum_{s \in F(\vec{x}_t)} P_s(\vec{d})$. The model described above are used, which states that these probabilities follow the equations below:

$$P_s(\vec{d}) = \frac{\exp(\beta_0 + \beta_s + \beta_d d_s)}{\sum_{k \in F_a(\vec{x}_t, \vec{f}_t)} \exp(\beta_0 + \beta_k + \beta_d d_k) + 1}, \tag{3.3.7}$$

where β_0 is the base utility on all choices, β_s is the utility of slot s itself, and β_d holds the utility sensitivity to delivery charge d_s. These β values are found by numerical optimisation for the historical data of purchases made and reflect the popularity of different times of day and the inferred price elasticity of demand. Similarly, the probability of no-booking under the delivery charge $\vec{d}$, $P_0(\vec{d})$, is given by:

$$P_0(\vec{d}) = \frac{1}{\sum_{k \in F_a(\vec{x}_t, \vec{f}_t)} \exp(\beta_0 + \beta_k + \beta_d d_k) + 1} \tag{3.3.8}$$

As Dong et al. (2009) show, under this choice model, given $OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, the optimal solution d_s^* to the online pricing problem can be achieved for $s \in F_a(\vec{x}_t, \vec{f}_t)$ with:

$$d_s^* = OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) - r_i - \frac{h}{\beta_d}, \tag{3.3.9}$$

where h is the unique solution to:

$$(h-1)\exp(h) = \sum_{s \in F_a(\vec{x}_t, \vec{f}_t)} \exp(\beta_0 + \beta_s + \beta_d(OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) - r_i)) \tag{3.3.10}$$

In Section 3.3.3 we give more details about how the order to remove, j^*, can be identified and how the order replacement is carried out in an online setting.

3.3.3 Insertion-cost evaluation and order replacement

Upon a new arrival in area a, the insertion feasibility and insertion cost have to be evaluated for fulfilling this order in time slot s, based on what orders have been accepted so far. As mentioned earlier, in this study, we maintain a delivery plan of all accepted orders and forecast orders dynamically and assume that at any interim stage, a "current best route" is available from the company's CVRPTW solver. As time goes by, we aim to replace forecast orders with actual incoming orders, one by one. For every potential replacement in a particular time slot s, we identify the best forecast order to remove, i.e., j_s^*, and calculate the incremental delivery cost involved in the replacement, and make it an estimate of the insertion cost for this time slot, i.e., $IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$. The methodology is summarised in Algorithm 1.

Algorithm 1 is called when we need to calculate the opportunity cost approximation to display available slots with the optimised pricing to customers. Let us consider a customer request i from area a to book a delivery slot. For this request, we can easily identify a subset of neighbouring forecast orders to replace by measuring their distance to the customer's location. To cover different implementation scenarios, we can apply different rules to calculate the subset, such as radius, road distance and/or postcode sectors. Let us denote the subset of forecast orders to remove by $\vec{f}_{rad}$. In this study, to find $\vec{f}_{rad}$, the algorithm will check available forecast orders within radius rad. If there are no forecast orders within radius rad, we increase the radius gradually with pre-set radius bands until forecast orders are found. The algorithm's performance is affected by the largest possible radius considered in implementation. If it is too small, the availability of slots might be restricted; if it is too large, the delivery cost prediction will require significant processing time to explore all forecast orders in the range. The

Algorithm 1: Opportunity-Cost Estimation for a Potential New Order i in Area a.

Compute set $\vec{f}_{rad}$ (set of candidate forecast orders, within radius rad, to be replaced by the new order) for new order i in area a.
for *each time slot* $s \in S$ **do**
  Compute insertion cost, $IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, according (3.3.2), denote the best order to remove as j_s^*.
  if $IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) \neq \infty$, *i.e., the insertion into slot s is feasible* **then**
    Record the best forecast order to remove, j_s^*, for later use in Alg. 2
    Calculate the potential revenue loss of accepting one more order for slot s, $RL_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, according to (3.3.3)
    Calculate opportunity cost for slot s, $OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, according to (3.3.5)
  else
    Set unavailable for slot s and area a
Solve (3.3.6) with the opportunity costs $OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as})$, and display available slots with the optimised pricing (3.3.9) to customers.

trade-off between processing time and forecast accuracy is presented in Section 3.4.5 with numerical results.

When all forecast orders are removed already, i.e., $\vec{f}_{rad} = \emptyset$, the algorithm will check the feasibility of inserting the new order without replacement. This is the same as the typical insertion heuristics carried out by, for example, Campbell and Savelsbergh (2006). This helps mitigate forecasting errors. Indeed, forecast orders are particularly important at the start of the booking horizon in guiding orders to suitable time slots. They become less important towards the end when the accepted orders with time windows nearly fix the routing plan. When the number of forecast orders is higher, we remove all remaining forecast orders in the end. Section 3.4.4 focuses on the forecast order levels and how the number of forecast orders may affect the performance.

Algorithm 2: Order replacement/insertion upon customer selection

1 Denote the slot which the customer has selected by $\hat{s}$; the new order being made in area a by i; and the forecast order to be replaced by $j^*_{\hat{s}}$
2 **if** $j^*_{\hat{s}} \neq \emptyset$ **then**
3   Remove $j^*_{\hat{s}}$ from the route
4   Update $\vec{f}_t \leftarrow \vec{f}_t - j^*_{\hat{s}}$
5 Insert order i into slot s of the route, Update $\vec{x}_t \leftarrow \vec{x}_t + \mathbf{1}_{a\hat{s}}$
6 Re-optimise the route for order set $(\vec{x}_t, \vec{f}_t)$

After the customer selects their preferred slot, Algorithm 2 will perform the actual replacement/insertion of the new order i in the selected slot $\hat{s}$. Note that the resulting approach is more robust if traffic conditions, schedule tightness, potential lateness, etc., are considered in detail in the dynamic routing system. Upon the acceptance of a new order, the CVRPTW is re-optimised/updated for a better delivery schedule until the next arrival comes to the system. This means that the routing optimisation could adjust the delivery time of the remaining forecast orders in the system to re-optimise the route.

3.4 Numerical results

To investigate how our approach performs in practice, we test the proposed methodology on four typical delivery areas, each characterised by different customer densities and distribution patterns. Actual customer locations and historical booking data are used in the tests, with essential manipulations to protect commercial information and customer privacy. The selection of customer locations, whether actual or forecast, is performed using (3.3.1) randomly from the previous seven periods to ensure both recency and stochasticity.

3.4.1 Routing Package

As explained in Section 3.3.2, the designed approach can be implemented with any routing packages for CVRPTW. This includes dynamic-routing packages with automatic updating schemes based on accepting new orders and static routing packages that allow warm-starting from known solutions obtained from replacing/inserting the new orders in the current best routes. Such routing packages are understood to be standard tools currently used by delivery companies. In our experimental tests, however, we do not rely on any company's specialised routing package but use a generic meta-heuristic, i.e., Simulated Annealing (SA), to solve the CVRPTW. The application of SA to the VRPTW was introduced and proven effective in terms of accuracy and execution time by Chiang (1996).

Implementing the SA in the dynamic routing can be characterised into offline and online phases. The offline phase is a Capacitated Vehicle Routing Problem (CVRP) without time windows that occur entirely before the booking horizon, which consists of only forecast orders. The best possible route for the forecast orders (and their best time slots) are found using SA, preliminary to the order acceptance process. The online phase covers the entire booking horizon when actual orders are collected, and forecast orders are replaced as time goes by, according to Algorithm 1 and Algorithm 2. SA is called to re-optimise the current route after each replacement is done until the arrival of the following order.

3.4.2 Experiment Settings

We test our model on four typical area settings to investigate different scenarios regarding the spread of orders. These areas are a Rural area, which reflects the countryside and small villages; a Semi-rural area, which represents towns; a Suburb area, which represents the outskirts of big cities where people live in disjoint but not far-way satellite communities; and a City area, where people live with the highest density, e.g., in apartments. Each area has a depot, whose location/distance to the primary service area is set to capture the business settings. See Figure 3.3.

As what has been used by Yang et al. (2016), we fit a non-homogeneous Poisson distribution to historical data and use it to generate customer arrival times in simulation. Order details, including customer addresses and order sizes, are randomly simulated from real orders in the past. We also borrow the time slot and MNL parameters from Yang et al. (2016), i.e., 27 slots per day, one hour each; some overlap with others. Please refer to (3.3.7) and (3.3.8) for the detailed MNL customer-choice model parameters.

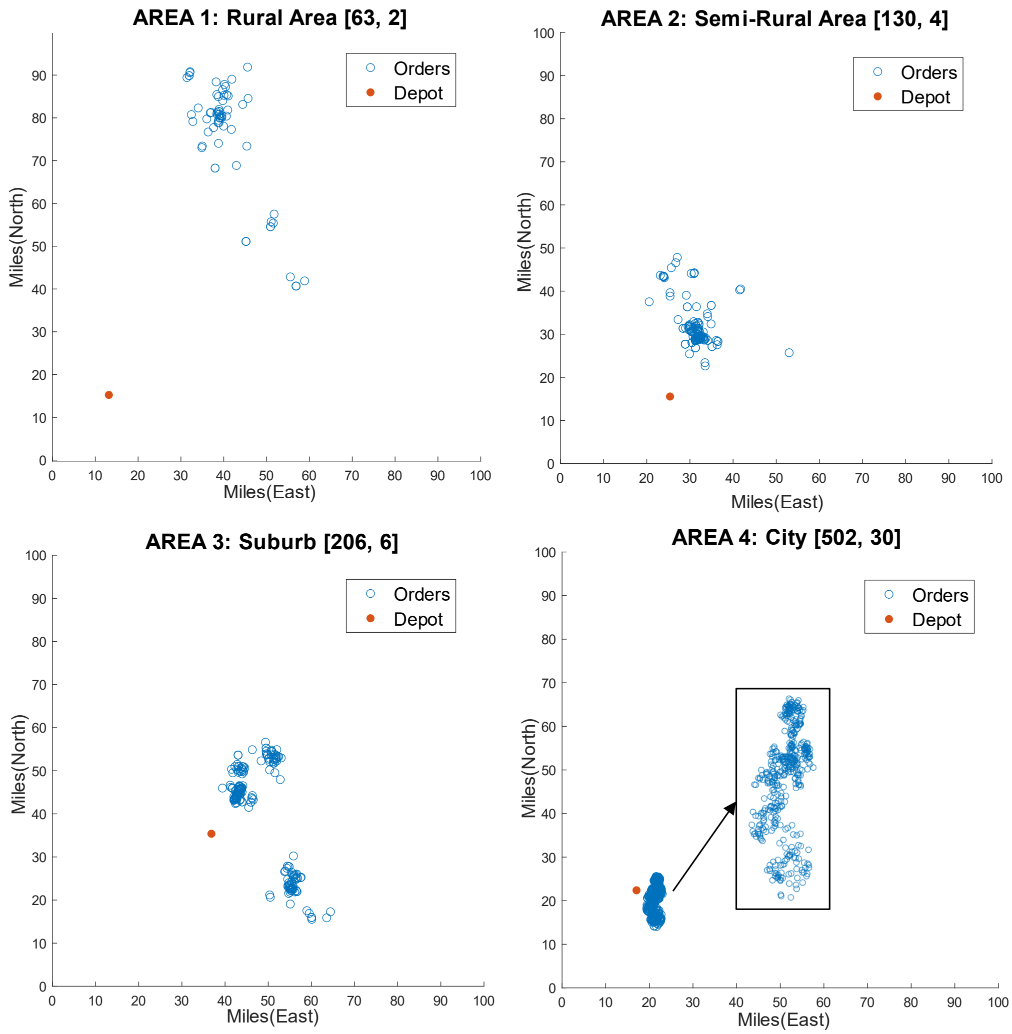

Figure 3.3: Different areas with different densities and spread patterns of orders. A detailed number of forecast orders and vehicles for each area is given as [#forecast orders, #vehicles] in the subtitles.

Continuous slot prices, in the range £[-10, 10], are considered in the simulation (i.e. these are the slot prices offered to customers), with negative prices indicating discounts offered to customers as an incentive to purchase. The commercial partner we are collaborating with has approved this pricing scheme for this project. Detailed revenue/profit and order size information cannot be published due to commercial concerns, but for a meaningful interpretation of the results, the ratio between order revenue and variable delivery cost is set to 40.5 to reflect the real situation.

3.4.3 Experiment Results

This section aims to test the effectiveness of the proposed dynamic-pricing approach by maintaining a set of forecast orders without time windows in the routing system. The studied methods are as follows:

1. **Static Pricing (SP)**: Implements a fixed pricing strategy where each slot has a static price of £3 throughout the booking horizon.

2. **Dynamic Pricing with Insertion Cost without Forecast Orders (DP-IC):** This method employs a short-sighted approach in estimating OC by considering only the marginal insertion cost of the new order into the current route plan based on orders accepted up to the present time. Consequently, at the beginning of the booking horizon, the insertion cost is high and gradually decreases as more orders are accepted. This short-term perspective guides the dynamic pricing strategy to dynamically adjust slot prices to maximise immediate profit. Studies such as Yang et al. (2016) and Yang and Strauss (2017) use this approach as a benchmark method.

3. **Dynamic Pricing Fixed Routing Forecast (DP-FR-F)**: is proposed by Yang et al. (2016) as a foresight policy that employs marginal insertion costs to estimate OC. Unlike DP-IC, however, this approach estimates marginal insertion costs using fixed, time-windowed forecasted orders derived from historical routes. The forecast route is constructed at the beginning and fixed over the entire booking horizon, i.e., the forecast route is not updating with new order acceptance. This approach allows for more effective OC calculations by basing them on a forecast route plan.

4. **Dynamic Pricing with Dynamic Routing of Time-Windowed Forecast**

Orders (DP-DR-TWF): Integrates forecast orders that have predefined time windows based on historical data, utilising updates in routing based on new order acceptance.

5. **Dynamic Pricing with Dynamic Routing of Forecast Orders without Time Windows (DP-DR-F):** Estimates OC based on both marginal insertion cost and displacement cost (revenue loss). It considers a forecast route generated from forecast orders without time windows as well as actual orders. The forecast route continually updates as new orders are accepted. Consequently, the OC estimation relies on a dynamic route plan, which incorporates the latest route changes. This has been proven to be an effective approach.

These methods are tested against each other using the same routing package and customer choice model (MNL), ensuring consistency in the experimental conditions. All other assumptions are the same for the tests, including:

- The same routing package is used for each paired test, i.e., the Simulated Annealing (SA) for CVRPTW as described in Section 5.1 with the same tunable parameters;
- All implement dynamic routing where updating (re-optimisation) of the current best route is performed after the acceptance of every new order until the next order arrives;
- All deploy the same MNL customer-choice model estimated from real data for customer selections;
- All deploy the same approach for insertion cost and revenue loss (where forecast orders exist) estimation according to 3.3.2 and 3.3.3. The only difference is the forecast routes (or whether there is a forecast route) used.

To have a fair test on the performance of the pricing policy alone, the feasibility of placing an order in a slot under these two benchmark policies is also informed by the dynamic-routing package.

Experiments are carried out using MATLAB on an Intel Core i9-7940X 3.1GHz machine. Since we deployed a meta-heuristic approach to solve the CVRPTW, we conducted 30 independent runs. We reported the average and the standard deviation (mean(s.d.)) to minimise the influence of the randomness involved in the solution approach. Performance on crucial indicators for profit and efficiency are presented in Table 3.1, with the best one in every row highlighted in bold. The results show that the DP-DR-F outperforms the other approaches in all the measurements, which confirms the effectiveness of the proposed approach.

Table 3.1: Results of SP, DP-IC, DP-FR-F, DP-DR-TWF and DP-DR-F methods

		SP	DP-IC	DP-FR-F	DP-DR-TWF	DP-DR-F
Rural	Total mileage	578.93(27.61)	575.65(22.84)	577.45(23.27)	562.46(22.10)	**543.08(24.49)**
	Mileage/Order	9.48(1.08)	7.70(0.89)	7.64(0.68)	7.84(0.84)	**6.75(0.49)**
	Number of orders	60.60(5.25)	74.40(5.97)	75.00(4.69)	71.37(6.41)	**79.66(3.71)**
	Total order size	213.38(21.20)	273.07(23.23)	273.35(19.90)	267.38(23.38)	**296.60(15.98)**
	Accepted price	3.00	-4.13(0.65)	-3.86(0.54)	-1.47(0.34)	-3.15(0.63)
	Total Profit	2109.36(206.98)	2168.66(189.12)	2187.86(157.03)	2226.23(208.64)	**2433.46(133.50)**
	Improvement on SP		2.81%	3.72%	5.54%	15.36%
Semi-rural	Total mileage	459.41(18.35)	446.16(17.39)	450.14(17.61)	419.76(21.10)	**415.79(16.38)**
	Mileage/Order	3.66(0.36)	3.09(0.26)	3.09(0.22)	2.94(0.26)	**2.67(0.14)**
	Number of orders	125.43(8.83)	144.40(9.84)	145.03(7.04)	142.33(7.89)	**152.00(6.22)**
	Total order size	448.47(30.41)	539.69(34.84)	545.63(28.63)	545.02(27.27)	**583.50(20.29)**
	Accepted price	3.00	-3.79(0.36)	-3.41(0.39)	-0.90(0.13)	-1.11(0.55)
	Total Profit	4431.28(300.94)	4341.81(270.34)	4447.90(231.75)	4685.57(239.68)	**5108.86(171.81)**
	Improvement on SP		-2.02%	0.38%	5.74%	15.29%
Suburb	Total mileage	640.77(24.33)	622.74(20.18)	622.67(23.20)	613.99(20.11)	**577.59(25.82)**
	Mileage/Order	3.18(0.35)	2.65(0.17)	2.65(0.18)	2.88(0.18)	**2.35(0.11)**
	Number of orders	202.03(15.65)	234.73(9.05)	234.20(8.90)	212.97(9.16)	**245.27(5.03)**
	Total order size	736.01(57.25)	891.57(32.52)	893.51(32.36)	817.10(36.21)	**927.40(20.93)**
	Accepted price	3.00	-3.80(0.26)	-3.50(0.30)	-0.00(0.34)	-0.59(0.50)
	Total Profit	7263.83(564.16)	7182.95(291.06)	7273.36(280.74)	7396.60(358.20)	**8249.52(230.96)**
	Improvement on SP		-1.11%	0.13%	1.83%	13.57%
City	Total mileage	453.07(7.63)	449.66(7.74)	449.38(8.63)	432.16(11.60)	**407.63(10.15)**
	Mileage/Order	0.99(0.04)	0.89(0.04)	0.89(0.03)	0.95(0.04)	**0.76(0.03)**
	Number of orders	458.43(14.20)	505.07(14.65)	504.43(13.94)	455.33(13.53)	**531.27(13.82)**
	Total order size	1702.06(57.96)	1966.86(57.56)	1966.51(58.29)	1794.86(72.67)	**2060.32(58.49)**
	Accepted price	3.00	-3.86(0.21)	-3.45(0.20)	1.53(0.16)	3.25(0.60)
	Total Profit	16777.67(562.93)	15858.71(435.56)	16066.44(463.20)	17676.47(667.65)	**20372.63(355.69)**
	Improvement on SP		-5.48%	-4.24%	5.36%	21.43%

One important observation here is that dynamic pricing (DP-IC) is not necessarily better than static pricing (SP), mainly when a poor estimation of the opportunity cost is used. Furthermore, the short-sighted incremental delivery cost based on accepted orders alone is insufficient for opportunity-cost estimation. This insight is in line with the conclusions of Yang and Strauss (2017), which emphasise the importance of "incorporating the impact of future profit opportunities from orders".

As this work proposes, incorporating forecast orders (without time windows) provides an easy method for future-profit estimation. Together with the better marginal delivery-cost estimation, obtained by maintaining a hybrid route with both forecast and actual orders, the approach achieves a 13.57-21.43% profit increase over the static-pricing method, which is better than that of 2.2-2.5% in Yang and Strauss (2017) and that of 2.6-6.2% in Yang et al. (2016). We also performed paired sample t-tests on the total profits for all the approaches in the studied areas. The p-values produced were all less than 0.05, indicating that the differences were statistically significant at a 5% significance level.

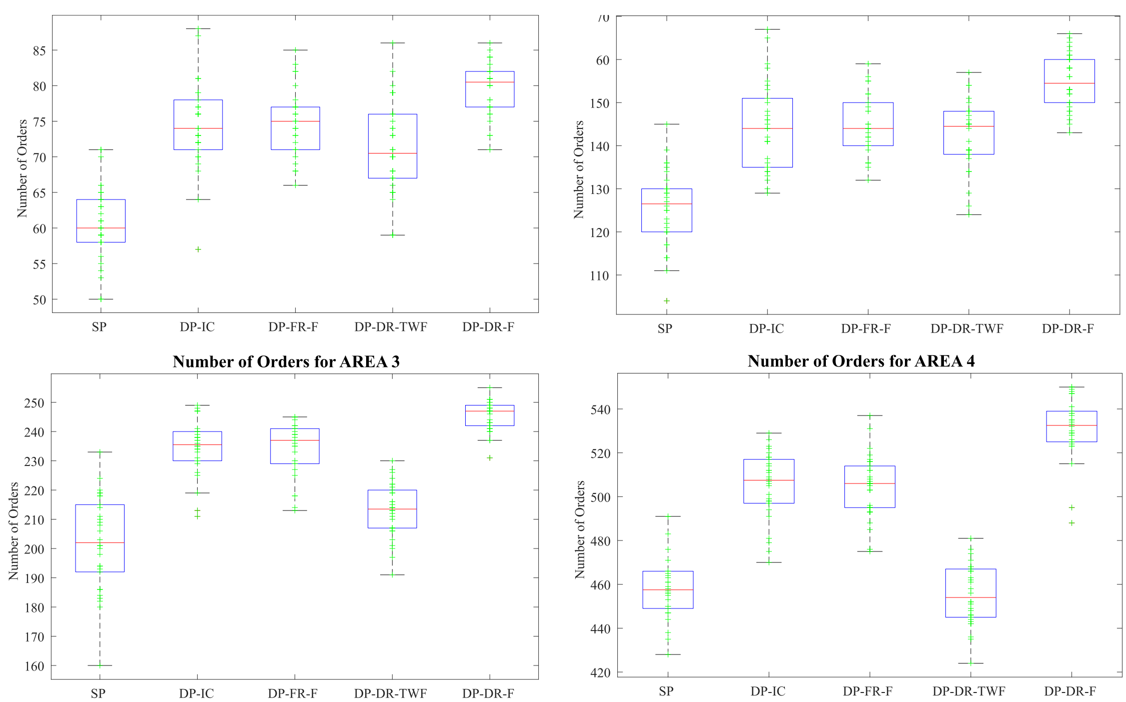

Figure 3.4: Total number of orders accepted by the methods

We look closely at key-related components to understand where the additional profits come from in the DP-DR-F approach. Figures 3.4 and 3.5 show a graphical comparison of the number of order commitments and the average travelling distance per order across all approaches. These elements demonstrate the efficiency of the final routes we end up with using the DP-DR-F approach so as to justify the capability of DP-DR-F in recognising the best time slots to offer and promoting them via dynamic pricing.

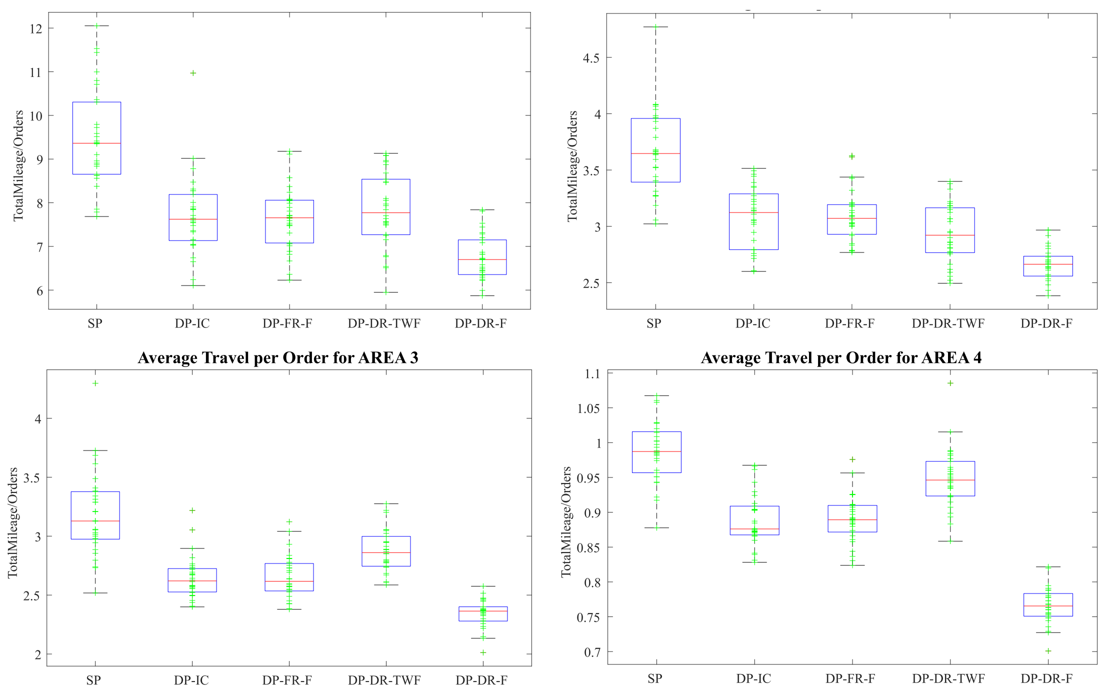

Figure 3.5: Average travel distance to satisfy an order

Compared to DP-DR-TWF, which uses time-windowed forecast orders, the DP-DR-F approach omits the time windows from the forecast orders, which thus allows the dynamic route planning to adjust the delivery time of a forecast order to the best possible time slot according to its location and fitness to the current best route. This feature provides extra flexibility for dynamic route planning to create more feasible slots over the booking horizon (more information to follow in Figure 3.7) and guide the overall booking process to a more compact route in the end. Higher profits are therefore achieved through selling more goods. These are believed as the key reasons for higher order commitments provided by DP-DR-F with the fixed fleet and time window capacity.

Figure 3.6 shows samples of the final routes obtained by the five approaches in Area 1. Comparing the plots in Figure 3.6, we can see that the DP-DR-F approach gives the most efficient routes, with noticeably fewer long links than the others. It performs especially well in the "remote area", which directs the orders in this area to adjacent slots, so as to avoid the van coming back to this area multiple times to meet demands at different times of the day. The improved route plan shows DP-DR-F's ability to promote the correct time slot that complies with the optimal route to reduce unnecessary travel to meet customer needs.

Figure 3.7 is related to how many slots are available on average as the booking process progresses for all methods. Based on this plot, we can see that all methods begin with a high slot availability, which decreases as time passes. Decreasing rates for all the other four approaches are similar, whereas, for DP-DR-F, the slope magnitude is lower. The higher availability of time slots with the DP-DR-F method stems from the booking process, where incoming actual orders are allocated to more efficient time slots for routing. This efficiency is achieved through the use of forecast orders and improved slot pricing, which is informed by better OC calculations. These factors together enhance routing efficiency and result in improved slot availability. This outcome justifies that DP-DR-F works well in reserving resources for later usage to provide more stable slot availability over the entire booking horizon than the benchmark approaches. The higher availability leads to a higher selection rate on average and therefore conveys more orders.

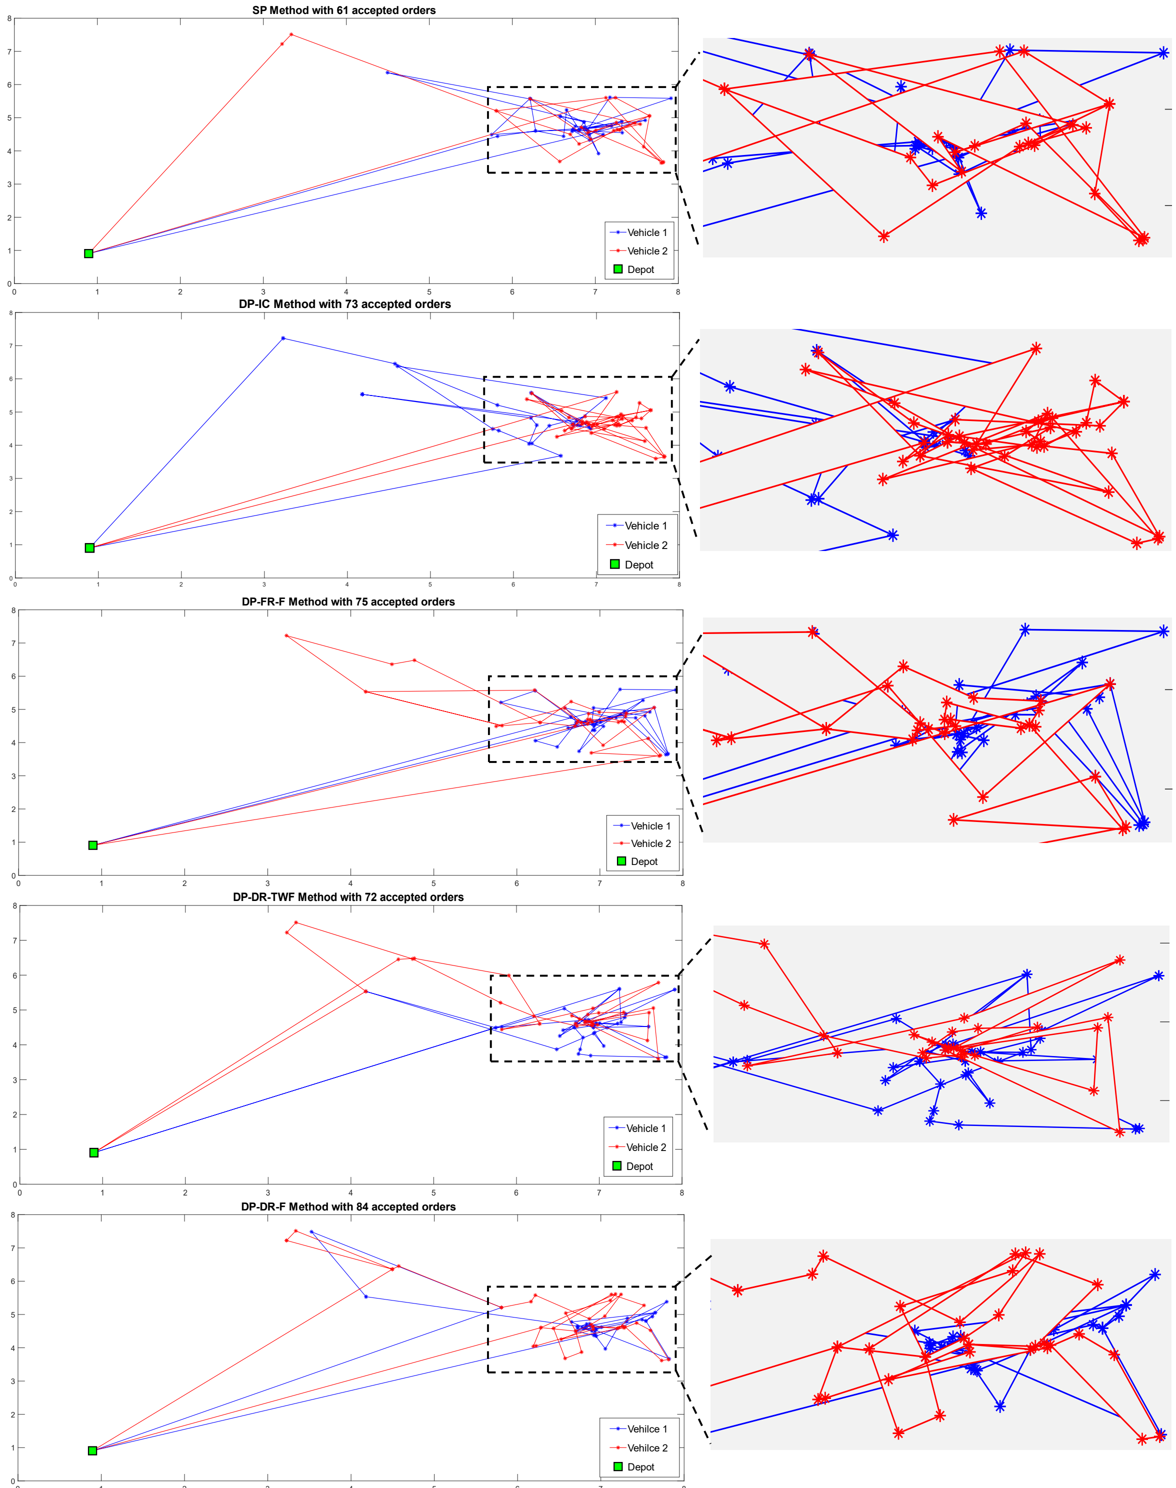

Figure 3.6: Final route for Area 1 using SP, DP-IC, DP-FR-F, DP-DR-TWF and DP-DR-F methods

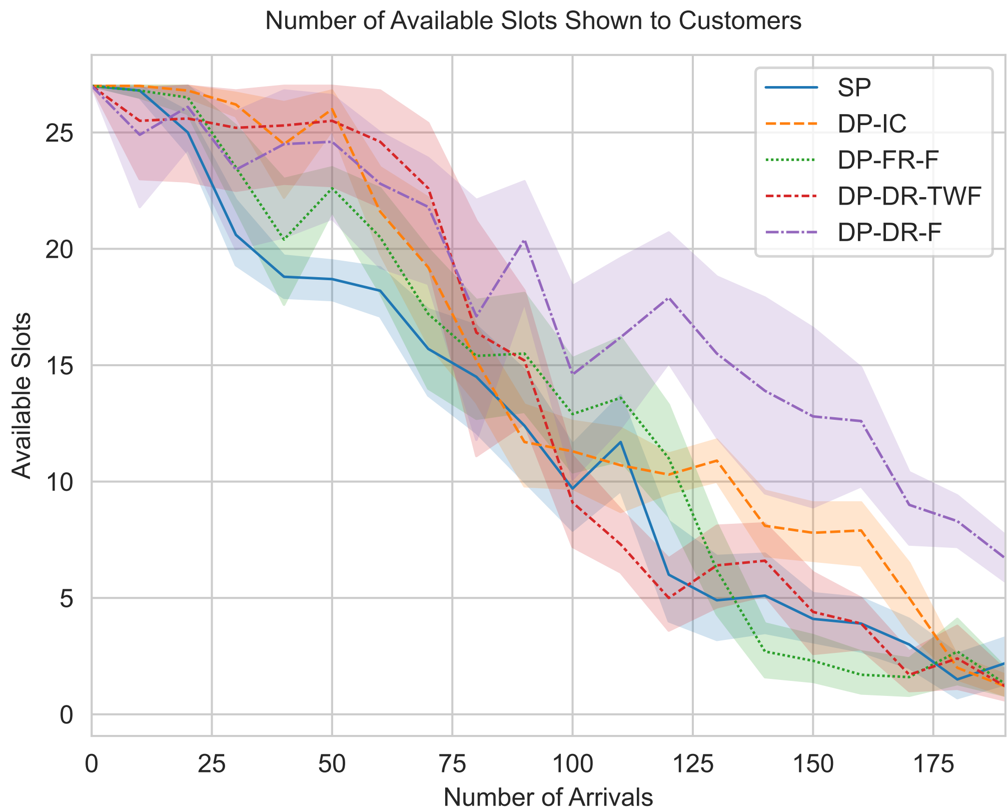

Figure 3.7: Comparison of slots availability over time for Area 1

In addition, by applying DP-DR-F, a significantly higher number of slots are still available when the booking horizon reaches its end. However, this approach has committed a notably higher number of orders than others. This upshot shows further profit growth potentials of this approach suppose a more extensive market can be reached by the firm, which is not possible with any other approaches.

Another critical term influencing the total profit is slot price, which is the fee customers pay for the delivery service. Figure 3.8 shows the average slot prices offered to the time slots that customers eventually select with every approach. It is not hard to see that the DP-IC approach outperforms the static approach in the number of accepted orders (Figure 3.4) and the per-order delivery costs (Figure 3.5). However, there is no significant improvement in the overall profit due to the low average price it charges.

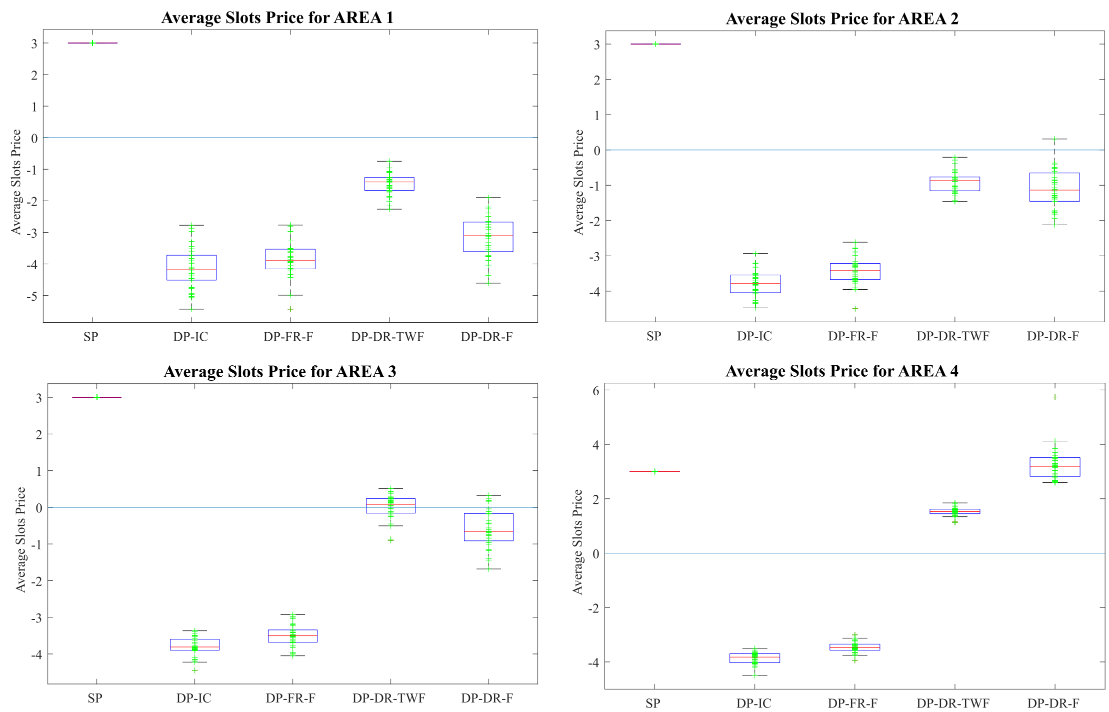

Figure 3.8: Average prices for the slots that are booked by customers

Upon the arrival of a new order, there is a trade-off between long-term profit and immediate gain. This balance explains why the prices offered by the DP-IC and DP-FR-F approaches are constantly lower than those of DP-DR-TWF and DP-DR-F. Without considering the expected revenue loss in the opportunity-cost estimation, the DP-IC and DP-FR-F approaches only focus on the immediate gain brought by the order under consideration and try very hard to persuade this customer to buy by lowering the price offered. However, with the DP-DR-TWF and DP-DR-F approach, as they still expect more future orders, the pressure of conveying an order now is lighter, so the price offered is higher on average. Comparing DP-DR-TWF and DP-DR-F, however, the prices charged by DP-DR-F are slightly lower in most cases. This is because, in DP-DR-F, the best time slot is purely identified through the dynamic routing system without time windows, which takes no consideration of the original popularity of time slots. This puts higher pressure on lowering the price to persuade a customer to book an undesirable time slot if their location is deemed the best fit for that time slot.

Note that the average prices offered by dynamic approaches are, in most cases, lower than zero. The finding is consistent with the previous study using the same choice model, i.e., Yang et al. (2016) and Yang and Strauss (2017). As claimed in both previous works, the profit from selling an order is much higher than the profit from making a delivery. The system offers discounts to encourage customers to buy, rather than highly charging them for delivery.

Figure 3.9 displays the offered prices change over time on a sample run in Area 1. Prices offered by all dynamic pricing schemes are decreasing over time.

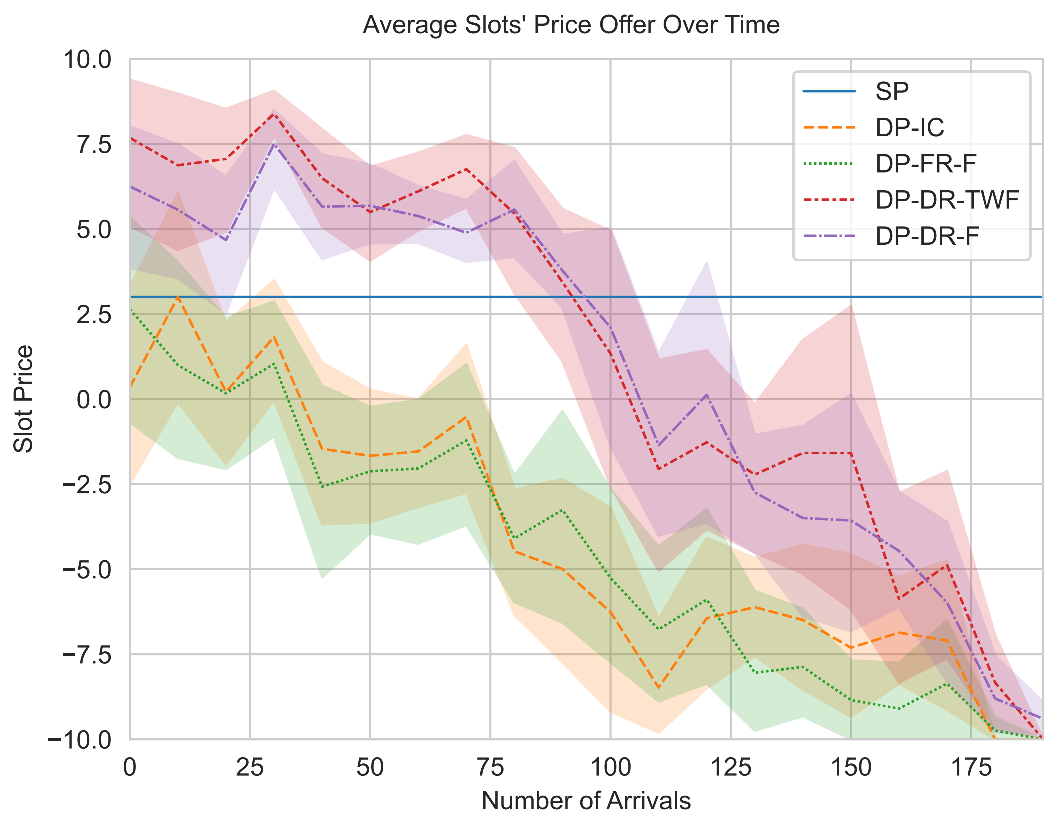

Figure 3.9: Average slots' price offered over time for Area 1

Such a trend is understandable, as when slot availability decreases, the pricing problem expects a higher no-booking rate, so it tends to lower the price to persuade customers to buy. Also, as explained in Yang et al. (2016), popular slots are filled in earlier than unpopular ones in the booking horizon, so lower prices must be offered further to promote unpopular slots towards the end of booking time.

3.4.4 Analysis of the number of forecast orders

As explained in section 3.3, we use a moving average to estimate the number of forecast orders in each area, which may lead to, sometimes significant, forecast errors. This subsection investigates how sensitive the DP-DR-F approach is to forecasting errors. To this aim, we create scenarios where the number of forecast orders is significantly (20%) higher or lower than the moving average and test the DP-DR-F on them. Table 3.2 shows the obtained results for 30 independent runs in the format of the average and the standard deviation (mean(s.d.)) for the four studied areas, based on the same set of random examples that have been used in Table 3.1 to make results comparable across tables.

Table 3.2: Comparison of different forecast order levels for the 4 studied areas

		80% Forecast	**100% Forecast**	**120% Forecast**
Rural	Total mileage	548.04(24.81)	543.08(24.49)	522.15(21.05)
	Mileage/Order	6.94(0.53)	6.75(0.49)	6.82(0.61)
	Number of orders	78.23(4.34)	**79.67(3.71)**	76.03(5.90)
	Total order size	288.49(17.93)	**296.60(15.98)**	282.80(19.53)
	Accepted price	-3.26(0.64)	-3.15(0.63)	-1.70(0.88)
	Total Profit	2356.21(155.07)	**2433.47(133.50)**	2428.21(151.38)
Semi-rural	Total mileage	417.00(18.27)	415.79(16.38)	375.37(14.96)
	Mileage/Order	2.70(0.17)	2.67(0.14)	2.59(0.23)
	Number of orders	153.77(5.73)	**155.00(6.22)**	144.67(12.06)
	Total order size	577.47(20.33)	**583.50(20.29)**	540.01(47.61)
	Accepted price	-2.22(0.49)	-1.11(0.55)	1.10(1.03)
	Total Profit	4885.20(170.11)	**5108.86(171.81)**	5040.52(382.49)
Suburb	Total mileage	585.34(26.56)	577.59(25.82)	554.94(23.39)
	Mileage/Order	2.40(0.18)	2.35(0.11)	2.45(0.15)
	Number of orders	243.43(9.24)	**245.27(5.03)**	226.13(9.81)
	Total order size	921.44(33.65)	**927.40(20.93)**	847.16(38.49)
	Accepted price	-1.98(0.46)	-0.60(0.50)	2.26(0.74)
	Total Profit	7858.11(276.22)	**8249.52(230.96)**	8175.12(322.48)
City	Total mileage	413.09(10.44)	407.63(10.15)	389.75(8.86)
	Mileage/Order	0.80(0.03)	0.77(0.03)	0.83(0.04)
	Number of orders	516.10(12.05)	**531.27(13.82)**	469.03(19.45)
	Total order size	2017.98(41.48)	**2060.32(58.45)**	1793.04(86.05)
	Accepted price	1.85(0.77)	3.35(0.60)	6.58(0.45)
	Total Profit	19229.16(391.95)	**20372.63(355.70)**	19315.16(799.10)

Concerning the results obtained from the simulation in Table 3.2, we can infer that if the moving average estimate (3.3.1) decides the number of forecast orders, we can expect higher performance in total profit. However, if we underestimate or overestimate the number of forecast orders compared to what the moving average method suggests (in our study, 20% lower or higher than the moving average estimate), the performance undergoes a decrease in efficiency. We can notice this point in all studied areas, which affirms the robustness of the proposed method to the number of order estimations. Moreover, the results for a higher number of forecast orders are slightly better than those for a lower number of forecast orders. This fact emphasises the critical role that forecast orders play in the booking process, even towards the end when actual orders with time windows become the majority of the group and the routes are relatively fixed.

3.4.5 Impact of radius size on performance and run-time

The proposed approach is online, so the time it takes to find feasible time slots and optimise their prices is crucial for successful implementation. One way of accelerating the decision process is limiting the number of forecast orders to replace when evaluating the insertion cost. This subsection presents how various forecast order radii can affect the proposed method's functionality and run-time. More precisely, when the DP-DR-F method wants to find the forecast orders for replacement with a new arrival, only the forecast orders located within a specific radius of the new arrival will be considered. The greater this radius is, the higher the number of neighbouring forecast orders will be, resulting in more computation time.

Table 3.3: The effect of different radii on performance and run-time of DP-DR-F method in Area 3. ADO stands for Average Distance between Orders.

Radius	Total mileage	Total Order size	Number of Orders	Total Profit	Run-time(sec)
0.01×ADO	580.80	891.90	234.37	8214.05	0.0041
0.10×ADO	577.59	909.20	240.43	8214.47	0.0104
0.50×ADO	581.34	914.99	241.50	8160.30	0.0336
1.00×ADO	573.67	922.15	244.40	8240.20	0.0519
∞	579.47	923.15	244.90	8251.86	0.0887

In Table 3.3, we define experiments with different radii based on order density, i.e., as a ratio to the average distance between orders (ADO). When the method tries to find the candidate forecast orders, all the forecast orders within a specific radius will be explored, and replacement feasibility will be conducted for each of them. If no forecast orders exist in this radius, the method will double the size of the radius to search for forecast orders. This process will continue until at least one forecast order is found.

The average running time to obtain a list of feasible slots, with their optimised prices, is reported in the last column of Table 3.3. This measure can be seen as the average online reaction time upon a customer's arrival. According to this table, there is an increasing trend in performance (e.g. profit) when we enlarge the area searched for a replacement. At the same time, the execution time increases more sharply, which is in line with our expectations. Compared to the largest possible radius, the reduced search range, e.g., to $0.01 \times ADO$, only sacrifices less than 0.5% of the total profit while reducing the average reaction time by 95.38% compared to the full search. The obtained results justify the effectiveness of the proposed simplification. In practice, the online grocer can choose a radius according to the maximum affordable run-time to have the best achievable results within the time limit.

3.5 Conclusion and Future Work

This chapter introduces a novel dynamic pricing method for attended home delivery using forecast orders without time windows. The approach maintains a dynamic route of actual orders (with time windows) and forecast orders (without time windows). It estimates opportunity costs online using the most up-to-date information in the dynamic route. No extra learning is needed. The approach can be integrated with any dynamic-routing package a company is using, which allows the company's specific routing needs and restrictions to be considered as well. The approach is tested on real data with an MNL customer-choice model.

This study operates under the assumption that dynamic pricing effectively balances customer demands, allowing uncommitted orders to be flexibly distributed across different time slots. However, a notable limitation lies in the practical aspect of customer availability, which can influence demand and is not currently accounted for in our model. Future research in Chapter 4 aims to address this by incorporating the original slot popularity into the forecast route planning of orders without specific time windows. This integration seeks to harness the advantages of both dynamic pricing and customer preference considerations, offering a more holistic and practical approach to managing order assignments and delivery efficiency.

By conducting experiments on four different areas, the advantages of employing forecast orders without time windows are observed in higher-order commitments, lower delivery costs and higher overall profits compared to all benchmarking approaches. The improvement of 13.57-21.43% profit on Static Pricing is better than the results from former approaches in Yang and Strauss (2017) and Yang et al. (2016) and of an amount likely to be of commercial interest to those managing AHD operations. The robustness of the DP-DR-F approach is also tested through experiments when the number of forecast orders is overestimated or underestimated. Potential accelerations of the approach through restricting the exploration radius are discussed and tested to improve running efficiency and suitability for online implementations. As indicated, Table 3.3 shows that reducing the search radius to $0.01 \times ADO$ cuts reaction time by 95.38% with minimal impact on profit (less than 0.5% loss), underscoring the efficiency of our approach. Additionally, variations in the ratio between the search radius r_i and delivery charges could affect these outcomes, highlighting the need for adaptable pricing and slot allocation strategies.

4

Efficient Forecast-Based Routing and Dynamic Time Window Management for Attended Home Deliveries

4.1 Introduction

By using forecast orders without time windows proposed in Chapter 3, the route planning can be more adaptable to accommodate actual orders as they emerge during the booking period. The absence of strict time constraints associated with specific time windows allows for greater flexibility in adjusting the routes. However, it is important to acknowledge that this approach may lead to predicted routes that are impractical, as the distribution of planned delivery times for forecast orders may not align with historical customer preferences. The routing package, without time-window constraints for forecast orders, focuses solely on minimising distances between them, leading to the potential of shuffling them freely around the time-line for delivery. In this study, the highest committable order numbers for each time slot are determined and used as realistic upper bounds to guide dynamic routing optimisation. This factor can enhance the opportunity cost calculation compared to the work in Chapter 3.

The availability of time slots emerges as a critical factor, influencing customer order placement and system profitability. This chapter also introduces extended-duration time slots, augmenting the fixed-length time slots in the service offerings. This enhancement aims to boost the flexibility of the CVRPTW and enhance slot availability during dynamic route planning. To encourage the selection of these longer time slots, customers are offered incentive discounts. When a customer opts for an extended time slot, they will also receive a notification detailing the expected delivery time on the day of delivery, specifying a time within the committed time window range.

The subsequent sections of the chapter are summarised as follows: Section 4.2 introduces a model for the AHD problem within the context of integrated demand management and dynamic routing. Further, contributions to the problem are proposed from two aspects in sections 4.3 and 4.4. The sections detail a solution approach to the AHD problem and present a novel strategy to enhance routing performance, respectively.

Section 4.5 explains the calculation of the optimal delivery charges. The numerical results obtained from the two suggested approaches are discussed in Section 4.6. These results are further analysed and summarised in Section 4.6, followed by insights that were uncovered and future research directions. Additionally, Section 4.7 provides a comprehensive summary of the key findings and conclusions of this study.

4.2 Dynamic pricing formulation in attended home delivery

This chapter builds upon the problem statement outlined in Chapter 3. The problem context remains consistent, focusing on an e-retailer managing an online booking system within a discrete and finite booking horizon. However, there are specific variations for this chapter. Notably, the number of time windows is reduced to 17, and time windows are non-overlapping, distinguishing this formulation from the previous one. The core problem and its elements, such as customer arrivals, the state representation, probability models, and the objective function, remain the same. For detailed context, readers are referred to Chapter 3.

4.3 Demand model and opportunity-cost approximation

This section builds upon the methods developed in Chapter 3 and is reproduced with some changes from Abdollahi, M., Yang, X., Nasri, M. I., and Fairbank, M. (2023). Demand management in time-slotted last-mile delivery via dynamic routing with forecast orders. European Journal of Operational Research.

The term "Opportunity Cost" (OC) refers to the potential profit loss that occurs when a time slot is assigned to the current customer instead of reserving it for future customers. In Equation (3.2.1), the term $\left(V_{t+1}(\vec{x}_t) - V_{t+1}(\vec{x}_t + \mathbf{1}_{as})\right)$ represents the difference in profit at time step $t+1$ between not allocating slot s in area a to the current customer and allocating it. This difference illustrates the potential profit loss, which is commonly referred to as OC. A substantial challenge in dynamic pricing problems is accurately approximating the OC, which critically relies on a crucial factor known as the displacement cost or revenue loss (RL). Much prior research has focused on addressing the estimation of RL, as seen in works by [Yang and Strauss (2017), Klein et al. (2018), and Strauss et al. (2020)].

To address the demand challenges inherent in the AHD problem, the methodology proposed in Chapter 3 was adopted. This method involves the creation of a forecast delivery route plan using historical data as a warm start at the beginning of the booking period for the dynamic routing system. The dynamic routing system continuously updates this route plan in the absence of customer arrivals, striving to optimise the routes in response to new orders. Additionally, this predictive plan serves as an initial approximation of the final route plan and is subsequently employed to compute the Opportunity Cost (OC). The overarching process of this approach encompasses customer arrivals within the booking window, driving the system to dynamically offer priced time slots for deliveries, as depicted in Figure 4.1.

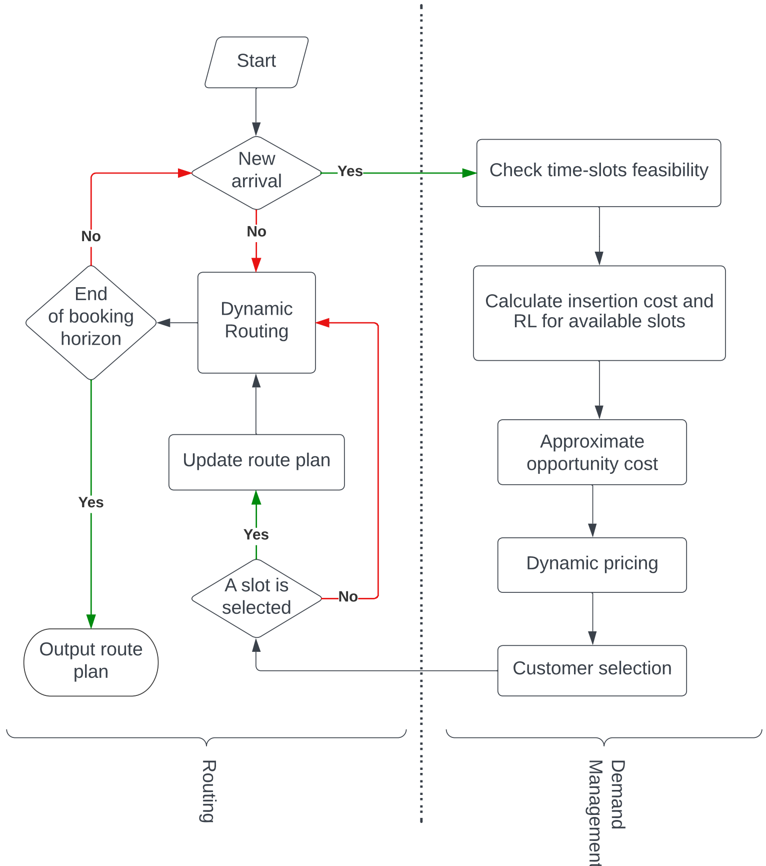

Figure 4.1: Implication of the integrated dynamic routing and demand management problem.

To ensure that the forecast route plan remains unaffected by any fixed pricing policy, the forecast route plan was built using forecast orders without predefined time windows, derived from the historical purchasing data. This approach enhances route planning by enabling the system to focus solely on optimising the delivery sequence, unencumbered by the constraints imposed by time windows. The use of forecast orders without predetermined time windows has demonstrated its efficacy in introducing flexibility and capacity to handle a larger volume of orders as shown by Abdollahi et al. (2023). Furthermore, this approach operates on the assumption that forecast orders will eventually be replaced with actual customer orders. Additionally, customers situated in close proximity to the forecast order's location can be incentivised (through pricing strategies) to choose the final time slot allocated by the route-planning software for each forecast order.

However, due to the ability to freely allocate forecast orders across time windows, the resulting distribution of orders might significantly deviate from the original popularity of time slots. This deviation can render the forecast delivery plan unrealistic, especially when employing dynamic pricing strategies with predetermined price limits. To address this concern, this section introduces a novel mechanism that integrates the realistic popularity of time slots into routing considerations while retaining the routing flexibility provided by employing forecast orders without predefined time windows.

4.3.1 Orders CAP for every time slot

Assuming knowledge of the optimal dynamic pricing policy, predicting the slot choices made by customers under this policy becomes relatively straightforward, thus enhancing the accuracy of the final route predictions. However, a significant challenge arises from the fact that the optimal dynamic pricing policy is not known in advance, and is inherently uncertain. In practical scenarios, retailers commonly impose specific limitations on the delivery prices they offer. These limitations can take the form of a defined continuous pricing range, like $[d_s^{min}, d_s^{max}]$, as observed by Koch and Klein (2020). Alternatively, retailers might choose to offer delivery prices from a predetermined set of discrete levels, such as $\{d_s^1, d_s^2, \ldots, d_s^l\}$, as explored in studies by Dong et al. (2009) and Yang and Strauss (2017).

The objective of this section is to determine the upper-bound on the number of orders that can feasibly be accepted for each time slot within the constraints of the choice model, considering the specified price limits. This is referred to as the "order cap" for each time slot s, denoted as CAP_s in subsequent discussions. Also, the $\overrightarrow{CAP}$ encompasses all the CAP_s values for each time slot ($\overrightarrow{CAP} = \{CAP_s | s \in S\}$). This determination allows us to evaluate the degree to which dynamic pricing, operating within specified upper and lower price limits ($[d_s^{min}, d_s^{max}]$), can influence customer demand. For each time slot s, computing CAP_s involves setting the price of slot s to its lowest possible value, d_s^{min}, while assigning the maximum possible price, d_s^{max}, to all other slots:

$$\overrightarrow{Price}(s') = \begin{cases} d_s^{\min} & \text{if } s' = s, \\ d_s^{\max} & \text{otherwise.} \end{cases} \tag{4.3.1}$$

This setup allows for exploring highest potential number of orders for slot s to be chosen, considering the choice model used (in this study, the MNL model with selection probabilities defined by Equations (3.3.7) and (3.3.8)).

In essence, within each simulation run (Sim_i), a complete booking period is emulated using the time-dependent Poisson process to model stochastic customer arrivals. The previously mentioned fixed pricing policy (4.3.1) is applied to the available slots. The process of customer selection is determined through Equations (3.3.7) and (3.3.8). Following this, the number of accepted orders for each time slot s is logged to facilitate the subsequent calculation of the CAP_s value in Algorithm 3.

A series of such simulations, as outlined in Algorithm 3, is performed for every slot s. The value of CAP_s is then calculated as the average number of accepted orders in slot s, rounded up to the nearest integer, by the conclusion of the booking horizon. This entire procedure is reiterated for all slots, with multiple simulations conducted to derive empirical upper-bound for each slot. The vector of resultant upper-bounds ($\overrightarrow{CAP}$) acts as a reference to guide the distribution of forecast orders among time slots.

Algorithm 3: Calculate upper-bound vector ($\overrightarrow{CAP}$) for all slots

```
Input  : Number of runs: nRun;
         Available slots: S;
         Slot price range: [d_s^min, d_s^max]
Output: Set of upper bounds: CAP→.
CAP→ ← [0, 0, . . . , 0] ;    // Initialise upper-bound vector of all slots
for each time slot s ∈ S do
    Initialise all slots' price using (4.3.1)
    #Orders→ ← [0, 0, . . . , 0] ;    // Initialise number of accepted orders for slot s
    for i ← 1 to nRun do
        Run simulation Sim_i for a full booking period ;    // Offering fixed prices from Price→ to all arriving customers, irrespective of their area a and arrival time t.
        #Orders→(i) ← final number of accepted orders in slot s during Sim_i
    CAP_s ← round(avg(#Orders→)) ;    // Average rounded-up orders for slot s
```

4.3.2 Formal statement of CVRPTW with order caps

In Section 4.3.1, the concept of CAP_s for each time slot s is introduced, functioning as a maximum order limit. The objective of incorporating CAP_s is to impose an additional constraint on the maximum allowable number of accepted orders during the solution of the CVRPTW problem. This adjustment addresses the need for a more realistic distribution of forecast orders without time windows, following the implementation of a pricing policy.

A formal formulation of the mathematical model is now presented which is the focus of this study. Within the booking horizon, $\mathcal{A}_t$ is defined as the set of actual orders and $\mathcal{F}_t$ as the set of forecast orders at time step t. The cost of travel from order i to j is represented by C_{ij}. The binary variable X_{ijk} takes a value of 1 when there exists a direct connection between orders i and j using van k on the proposed delivery route plan. l_{ij} denotes the distance between orders i and j. Each van in the problem has a specified capacity denoted by Q, and each order i has a specific demand represented by q_i. Additionally, Φ_{ik} denotes the calculated time at which van k either begins service at order location i or departs from/arrives at the depot (for $i = 0$ and $i = D$ respectively). The problem setting provides a list of pre-defined time windows, each of which has a designated start time denoted by U_s^B and an end time denoted by U_s^E. $[TW_i^B, TW_i^E]$ signifies the agreed-upon time window of order i; for forecast orders, this constraint is relaxed. This modified problem is named as CVRPTW-CAP, and its formulation is as follows.

Constraints (4.3.3) and (4.3.7) enforce the requirement that each order is visited only once and that the van's capacity is not exceeded. The van departs from the depot at location 0, visits the orders, and returns to the depot at location D, as indicated by constraints (4.3.4)-(4.3.6). Inequality (4.3.8) is designed to calculate the arrival times of orders, ensuring that if van k is travelling from location i to location j, it cannot arrive at j before a certain time. This time is calculated as $\Phi_{ik} + b_i + \frac{l_{ij}}{v} - \mathcal{M}(1 - X_{ijk})$. Here, $\mathcal{M}$ represents a large number. The term b_i denotes the service time required at the order location i, and v is the travel speed of the delivery vehicle. Constraint (4.3.8) also eliminates possible sub-tours. The binary variables α_{is} and δ_{is} ensure order allocation within a time slot's start and end limits, respectively. ξ_{is} acts as an overall indicator confirming correct order placement within these time boundaries. Constraint (4.3.9) guarantees that time windows are respected. Furthermore, constraints (4.3.10), (4.3.11), and (4.3.12) verify whether order i is serviced within time slot s. Constraint

(4.3.13) counts the number of orders in slot s and ensures that it does not exceed the CAP_s.

$$
\begin{aligned}
\min \quad & \sum_{i\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{0\}}\ \sum_{j\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{D\}}\ \sum_{k} C_{ij}X_{ijk} && && (4.3.2)\\
s.t. \quad & \sum_{j\in\mathcal{A}_t,\mathcal{F}_t,\{D\}}\ \sum_{k\in V} X_{ijk} = 1, && \forall i\in\mathcal{A}_t\cup\mathcal{F}_t && (4.3.3)\\
& \sum_{j\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{D\}} X_{0jk} = 1, && \forall k\in V && (4.3.4)\\
& \sum_{i\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{0\}} X_{iDk} = 1, && \forall k\in V && (4.3.5)\\
& \sum_{i\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{0\}} X_{igk} = \sum_{j\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{D\}} X_{gjk}, && \forall g\in\mathcal{A}_t\cup\mathcal{F}_t, \forall k\in V && (4.3.6)\\
& \sum_{i\in\mathcal{A}_t\cup\mathcal{F}_t} q_i \sum_{j\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{D\}} X_{ijk} \le Q, && \forall k\in V && (4.3.7)\\
& \Phi_{ik} + b_i + \frac{l_{ij}}{v} - \mathcal{M}(1-X_{ijk}) \le \Phi_{jk}, && \forall i\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{0\}, \forall j\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{D\}, \forall k\in V && (4.3.8)\\
& TW_i^B \le \Phi_{ik} \le TW_i^E, && \forall i\in\mathcal{A}_t, \forall k\in V && (4.3.9)\\
& U_s^B \ge \Phi_{ik} - \mathcal{M}\delta_{is}, && \forall i\in\mathcal{A}_t\cup\mathcal{F}_t, \forall k\in V, \forall s\in S && (4.3.10)\\
& U_s^E \le \Phi_{ik} + \mathcal{M}\alpha_{is}, && \forall i\in\mathcal{A}_t\cup\mathcal{F}_t, \forall k\in V, \forall s\in S && (4.3.11)\\
& \xi_{is} \ge \alpha_{is} + \delta_{is} - 1, && \forall i\in\mathcal{A}_t\cup\mathcal{F}_t, \forall s\in S && (4.3.12)\\
& \sum_{i\in\mathcal{A}_t\cup\mathcal{F}_t} \xi_{is} \le CAP_s, && \forall s\in S && (4.3.13)\\
& X_{ijk}\in\{0,1\}, && \forall i\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{0\}, \forall j\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{D\}, \forall k\in V && (4.3.14)\\
& \alpha_{is}\in\{0,1\}, && \forall i\in\mathcal{A}_t\cup\mathcal{F}_t, \forall s\in S && (4.3.15)\\
& \delta_{is}\in\{0,1\}, && \forall i\in\mathcal{A}_t\cup\mathcal{F}_t, \forall s\in S && (4.3.16)\\
& \xi_{is}\in\{0,1\}, && \forall i\in\mathcal{A}_t\cup\mathcal{F}_t, \forall s\in S && (4.3.17)\\
& \Phi_{ik} \ge 0, && \forall i\in\mathcal{A}_t\cup\mathcal{F}_t\cup\{0\}\cup\{D\}, \forall k\in V. && (4.3.18)
\end{aligned}
$$

Based on the provided formulation, it becomes apparent that this does not strictly fit the definition of a CVRPTW problem. This is due to the fact that the study introduces constraints solely on the maximum order count within each time slot, without directly assigning time slots to specific forecast orders. Since forecast orders lack specific time windows, they have the flexibility to be placed in any available slot based on their respective locations.

As the formal mathematical model is not computationally tractable due to the large state space X and the fact that computing minimum cost solution $C(X)$ alone is intractable since it requires solving a large vehicle routing problem with time windows, it subsequently will be addressed through an approximate solution using a Simulated Annealing (SA) algorithm. The proposed CVRPTW-CAP is NP-hard due to the exponential increase in possible routes and schedules with the number of customers, vehicles, and constraints, which are characteristics of NP-hard problems.. Consequently, the problem must be solved iteratively, upon every update of sets $\mathcal{A}_t$ and $\mathcal{F}_t$ with the acceptance of new orders. Due its computational complexity, heuristic methods will be employed to find a near-optimal solution for practical purposes.

4.3.3 Simulated annealing (SA) for CVRPTW-CAP

To efficiently solve the CVRPTW-CAP, Simulated Annealing (SA) as a heuristic has been employed. SA is used to find the initial (sub-)optimal routes with forecast orders, and is called upon the acceptance of every actual order to re-optimise the route. To impose the order limit CAP_s as specified by constraint (4.3.13), a forecast-order distribution step is introduced using Algorithm 4 into SA. During the operation of SA, when generating a random tentative route plan, Algorithm 4 (re-)assigns the earliest arrival time to each order in the route plan, considering the CAP_s values. Note that the earliest arrival time ($\vec{E}$) signifies the earliest moment when a delivery can be made to a customer, according to the current working route plan, ensuring that it occurs after this point. On the contrary, the latest arrival time ($\vec{L}$) denotes the latest allowable time by which the delivery must be completed, ensuring it happens no later than this specified time. For a given set of orders, each order is associated with a designated van number and defined $\vec{E}$ and $\vec{L}$.

In case when constraint (4.3.13) cannot be met after (re-)assigning the visiting times, penalty terms are added to the objective to push the SA search to feasible directions. In Section 4.3.3.1, the employed forecast order distribution approach is explained in details. Figure 4.2 outlines the integrated dynamic routing and demand management problem as introduced in Chapter 3, along with the integration of distribution mechanisms into the dynamic routing system to refine the estimation of the opportunity cost (OC).

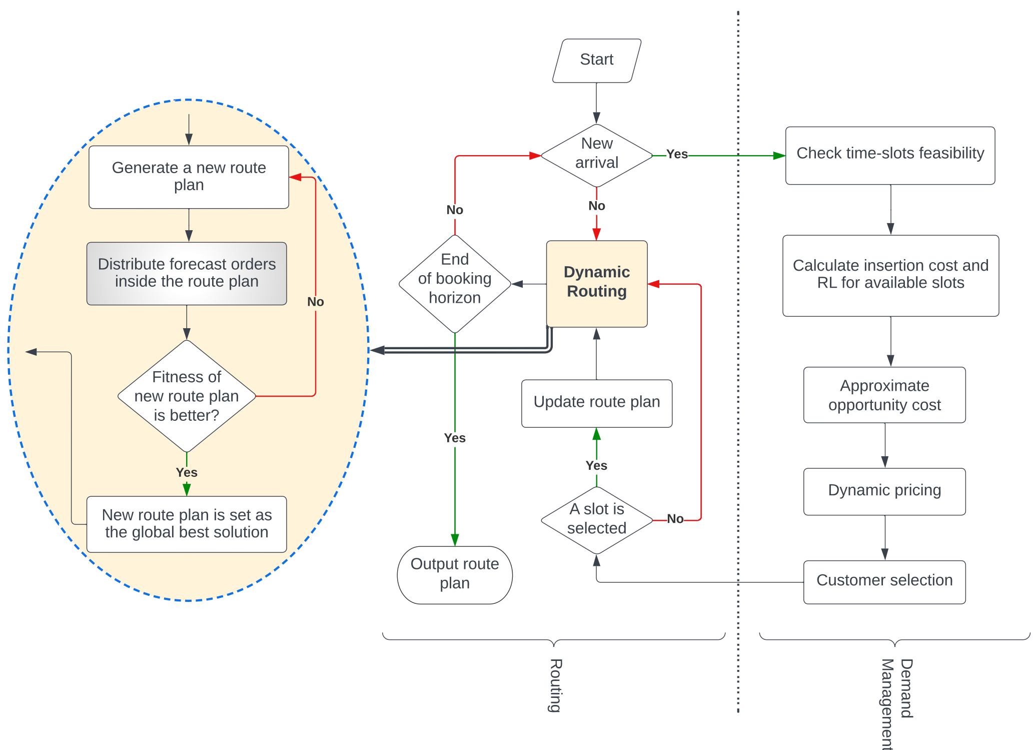

Figure 4.2: Incorporation of distribution mechanisms into the dynamic routing system as detailed in Section 4.3.

4.3.3.1 Forecast order distribution of a route plan

Recall that CAP_s represents the maximum number of orders expected in slot s, to be serviced by multiple delivery vans. Nevertheless, during the distribution of orders across slots in the SA process, the route of each delivery van is addressed sequentially. To achieve a more balanced distribution across all delivery vans, a refined CAP_s per van, denoted by $CAPpV_s(k)$, is introduced to reflect the slot-specific order limit in the route of delivery van k, based on the number of orders currently assigned to that particular van:

$$CAPpV_s(k) = \lceil CAP_s \cdot \frac{\text{length}(\overrightarrow{order}(k))}{\sum_s CAP_s} \rceil \tag{4.3.19}$$

where length$(\overrightarrow{order}(k))$ represents the number of orders for van k at time step t. Ceiling function $\lceil \cdot \rceil$ is applied to ensure that $CAPpV_s(k)$ takes integer values. $CAPpV_s(k)$ will not exceed CAP_s because the length$(\overrightarrow{order}(k))$ is never higher than $\sum_s CAP_s$. To clarify, CAP_s represents the upper-bound for the total orders that can be delivered in slot s, considering multiple vans. However, $\sum_s CAP_s$ is considerably higher than the actual number of orders that will be received at the end of booking period. $CAPpV_s(k)$ is introduced, which mirrors the structure of CAP_s while tailoring it to van k. This approach captures the highest expected order count for slot s on van k, adhering to the same pattern as in CAP_s, contingent on the current order count.

Algorithm 4 employs Equation (4.3.19) to guide re-assignment of earliest arrival time of orders denoted as $\vec{E}_b$ on a given route plan in SA. Delving into further details, when the current time slot approaches its capacity threshold ($CAPpV_s$), Algorithm 4 assesses the feasibility of shifting the next forecast order to the adjacent time slot. This process is outlined within lines 9 to 18 of the algorithm.

Algorithm 4: Distribution of forecast orders in time slots based on $\overrightarrow{CAP}$

Input : For each van k: order sequence $\overrightarrow{order}(k)$;

Earliest and latest arrival times for each van k and order i: $\vec{E}(k,i)$ and $\vec{L}(k,i)$;

Slots upper-bound: $\overrightarrow{CAP}$;

Available slots: S;

Start time of all slots: $\vec{U}^B$

Output: Balanced earliest arrival times $\vec{E}_b$.

Initialise $\vec{E}_b \leftarrow \vec{E}$
for $k \leftarrow 1$ *to number of vans* **do**
  Initialise $RS \leftarrow [0, 0, \ldots, 0]$; `// Record number of orders in each slot`
  Calculate refined $CAPpV_s(k) \leftarrow \lceil CAP_s \cdot \frac{\text{length}(\overrightarrow{order}(k))}{\sum_s CAP_s} \rceil$; `// Refined` CAP_s `for van` k
  for $i \leftarrow 1$ *to length*$(\overrightarrow{order}(k))$ **do**
    $s \leftarrow$ find_slot_number$(\vec{E}_b(k,i))$; `// Find the corresponding slot number for` $\vec{E}_b(k,i)$
    if i *is a forecast order* **then**
      if $RS(s) > CAPpV_s(k)$ **then**
        if s *is not the last slot in* S **then**
          for $n \leftarrow i+1$ *to length*$(\overrightarrow{order}(k))$ **do**
            $\text{Slack}_\text{n} \leftarrow \vec{L}(k,n) - \vec{E}_b(k,n)$; `// Calculate slack of subsequent orders`
          $\text{Min\_slack} \leftarrow \min_{i<n\leq \text{length}(\text{order}(k))} \text{Slack}_\text{n}$
          $\text{Shift} \leftarrow U^B_{s+1} - \vec{E}_b(k,i)$; `// Calculate required shift size`
          if Shift $\leq$ Min_slack **then**
            $RS(s) \leftarrow RS(s) - 1$; `// Remove order` i `from slot` s
            $RS(s+1) \leftarrow RS(s+1) + 1$; `// Add order` i `to slot` $s+1$
            for $h \leftarrow i+1$ *to length*$(\overrightarrow{order}(k))$ **do**
              $\vec{E}_b(k,h) \leftarrow \vec{E}_b(k,h) + \text{Shift}$; `// Delay following orders by Shift`

The feasibility of this shifting operation hinges on the flexibility of subsequent actual or forecast orders, which is determined by the range between their earliest and latest arrival times. In line 13, the required shift size is calculated by determining the time difference between the beginning of the next time slot and the earliest time at which the forecast order in consideration could be accommodated.

Subsequently, the algorithm evaluates the flexibility of all subsequent orders. If it is feasible to shift these orders by at least the calculated the required shift size, as indicated in lines 14 to 18, then the current shifting of the forecast order is deemed valid. However, if the shifting is infeasible due to the presence of actual orders, the earliest arrival time of the forecast order remains unaltered.

Towards the end of the algorithm, after examining all delivery vans for order distribution, if the total number of orders in a specific time slot violates the pre-defined CAP_s constraint, a penalty is introduced into the objective function. The assumption is that the index ordering used for labelling time slots s aligns with the actual delivery times over the day.

4.4 Augmented Time Window

The study in the previous chapter presented a partially time-windowed dynamic routing technique using forecast orders to assist the dynamic routing system in recommending more lucrative time slots for new actual orders. This method involves generating preliminary routes based on forecast orders, which are initially devoid of specific time windows. As the booking horizon progresses, these forecast orders play a pivotal role in guiding the system towards identifying favourable time slots for the inclusion of new, actual orders.

The process unfolds as follows: when a new actual order is received, the system evaluates the forecast orders and selects one for removal. Subsequently, the dynamic routing algorithm is invoked to reconfigure the delivery plan. With this replacement, the system transitions from having a forecast order without a predetermined time window to gaining an actual order with a fixed, one-hour time window. However, it is important to recognise that an order without a designated time window offers more flexibility in route optimisation, whereas fixed time windows tend to constrain this flexibility.

To address this diminishing flexibility concern, the present study proposes a solution aimed at encouraging customers to choose extended-duration time windows. By incentivising customers to opt for these broader time slots, the system can maintain and even enhance its ability to accommodate actual orders within flexible routing plans. This approach attempts to strike a balance between customer preferences and the operational efficiency of the dynamic routing system.

The existing literature on variable or extended-duration time windows in the context of the AHD problem has mainly explored two strategies: static short and long time slots (as shown by Klein et al. (2020)), and the combination of less popular slots with more popular ones to improve slot utility (as discussed by Strauss et al. (2020)). Differing from these methods, this study introduces a dynamic approach to slot combination, building upon the foundation laid in Chapter 3. This approach leverages the flexibility inherent in forecast orders without fixed time windows.

In this approach, optimal extensions for the existing set of Standard Time Windows (STW) is determined every time a new customer places an order. These extensions, termed Augmented Time Windows (ATW), are tailored to optimise the tentative route by considering both actual and forecast orders. This dynamic integration of ATW stands apart from the traditional static approach of merging time windows. Figure 4.3 depicts the integration of dynamic routing and demand management, including forecast order distribution into the dynamic routing system (left oval) and ATW calculation, alongside the existing STW to calculate OC (right oval). This integration results in an improved effectiveness of the routing process.

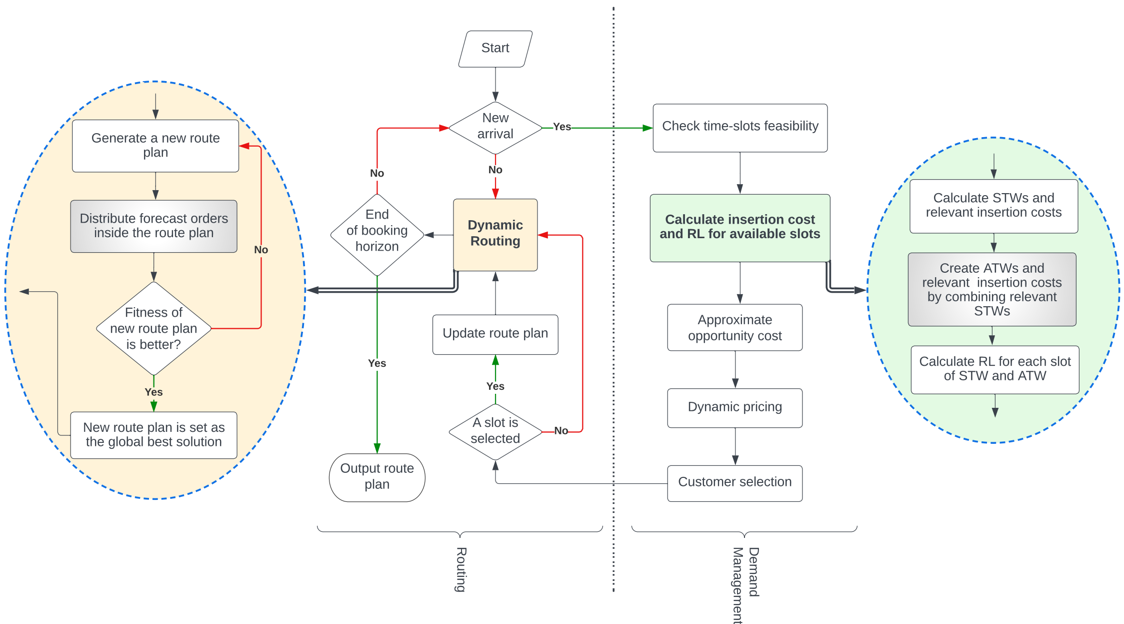

Figure 4.3: Creation of ATW using the existing STW to enhance the efficiency of the routing process as shown in the right oval and discussed in Section 4.4).

Furthermore, Figure 4.4 provides a clear example of how ATWs are formed from adjacent STWs. Let us consider a scenario where a new actual order j is inserted into the tentative route during the dynamic routing process. In Figure 4.4, j_s^* represents the index of the most appropriate forecast order to remove when inserting order j into slot s. Here, i and $i+1$ represent the indices of the preceding and succeeding orders on the route, respectively, between which order j is being inserted. The marginal cost of replacing j_s^* with j is denoted as IC_s. For further clarification on the values in Figure 4.4, let us consider that for slots 1 to 3, removing forecast order number 39—currently the second order in van 1—is optimal. The most suitable position for the incoming actual order is between orders 3 and 4 in van 2, which are indexed 45 and 58, respectively, with an insertion cost of 13.42.

It is evident that if the new order j is aimed to be serviced in slots 1, 2, or 3, the best forecast order to remove and the optimal insertion position on the route remain the same. This implies that even after this replacement and insertion, actual order j retains the flexibility to move across these three slots. However, if only STWs are available and the customer later selects slot 2, for instance, order j becomes constrained by the start and end time of slot 2 (i.e., [7 am, 8 am]). Consequently, this reduces the route's flexibility as any future new orders to be serviced in slot 3 cannot be inserted into gaps before order j, even if travel time allows both orders to be accommodated within a one-hour slot.

To maintain routing flexibility, slots 1, 2, and 3 are consolidated into an ATW that spans from 6 am to 9 am, as highlighted in green in Figure 4.4. If selected, order j can be fit into any time between [6am, 9am], making the tentative route more adaptable for absorbing a higher number of future orders. This strategy helps enhance the overall efficiency and responsiveness of the routing system.

	STW										*ATW*		
slot	1	2	3	4	5	6	7	8	9	10	11	12	13
j*	39	39	39	20	48	48	11	6	6	6	39	48	6
i	45	45	45	5	24	24	20	44	44	44	45	24	44
i+1	58	58	58	61	22	22	41	29	29	29	58	22	29
IC	13.42	13.42	13.42	12.49	6.80	6.80	4.51	-3.03	-3.03	-3.03	13.42	6.80	-3.03
$begin_s$	6 am	7 am	8 am	9 am	10 am	11 am	12 pm	1 pm	2 pm	3 pm	6 am	10 am	1 pm
end_s	7 am	8 am	9 am	10 am	11 am	12 pm	1 pm	2 pm	3 pm	4 pm	9 am	12 pm	4 pm

Figure 4.4: Constructing augmented time windows from standard time windows.

To generate the ATW set from neighbouring slots in the STW set, the indices of preceding and succeeding orders (i and $i+1$) are examined during the order insertion or replacement process. The time slots in STW that have matching values for i and $i+1$ are identified as the desired time slots to combine and form a new extended time window to be added to ATW. Notably, since ATW are generated from adjacent time slots, their inclusion in subsequent routing (re-)optimisation does not burden the dynamic routing package with maintaining disconnected aggregated time windows. Instead, an order selecting ATW simply becomes a standard order in the follow-up routing (re-)optimisation with a longer time window.

One final note concerns the utility of ATW. It was initially assumed that no utility decrement would occur when combining slots; that is, the utility of ATW is equal to the average utility of all individual slots involved. However, in Section 4.6.4, the updated MNL model that incorporates ATW is extensively explored, evaluating different utility decrement bands and their impact on the performance of the method.

4.5 Calculate Optimal Delivery Charges

Building upon the dynamic routing methodology discussed in Chapter 3, optimal delivery charges will be calculated in this chapter. After aligning the earliest arrival times of all orders ($\vec{E_b}$) using Algorithm 4 from Section 4.3, we calculate the revenue loss (RL). It is crucial to note that the number of time slots now encompasses both scheduled time windows (STW) and actual time windows (ATW), a key distinction from the methodologies described in Chapter 3.

4.5.1 Extending Revenue Loss Calculations

Introduced in Chapter 3, the concept of calculating RL is enhanced in this chapter to include both actual and forecast orders remaining in the system at any given time t. This extension is important for assessing the monetary value of time consumed by new orders in various time slots.

$$eRL_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) = \sum_{s' \in S = \{STW \cup ATW\}} \theta_{t,s'} \left(w_{s'}(\vec{x}_t, \vec{f}_t) - w_{s'}(\vec{x}_t + \mathbf{1}_{as}, \vec{f}_t - j_s^*) \right) \quad (4.5.1)$$

Here, $\theta_{t,s'} \in \mathbb{R}$ denotes the monetary value associated with the time consumed by the new order 1_{as} in time slot s'.

$\theta_{t,s'}$ is computed using the best route data generated by the dynamic routing package at time t:

$$\theta_{t,s'} = \frac{\sum_i \{r_i | i \in (\vec{x}_t \cup \vec{f}_t),\ U_{s'}^B \leq \Phi_i \leq U_{s'}^E\}}{U_{s'}^E - U_{s'}^B} \quad (4.5.2)$$

This calculation reflects a more comprehensive scenario than previously modelled, considering both actual and forecast orders.

4.5.2 Calculation of Insertion Costs

Following Equation (3.3.2) from Chapter 3, we calculate insertion costs (IC) for adding a new delivery. If it becomes infeasible to fulfil all orders $(\vec{x}_t, \vec{f}_t)$ due to capacity and time constraints, the delivery cost $DC_t(\vec{x}_t, \vec{f}_t)$ is considered infinite.

4.5.3 Opportunity Cost Estimation

The extended revenue loss eRL and insertion cost IC are combined to estimate the opportunity cost:

$$OC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) = IC_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) + eRL_t(\vec{x}_t, \vec{f}_t, \mathbf{1}_{as}) \quad (4.5.3)$$

This metric approximates the term $(V_{t+1}(\vec{x}_t) - V_{t+1}(\vec{x}_t + \mathbf{1}_{as}))$ from Equation (3.2.1), essential for determining the optimal slot pricing.

Using the derived opportunity cost, we calculate the optimal slot price d_s^* in accordance with the dynamic pricing model advocated by Dong et al. (2009), previously applied in Chapter 3 and delineated in Equations (3.3.9) and (3.3.10).

4.6 Experimental Results

In the forthcoming section, an outline of the study areas and the methodologies used to compare the proposed method with several benchmarks will be provided. Results obtained from diverse experiments will be presented, evaluated against various measures. Additionally, a detailed explanation concerning the implementation aspects related to the distribution of forecast orders will be given. The potential of augmented time windows to enhance the effectiveness of the methods will also be explored. Performance improvements achieved through this approach will be illustrated and compared to scenarios exclusively using standard time windows.

4.6.1 Data Specification

To examine the impact of forecast order distribution and augmented time windows, experiments were conducted across seven distinct delivery areas. These areas were chosen to represent different customer densities and distribution patterns. The experiments were carried out using actual customer locations and historical booking data, while ensuring the protection of commercial information and customer privacy through appropriate measures. Table 4.1 provides a comprehensive overview of the selected areas, including their specific characteristics and the number of vans deployed to fulfil the final accepted delivery orders.

Table 4.1: Features of seven areas under study

Area	1	2	3	4	5	6	7
Geographical spread	Rural	Suburb	Semi-Rural	Semi-Rural	Suburb	City	City
Average number of customers per day	63	130	206	308	398	502	611
Average distance between customers (mile)	11.78	5.87	13.43	11.98	8.81	2.67	2.32
Number of delivery vehicles	2	4	6	9	11	14	16
Type of vehicle	Van	Van	Van	Van	Van	Van	Van

Chapter 4 builds upon the dynamic pricing methods introduced in Chapter 3 which are Static Pricing method (SP), Dynamic Pricing with Insertion Cost without employing forecast order (DP-IC), Dynamic Pricing with Fixed routing of Forecast order with time window (DP-FR-F), and Dynamic Pricing with Fixed routing of Forecast order with time window (DP-FR-F) extending them with two new approaches introduced in this chapter:

- **Dynamic Pricing with Dynamic Routing of Distributed Forecast orders without time window (DP-DR-DF)**: This method considers dynamically adjusting the pricing based on forecast route plan entailing distributed forecast orders; forecast route is updated upon new order acceptance. This is the approach proposed in Section 4.3.
- **Dynamic Pricing with Dynamic Routing of Distributed Forecast orders without time windows entailing Augmented Time Windows (DP-DR-DF-ATW)**: This approach is same as DP-DR-F, but with the inclusion of augmented time windows as proposed in Section 4.4.

For a detailed breakdown of the components of each method, please refer to Table 4.2.

Table 4.2: Comparison of different methods based on key features

Method	Pricing policy	OC	Route planning	Forecast order type	Time window type
SP	Static	-	-	-	standard
DP-IC	Dynamic	IC	-	-	standard
DP-FR-F	Dynamic	IC	Fixed forecast orders	With TW	standard
DP-DR-F	Dynamic	IC + RL	Dynamic forecast orders	Without TW	standard
DP-DR-DF	Dynamic	IC + RL	Distributed dynamic forecast orders	Without TW	standard
DP-DR-DF-ATW	Dynamic	IC + eRL	Distributed dynamic forecast orders	Without TW	standard & augmented

Lastly, the experiments were conducted using MATLAB on a high-performance computing system equipped with an Intel Core i9-7940X 3.1GHz processor. To ensure reliable results, a total of 30 independent runs of the experiments is performed and the average is reported.

4.6.2 Comparative analysis of the methods based on various metrics

To evaluate the effectiveness of the new approaches introduced earlier, namely the employment of distributed forecast order and dynamically extending feasible time windows through the implementation of augmented time windows, various measures have been considered. Among these metrics, profit growth stands out as a crucial benchmark to assess the performance of the methods. Figure 4.5 presents the achieved profit for all the methods across different areas, along with the percentage improvement in profit compared to the static pricing method used as the baseline. This visualisation provides a comprehensive overview of how each method performs in terms of generating profit and highlights the extent to which they outperform the static pricing approach.

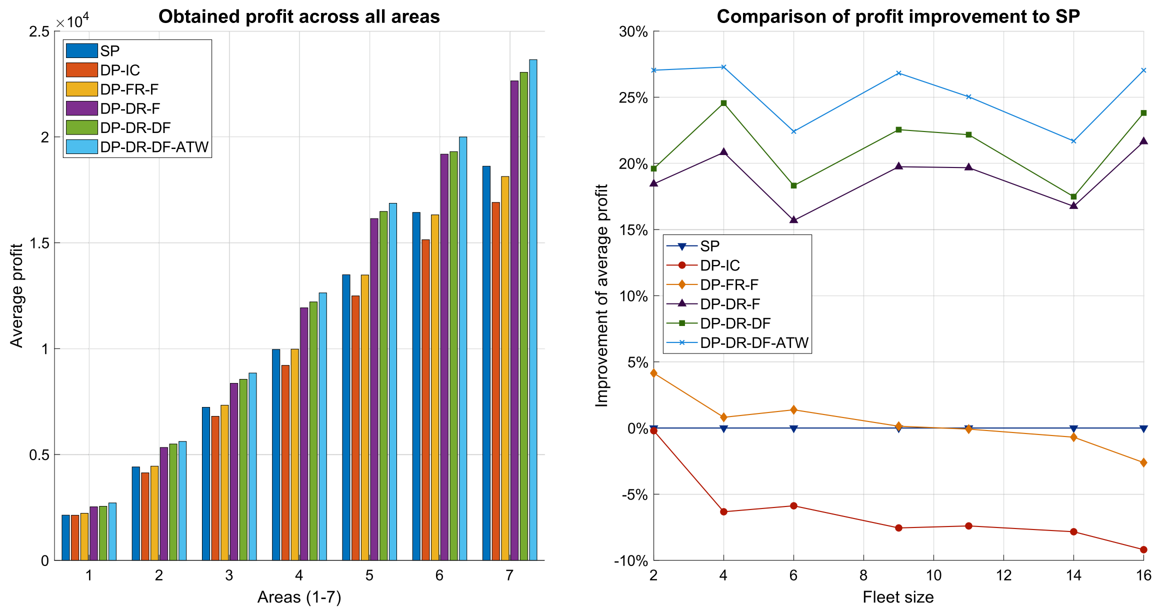

Figure 4.5: Comparison of the profit growth of the studied methods. Fleet sizes in Areas 1 to 7 correspond to 2 to 16, as depicted in the figure.

Analysing Figure 4.5, it is evident that DP-DR-DF outperforms DP-DR-F, highlighting the significance of distributing forecast orders. This observation emphasises the importance of accurately estimating revenue loss, as it enables a more realistic calculation of opportunity costs. By considering and accurately estimating revenue loss, businesses can make more informed decisions and achieve better financial outcomes by optimising their pricing policies. For a more detailed investigation into the distribution of forecast orders and the estimation of time budget as a component of revenue loss calculation, please refer to Section 4.6.3.

The integration of augmented time windows into DP-DR-DF to create DP-DR-DF-ATW leads to even more promising results, with a notable profit improvement ranging from 22% up to 27% compared to the static pricing method. This outcome underscores the effectiveness of augmented time windows in guiding customers towards time slots that have a higher potential to generate increased overall profit throughout the entire booking period. By offering customers more versatile choices in terms of time window length along with discounted delivery charges, the routing package can more efficiently manage accepted orders. This helps to mitigate the biased random selection of customers towards popular time windows. As a result, there is a higher slot availability throughout the booking period, increasing the chances of customers successfully booking a desired time slot. Note that the results as displayed in Figure 4.5 for DP-DR-DF-ATW assumes no utility decrement when combining slots, i.e., utility of augmented slots is equal to the average utility of all individual slots involved. It is acknowledged that this assumption may not always hold true in practice. In Section 4.6.4, the updated MNL model that incorporates augmented time windows is extensively explored, evaluating different utility decrement bands and their impact on the performance of the DP-DR-DF-ATW method.

Another advantage of the methodology proposed in Section 4.4 for augmented time windows (ATW) is its compatibility with various methods and approaches in dynamic pricing. This approach can seamlessly function as an add-on to any insertion-based dynamic routing system, facilitating profit improvement without requiring modifications to the specific method for dynamic pricing. This flexibility allows for easy integration into different methods with diverse approaches. Furthermore, the inclusion of ATW does not impose significant computational burdens on the methods. Instead, it simply redefines the standard time window based on the inherent characteristics of each method. Consequently, the integration of augmented time windows contributes to improved performance without introducing excessive computational complexities, as it involves working with and combining pre-calculated standard time windows.

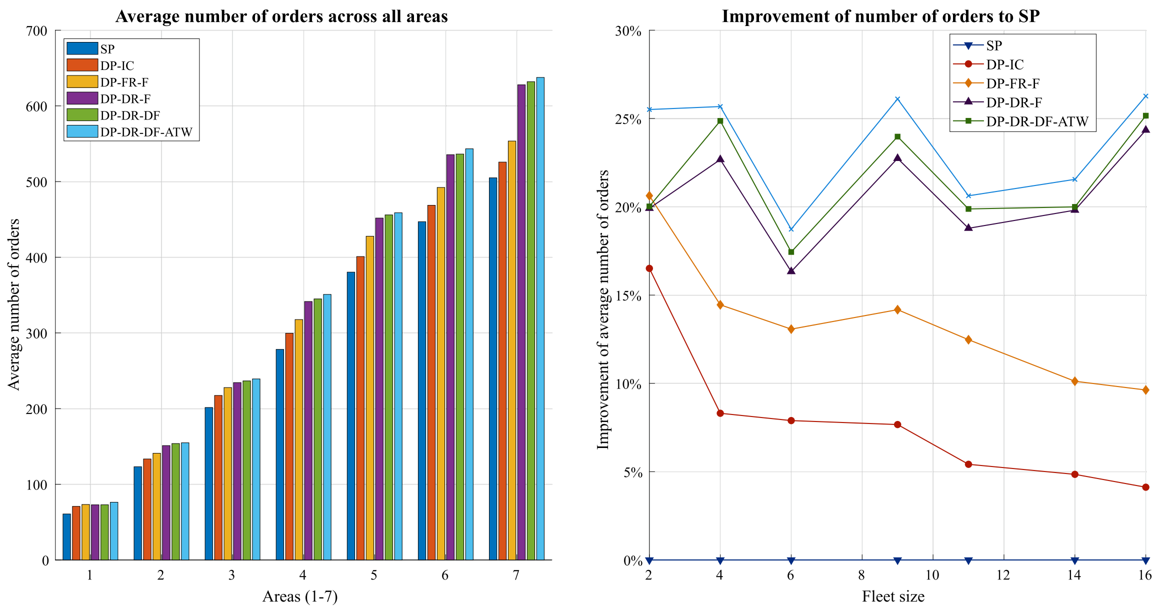

Figure 4.6: Comparison of the acceptance rate of the studied methods on the number of accepted orders. Fleet sizes in Areas 1 to 7 correspond to 2 to 16, as depicted in the figure.

Figures 4.6 and 4.7 highlight the improved efficiency the method brings to the dynamic routing system. Figure 4.6 showcases the average number of accepted orders across all areas. This figure highlights the effectiveness of new approaches in accommodating a greater number of orders within the same problem setting. By optimising the routing process, the DP-DR-DF-ATW excels in incorporating more orders into the routes, thereby maximising resource employment and meeting customer demand effectively. In addition, the DP-DR-DF approach achieves a slight increase in accepted orders while generating more profit through enhanced pricing of time slots and more accurate RL estimation.

In Figure 4.7, a graphical comparison is presented, focusing on the average travelling cost required to serve an order and the average profit gain per mile in all areas. These metrics directly reflect the efficiency of the final routes generated by the methods and their impact on the overall performance of the dynamic routing system. The figure underscores the effectiveness of new approaches in achieving cost-minimisation objectives while maximising profit gains per mile. This emphasises the importance of an efficient forecast route plan, enabling intelligent and strategic travel decisions, ultimately contributing to enhanced profitability.

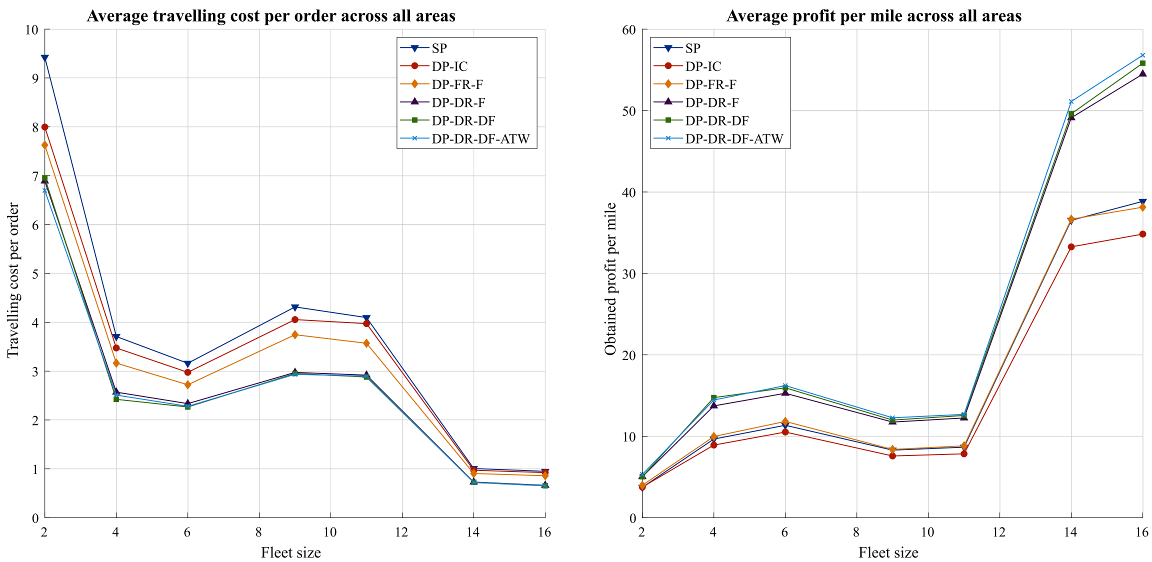

Figure 4.7: Performance of the studied methods with regards to travelling cost and profit gain

The availability of time slots and their corresponding prices over the booking period significantly impact the total profit. Figure 4.8 provides insights into the number of available time slots and the average advised price for representative areas (1 and 7). When a new order arrives, there exists a trade-off between immediate gain and long-term profit. This trade-off explains why the prices offered by DP-IC and DP-FR-F approaches are consistently lower than those of DP-DR-F, DP-DR-DF, and DP-DR-DF-ATW. Note that the DP-IC and DP-FR-F approaches, lacking consideration for expected revenue loss in the opportunity-cost estimation, primarily focus on the immediate gain associated with the order under consideration. To entice the customer to make a purchase, these approaches lower the price offered significantly. However, in the case of DP-DR-F, DP-DR-DF, and DP-DR-DF-ATW approaches, which anticipate future orders, there is less pressure to secure an immediate order. Consequently, the prices offered tend to be higher on average. The variation in prices across different approaches underscores the significance of incorporating an accurate assessment of revenue loss in the opportunity-cost estimation. This consideration enables approaches to strike a balance between immediate gains and long-term profitability, resulting in more strategic pricing decisions and ultimately maximising the total profit.

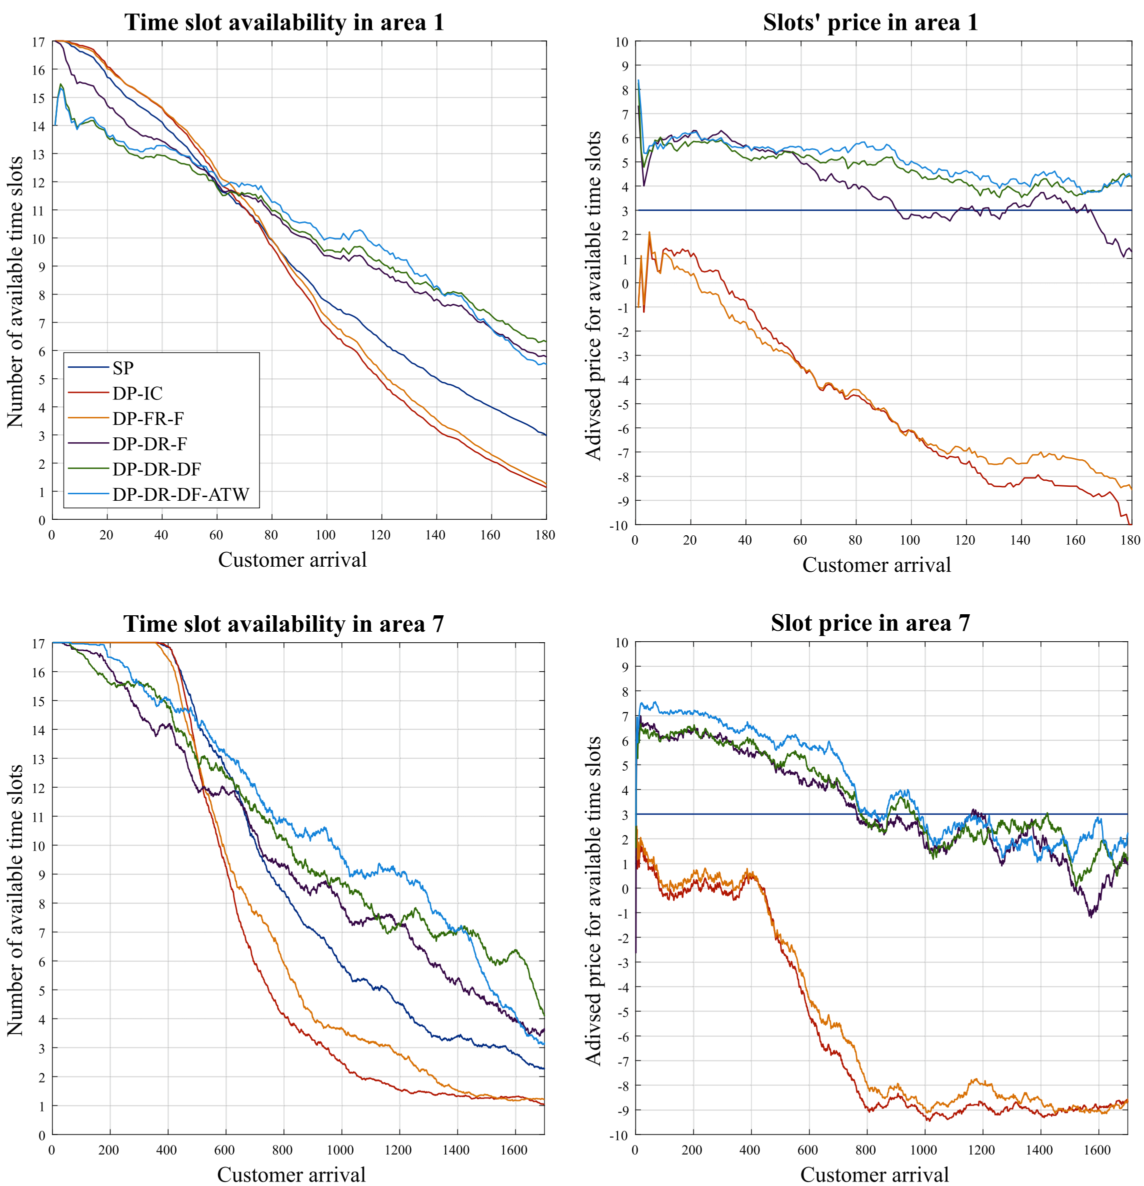

Figure 4.8: Comparison of time slot availability and pricing: areas 1 and 7

When comparing DP-DR-DF and DP-DR-DF-ATW with DP-DR-F, it is evident that the prices charged by DP-DR-F are generally lower. This is because in DP-DR-F, the best time slot is determined solely based on the dynamic routing system without considering time windows or the original popularity of time slots. Consequently, there is a higher pressure to lower the price in order to persuade a customer to book an undesirable time slot if their location is deemed the best fit for that particular slot. On the other hand, DP-DR-DF and DP-DR-DF-ATW offer more stable delivery prices compared to DP-DR-F by leveraging the benefits of forecast order distribution. Ideally, the pricing policy towards the end of the booking horizon should aim to treat customers fairly by providing consistent prices across all time slots. This equitable approach is particularly evident in low-density areas with lower complexity, where the demand system is less intricate. However, in areas with higher demand, this advantage may be somewhat diminished due to factors such as increased order volume, more delivery vehicles, and greater pressure on the dynamic routing system. Nevertheless, in real-world scenarios, it is common for prices in dynamic pricing schemes to decrease over time as slot availability dwindles. Typically, popular slots tend to be reserved earlier than less popular ones, resulting in lower prices being offered towards the end of the booking period to stimulate bookings for the less popular slots.

4.6.3 The Impact of Forecast Order Distribution on Time Budget

This section aims to demonstrate the impact of distributing forecast orders based on the CAP. Figure 4.9 serves as a visual representation, highlighting how the distribution of forecast orders contributes to aligning the time budget of each slot with the slot popularity predicted by the MNL model. This comparison provides valuable insights into how effectively the CAP-based distribution approach can optimise time slot employment and improve the overall match between customer demand and available time slots.

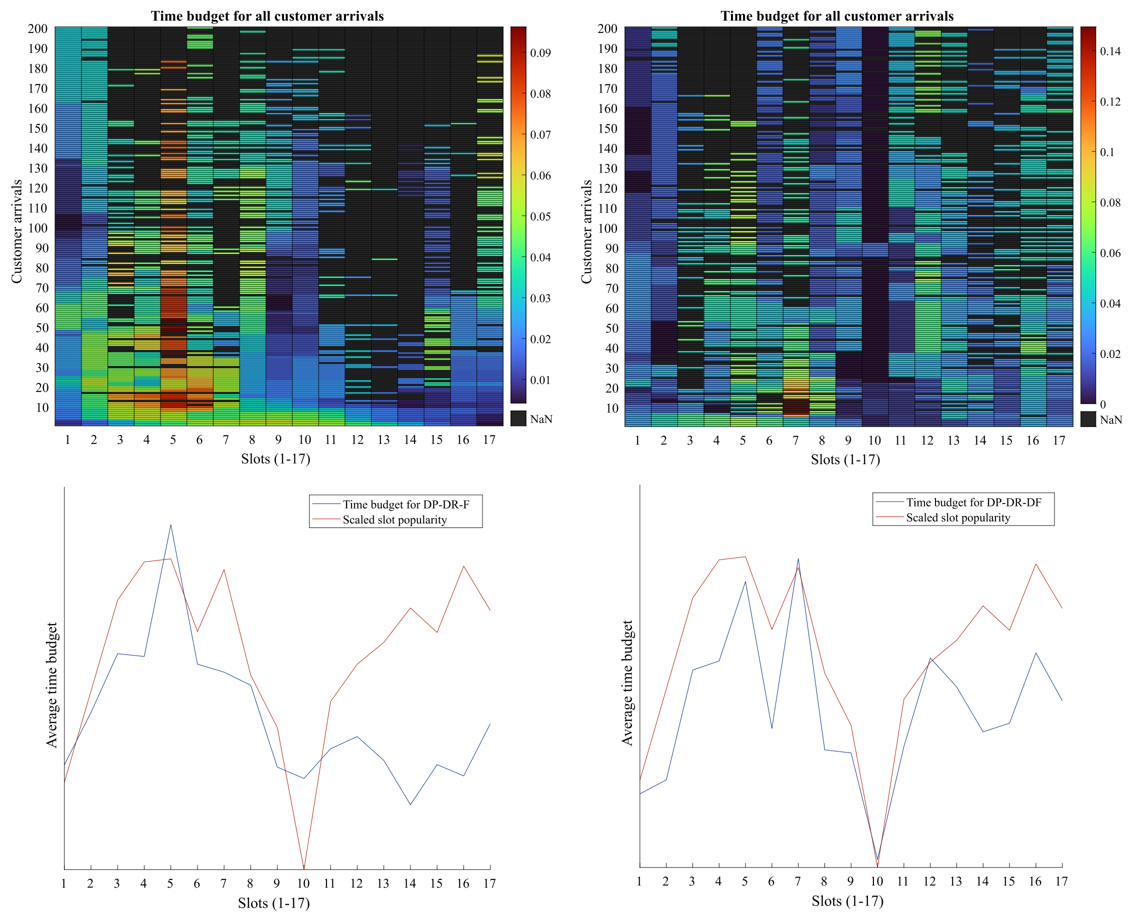

Figure 4.9: Estimated slot popularity based on slots' time budget during booking horizon of Area 1 compared to scaled slot popularity of the MNL choice model. The first column is for DP-DR-F and the second column is for DP-DR-DF.

The left heatmap represents the DP-DR-F method, which does not consider forecast order distribution. It shows that the time budgets for the first half of the time slots have higher values across all customer arrivals. In contrast, the right heatmap, corresponding to the DP-DR-DF method, which integrates forecast order distribution, reveals two distinct peaks in the time slot values. This pattern corresponds to slot popularity as observed in the MNL model. The upper heatmaps, which display the number of customer arrivals over time on the y-axis and the number of time slots (17) on the x-axis, illustrate which time slots received more customers. A hotter colour indicates higher customer acceptance. The top left heatmap shows customer acceptance that is incongruous with the observed slot popularity, whereas the top right heatmap displays a pattern more aligned with slot popularity, indicating higher acceptance in popular slots as represented by hotter colours.

The left heatmap corresponds to the DP-DR-F method, where forecast order distribution is not considered. It reveals that the time budgets for the first half of the time slots exhibit higher values across all customer arrivals. In contrast, the right heatmap representing the DP-DR-DF method, which integrates forecast order distribution, demonstrates two distinct peaks in the time slot values. This pattern aligns with the slot popularity observed in the Multinomial Logit (MNL) model. The top heatmaps which show the number of customer arrivals over time on y axis and the number of time slots (17) on the x axis, illustrate that which time slots received more customers. The hotter the colour the more customer acceptance. The top left heatmap shows incongruous customer acceptance compared to the shown slot popularity while the top right heatmap show more resemblance to slot popularity as it shows higher acceptance in popular slots shown be hotter colours.

The line graph positioned below each heatmap illustrates the average time budget across all customer arrivals, juxtaposed with the scaled slot popularity. Notably, the incorporation of forecast order distribution in the DP-DR-DF method leads to a more pronounced alignment between the slot popularity and the time budget, highlighting the impact of this approach on achieving a closer match between predicted popularity and actual time allocations.

4.6.4 Incentives on the ATW by calculation of new MNL parameters corresponding to ATW

This section focuses on adjusting the MNL choice model by incorporating the ATW parameters in addition to the previously known STW parameters. As mentioned in Section 4.2, the probability Equations of selecting ($P_s(\vec{d})$) and not-selecting ($P_0(\vec{d})$) an offered time slot s under the delivery charge vector $\vec{d}$ upon each customer's arrival are governed by Equations (3.3.7) and (3.3.8), respectively. The β_s parameter in these Equations represents slot utility capturing customer preferences when selecting a time slot. Since augmented time windows have not been implemented before, there is no direct historical data available to recover their β_s values. Therefore, the MNL parameters need to be drawn, specifically the β_s parameter originally defined for STW, for ATW. This involves using Equation (4.6.1) to calculate the β_s values for the ATW slot s'. It averages all the β_s values for the constituent slots that form s' and then decreases this value by a factor related to the length of s'. More specifically, if s'_i represents the augmented time window which is a combination of standard time windows, $\{stw_1, stw_2, \ldots, stw_{n_i}\}$, the following formula is used to calculate the β value for s'_i:

$$\beta_{s'_i} = \frac{\sum_{j=1}^{n_i} \beta_{stw_j}}{n_i} - ((n_i - 1) \times \gamma) \tag{4.6.1}$$

Since an augmented time window has a broader time range and requires customers to wait longer at the delivery location compared to a standard time window, its popularity among customers is expected to decline. Equation (4.6.1) is constructed based on the assumption that the utility of an augmented slot decreases linearly with its length, and the rate of reduction is captured by the parameter γ. During experiments, different values of γ are tested to examine the influence of combining time slots.

To determine suitable scales for γ values, γ can be interpreted in terms of monetary units. Recall that β_d represents the utility parameter associated with delivery price. Consequently, as the price per unit increases, the utility will decrease by an amount proportional to $|\beta_d|$ (where β_d is a negative value). For example, if extending a time slot by one hour leads to a utility reduction equivalent to increasing its price by £Z, then γ can be set as $Z \cdot \beta_d$.

In this research, a range of γ values spanning from 0 to 0.2 is considered, approximately corresponding to a price increment of £0 to £2.6. The summary of profit gains is visualised in Figure 4.10, which highlights the performance of the DP-DR-DF-ATW method across various γ bands in comparison to other approaches.

Analysing Figure 4.10, it is observed that areas with lower customer density exhibit higher profit gains when using the DP-DR-DF-ATW method; areas with a higher volume of customer requests show a lower profit gain. Nevertheless, the improvement in profit remains substantial. This indicates the efficacy of applying ATW in exploring and exploiting customers' flexibility in time slots.

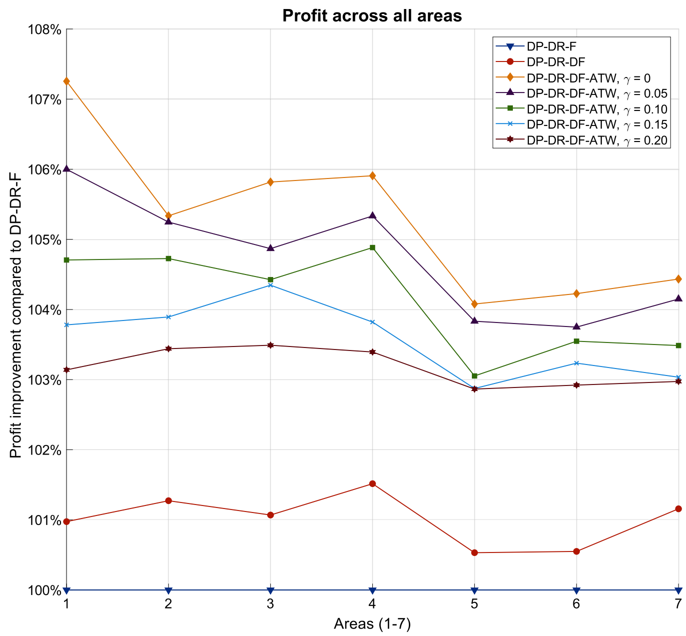

Figure 4.10: Impact of γ values on profit in all areas

A more comprehensive display of the impact of γ values is presented in Table 4.3, which focuses on representative areas 1, 3, and 6. This table provides a comparison of the performance metrics between DP-DR-DF and DP-DR-DF-ATW variants using different γ values. As explained before, higher γ values means lower preference of ATWs and therefore lower selection rates of these slots. As a result this leads to increased mileage per order and a decrease in the number of orders, which justifies ATWs' ability of maintaining routing flexibility and conveying more orders. As γ increases, the total profit decreases. However, even with the largest γ value, the DP-DR-DF-ATW variants outperform DP-DR-DF. This is particularly evident when considering the profit improvement, which indicates a larger deviation from the maximum profit for DP-DR-DF.

Table 4.3: Sensitivity analysis on γ for DP-DR-Df-ATW with comparison to DP-DR-DF

	γ **values**	Total mileage	Mileage/Order	Number of orders	Total order size	Accepted price	Total profit	profit improvement(%)
Area 1	0	510.74	6.62	76.27	284.63	1.80	2712.81	6.22
	0.05	517.50	6.62	75.7	285.79	1.67	2681.05	4.98
	0.10	507.21	6.70	75.1	282.067	1.51	2648.33	3.70
	0.15	516.42	6.88	74.26	276.80	1.37	2624.93	2.78
	0.20	513.90	6.86	74.03	276.43	1.39	2608.71	2.74
	DP-DR-DF	507.32	6.89	72.93	274.46	0.95	2553.96	

	γ **values**	Total mileage	Mileage/Order	Number of orders	Total order size	Accepted price	Total profit	Profit improvement(%)
Area 3	0	545.32	2.23	239.20	914.05	2.39	8847.95	4.70
	0.05	553.63	2.33	237.20	909.17	2.26	8768.50	3.76
	0.10	549.51	2.30	237.93	907.25	2.18	8731.50	3.32
	0.15	552.19	2.33	236.13	906.33	2.20	8724.89	3.24
	0.20	540.03	2.28	239.39	900.57	2.12	8653.31	2.40
	DP-DR-DF	537.05	2.28	236.60	899.25	1.32	8450.85	

	γ **values**	Total mileage	Mileage/Order	Number of orders	Total order size	Accepted price	Total profit	Profit improvement(%)
Area 6	0	391.22	0.72	543.30	2141.70	1.11	19996.41	3.66
	0.05	391.92	0.72	539.63	2135.04	1.06	19904.82	3.18
	0.10	390.73	0.72	540.67	2140.20	0.90	19866.18	2.98
	0.15	391.18	0.73	538.23	2134.21	0.89	19806.22	2.67
	0.20	392.70	0.73	538.73	2139.41	0.69	19746.09	2.36
	DP-DR-DF	392.58	0.73	535.37	2122.49	0.13	19291.04	

In the final analysis, the selection rate comparison between ATW and STW is considered. Figure 4.11 shows that a considerable proportion of customers, ranging from 8.5% to 13.2%, exhibited potential behavioural changes by opting for ATW instead of STW in the experiment conducted. This finding strongly suggests that a notable segment of customers offer their flexibility and potential cost savings offered by choosing from ATW, which typically have a lower delivery charge, as opposed to selecting specific predefined time slots with a higher delivery charge associated with STW.

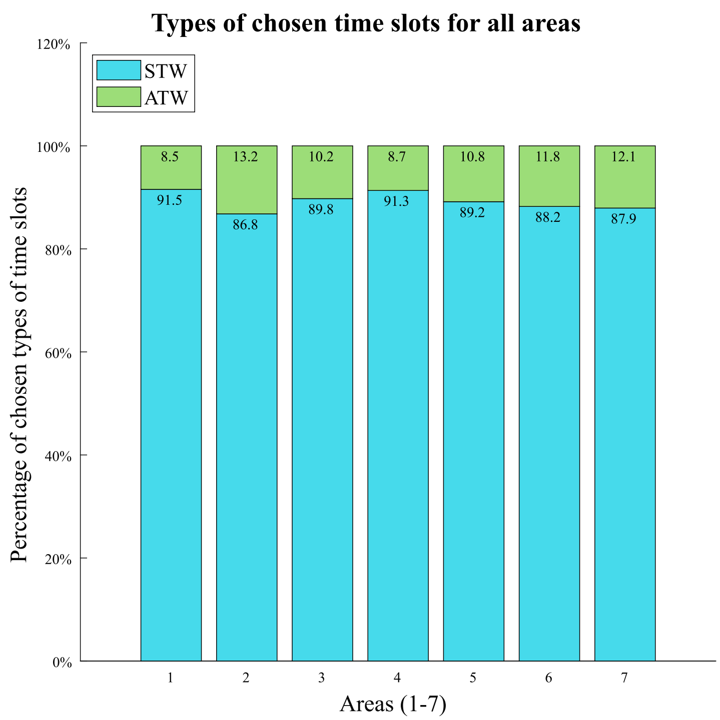

Figure 4.11: Percentage distribution: selection of STW versus ATW in all areas

4.7 Conclusion

In conclusion, this research makes two contributions to address the integrated demand management and dynamic vehicle routing problem. Firstly, by incorporating the distribution of forecast orders into the dynamic routing system with forecast orders without time window, the estimation of opportunity cost has been improved, resulting in enhanced profitability for the overall demand management system. This improvement demonstrates the potential for increased financial gains by leveraging forecast orders in dynamic routing decisions.

Secondly, the introduction of augmented time windows expands the range of time slot choices available to customers, leading to higher slot availability throughout the booking period. This, in turn, improves the overall performance of the dynamic routing system, resulting in better efficiency and effectiveness. this research outlines the process of creating augmented time windows and updating the choice model to align with the offered time slots, providing practical guidance for implementation.

The results obtained from this experiments showcase the promising impact of augmented time windows on profit-making, system simplicity, and integrability. However, there is still room for further exploration and enhancement. Future research can focus on managing standard time slots that are not adjacent and developing sophisticated incentive schemes to make augmented time windows more attractive and appealing to customers. These extensions have the potential to further optimise the performance and attractiveness of the proposed approach.

Overall, these findings highlight the potential benefits of integrating forecast orders and augmented time windows into the demand management and dynamic routing system. This research contributes to the advancement of the field and paves the way for future innovations in improving profitability, efficiency, and customer satisfaction in attended home deliveries.

5

Enhancing Opportunity Cost Approximation through Incorporation of Delivery Price Reduction

5.1 Introduction

This chapter focuses on improving demand fulfilment by introducing a crucial concept—the cost associated with assigning a current time slot to a customer, and how this affects the delivery prices for future customers. This approach centres on the idea of opportunity cost, which measures the potential revenue lost when one customer order is selected over another.

The main goal of the delivery model—maximising profit—relies on three factors: revenue from accepted orders, the prices customers pay for their chosen delivery slots, and the costs of fulfilling these orders. Understanding the trade-offs between immediate revenue and potential future gains or losses is critical, especially when considering the costs involved in accepting new orders. However, there is a gap in research on how lowering prices for future slots, in response to current customer choices, affects overall profitability. Addressing this could greatly refine the understanding of opportunity costs.

Research in this area has used various methods, including simulation and predictive models, as seen in studies by Yang and Strauss (2017) and Ulmer (2020). Other researchers, like Koch and Klein (2020) and Abdollahi et al. (2023), have explored anticipatory route planning. These studies help us understand components of opportunity cost such as marginal delivery cost and displacement cost. Yet, there is still much to learn about how changes in delivery prices for future slots influence overall costs and decision-making. This aspect is particularly relevant to dynamic pricing strategies and is further discussed by Koch and Klein (2020), who provide insights into how e-retailers adjust their strategies to maximise profits by optimising how orders and delivery slots are allocated.

This study begins to explore how to estimate price reductions for less popular remaining time slots using a dynamic pricing model based on customer choices. The research also aims to incorporate these estimates into broader demand management strategies, enhancing how delivery services are planned and executed.

The contributions of the work can be summarised as follows:

- **Enhanced Opportunity Cost Approximation:** This research introduces an innovative approach to refine opportunity cost estimation in AHD. By analysing the effects of accepting one order now on potential delivery price reductions for future orders, the estimation of opportunity cost becomes more accurate and comprehensive.

- **Integration with Existing Demand Management Strategies:** The approach developed in this research can be integrated with existing methods within the Automated Home Delivery (AHD) framework. By doing so, it enhances the calculation of opportunity costs, thereby improving the accuracy of determining delivery prices for available slots.

The subsequent sections of this chapter are structured as follows: Section 5.2 delineates the contribution, focusing on enhancing the estimation of opportunity cost by adjusting delivery prices for future time slots based on the reduction in delivery prices resulting from slot allocation at the present time. Section 5.4 elaborates on the experimental setting, encompassing area characteristics and benchmark approaches. Moving forward, Section 5.5 presents a detailed analysis of the proposed approach's performance, emphasising its potential for integration with other existing methodologies. The study culminates in Section 5.6, where a conclusive summary of the accomplishments and the insights gained throughout this investigation is provided.

5.2 Dynamic pricing formulation in attended home delivery

Similar to Chapter 4, this chapter extends the problem framework introduced in Chapter 3. The core structure of the problem remains unchanged, focusing on the management of a delivery system by an e-retailer within a discrete and finite booking horizon. However, specific adaptations unique to this chapter need to be highlighted:

- **Adjustments to Time Windows**: The number of time windows has been reduced to 17 from a previous count of 27, and these windows are now non-overlapping. This dual adjustment serves two critical purposes:
  - **Granular Analysis of Displacement Costs**: By treating each time slot individually with its own specific allocation of forecast orders, the analysis becomes more granular. This change ensures that each time slot can be analysed independently, which is crucial for reflecting a more realistic scenario where each slot is served at the earliest possible time.
  - **Accurate Representation of Demand**: Making the time windows non-overlapping is essential to prevent the distribution of forecast orders across overlapping slots. This avoids distortions in displacement cost calculations that can arise from inaccurately represented time-specific demands and operational constraints.

The central elements of the problem, such as customer arrivals, state representation, probability models, and the objective function, remain consistent with the prior chapters. For a comprehensive understanding of these foundational components, readers are directed to Chapter 3.

This study aims to enhance opportunity cost estimation by approximating the reduction in delivery prices for subsequent feasible time slots resulting from assigning an available time slot to the current customer. The dynamic routing will be used to control the order acceptance system. Moreover, the marginal delivery costs and potential revenue loss will be approximated by solving a CVRPTW dynamically over the booking horizon. Depending on the specific design of various methodologies, the calculation of opportunity cost may vary. Each approach may emphasise or incorporate different cost components in its estimation of the overall opportunity cost. Below are some potential costs typically considered in approximating opportunity cost:

1. **Marginal Insertion Cost**: This is the immediate cost associated with accepting an incoming customer request. It arises from the additional mileage required for delivery due to the inclusion of the new order. The estimation of this cost is a fundamental part of the opportunity cost calculation.

2. **Revenue Loss Due to Time Slot Allocation**: Another cost consideration pertains to the allocation of a time slot to the new arrival. When a time slot is allocated to a current customer, it becomes unavailable for future customers. In the existing literature, this cost is known as the "revenue loss due to time slot allocation". It is assessed by assigning a monetary value to each time slot, referred to as the "time budget", and calculating the amount of time the new arrival will occupy in the selected time slot. This cost is then reflected in the opportunity cost.

3. **Delivery Price Reduction of Remaining Time Slots**: The third factor relates to the prices offered by the company. To incentivise customers to book the remaining slots, which are, in most cases, less-popular ones, the company may reduce the delivery-slot prices towards the end of the booking horizon. This reduction aims to make the remaining time slots more attractive for selection, thus increasing overall profit. This cost, associated with reducing the delivery prices of the remaining time slots due to the allocation of the customer-selected time slot, is the focus of this study, which is referred to as "Delivery Price Reduction (DPR)".

Having delineated the potential costs involved in opportunity cost approximation, it becomes imperative to offer a comprehensive exploration of DPR. The subsequent section provides an in-depth analysis, unveiling the intricacies and nuanced impacts of DPR on the overall opportunity cost, and by extension, the decision-making processes in the context of Attended Home Delivery (AHD).

5.3 Delivery Price Reduction (DPR)

This section extends the methods discussed in Chapter 3 and Chapter 4. It is based on the published work in Abdollahi, M., Yang, X., Nasri, M. I., and Fairbank, M. (2023). Demand management in time-slotted last-mile delivery via dynamic routing with forecast orders. European Journal of Operational Research.

Most works in recent AHD context consider first two components in the opportunity cost approximation, i.e., the marginal insertion cost and the potential revenue loss due to time slot allocation. The aim is to improve the opportunity cost estimation by incorporating an approximation of the reduction in delivery prices for subsequent feasible time slots arisen from allocating an available time slot to the current customer, as denoted by DPR in this study.

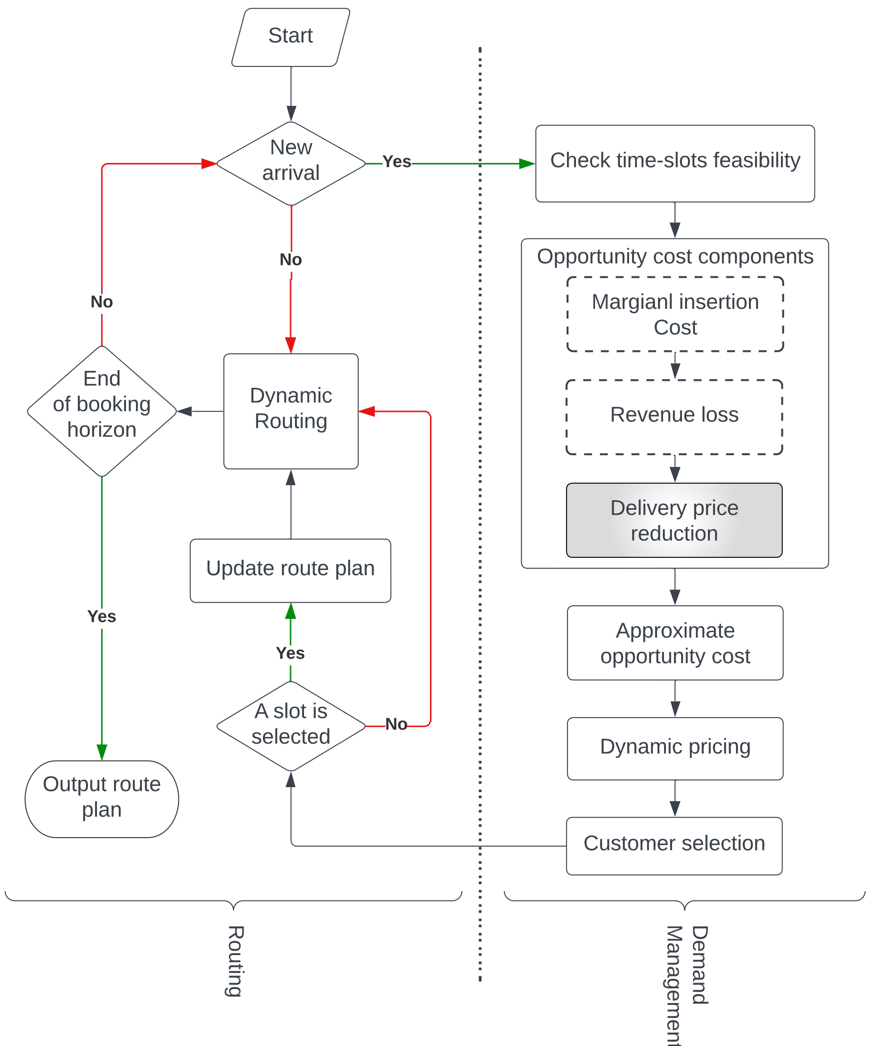

Figure 5.1: Implication of delivery price reduction (DPR) in the integrated dynamic routing and demand management problem of the AHD.

An overview of the incorporation of DPR as a new component in the opportunity cost within the context of integrated demand management and vehicle routing problem is depicted in the flowchart 5.1. An important point to note is that the DPR component can operate as an independent term in conjunction with other components within the opportunity cost. Consequently, its integration to enhance the estimation of opportunity cost is not restricted by any methodologies classified as non-learning-based solution approaches by Klein and Steinhardt (2023).

Section 5.3.1 begins by introducing the Multinomial Logit (MNL) model for customer choices and the closed-form solution for optimal pricing in in AHD applications, and then elucidate the process of estimating DPR in Section 5.3.2.

5.3.1 Multinomial Logit (MNL) Model and Conditions for Optimal Pricing

The dynamic pricing strategy outlined in this section leverages the foundational models and equations introduced in Chapter 3. Utilising the dynamic programming formulation (3.2.1), this chapter focuses on the application of these principles specifically to the time slot selection in the context of an online dynamic pricing problem.

As previously discussed, estimating the selection probabilities for each time slot or opting out entirely at each time step t is complex. For this, the Multinomial Logit (MNL) model, initially introduced by McFadden et al. (1973) and elaborated in Chapter 3, is employed. The model captures customer choice probabilities, which are crucial for formulating any effective dynamic pricing strategy.

The MNL model assumes that customers are utility maximisers and that their choices are influenced by a set of utilities, which include a base utility β_0, the utility specific to each time slot β_s, and the sensitivity to delivery charges β_d. These utilities, detailed in Chapter 3, are integrated into the selection probability equations (3.3.7) and (3.3.8), reflecting the complex interplay of pricing, slot availability, and customer preferences.

Given the reliable approximation of the expected opportunity cost $OC_t(\vec{x}_t, \mathbf{1}_{as})$ — as defined in Chapter 3 and approximating $V_{t+1}(\vec{x}_t) - V_{t+1}(\vec{x}_t + \mathbf{1}_{as})$ — we apply Equations (3.3.9) and (3.3.10) from Chapter 3 to calculate the optimal pricing d_s^*. This approach, suggested by Dong et al. (2009), integrates both theoretical modelling and empirical data to optimise pricing decisions effectively.

5.3.2 DPR formulation

To elucidate DPR, let us consider a system status (state) at time t, denoted as $\vec{x}_t$, upon receiving a new order i from a, with the feasible time slots $F_a(\vec{x}_t)$. The goal is to evaluate the potential price variance for a subsequent order, based on whether the time slot s is allocated to order i or not. This is denoted by $DPR(\vec{x}_t, \mathbf{1}_{as})$. Note that this approach primarily centres on evaluating the price reduction of a single future order at the time point $t+1$. This choice facilitates a focused analysis to ensure manageability and efficiency, as in Markov Decision Process (MDP) the impact of subsequent states diminishes with increasing temporal distances. To accomplish this, an artificial order, $i+1$, is created by averaging the longitude and latitude coordinates of the orders within all potential customers and similarly, its revenue is determined by the average revenues of these orders. Denote the location and the revenue of order $i+1$ by a' and r', respectively.

To calculate the $DPR(\vec{x}_t, \mathbf{1}_{as})$, two distinct scenarios are considered, relevant to with whether order i selects slot s at stage t or not. If order i does not select slot s at stage t, the set of feasible slots for the artificial order $i+1$ is denoted by $F_{a'}(\vec{x}_t)$. The opportunity cost of inserting the artificial order $i+1$ into slot w in the tentative route is denoted by $OC_{t+1}(\vec{x}_t, \mathbf{1}_{a'w})$, and the equation (3.3.10) is applied:

$$(h-1)\exp(h) = \sum_{w \in F_{a'}(\vec{x}_t)} \exp(\beta_0 + \beta_w + \beta_d(OC_{t+1}(\vec{x}_t, \mathbf{1}_{a'w}) - r')) \tag{5.3.1}$$

On the other hand, suppose order i selected slot s at stage t, the set of feasible slots for the artificial order $i+1$ in stage $t+1$ becomes $F_{a'}(\vec{x}_t + \mathbf{1}_{as})$. Note that, slot s might either remain available or become unavailable, depending on the extent to which order i occupies slot s. In this context, $F_{a'}(\vec{x}_t + \mathbf{1}_{as}) = F_{a'}(\vec{x}_t)$ if s remains available, and change to $F_{a'}(\vec{x}_t + \mathbf{1}_{as}) = F_{a'}(\vec{x}_t) \setminus \{s\}$ if it becomes unavailable. The opportunity cost associated with inserting the artificial order $i+1$ into the tentative route is denoted by $OC_{t+1}(\vec{x}_t + \mathbf{1}_{as}, \mathbf{1}_{a'w})$, and accordingly equation (3.3.10) is adapted to this scenario:

$$(h' - 1)\exp(h') = \sum_{w \in F_{a'}(\vec{x}_t + \mathbf{1}_{as})} \exp(\beta_0 + \beta_w + \beta_d(OC_{t+1}(\vec{x}_t + \mathbf{1}_{as}, \mathbf{1}_{a'w}) - r')) \quad (5.3.2)$$

The resulting difference in optimal delivery prices offered to order $i+1$ in stage $t+1$ for each slot $w \in F_{a'}(\vec{x}_t + \mathbf{1}_{as})$ is denoted by $DPR_w(\vec{x}_t, \mathbf{1}_{as})$, and defined as:

$$\begin{aligned} DPR_w(\vec{x}_t, \mathbf{1}_{as}) &= d^*_w - d'^*_w \\ &= OC_{t+1}(\vec{x}_t, \mathbf{1}_{a'w}) - OC_{t+1}(\vec{x}_t + \mathbf{1}_{as}, \mathbf{1}_{a'w}) - \frac{h - h'}{\beta_d} \\ &\qquad \forall s \in F_a(\vec{x}_t), \forall w \in F_{a'}(\vec{x}_t + \mathbf{1}_{as}). \end{aligned} \quad (5.3.3)$$

Note that r' terms have been cancelled from the expression within Equation (5.3.3) due to their equality. Note further that when a specific slot $\tilde{w} \in F_{a'}(\vec{x}_t)$ but $\tilde{w} \notin F_{a'}(\vec{x}_t + \mathbf{1}_{as})$, the DPR for $\tilde{w}$ is not defined by equation (5.3.3). Nevertheless, the influence of the different availability of $\tilde{w}$ on the price reduction is captured within the h and h' values across all slots $w \in F_{a'}(\vec{x}_t + \mathbf{1}_{as})$, as well as by the other factors of the opportunity cost, i.e., insertion cost and revenue loss stemming from the introduced unavailability.

Finally, the collective influences of having slot s occupied by order i in stage t, i.e., the $DPR(\vec{x}_t, \mathbf{1}_{as})$, is computed as the average of the $DPR_w(\vec{x}_t, \mathbf{1}_{as})$ across all feasible slots $w \in F_{a'}(\vec{x}_t + \mathbf{1}_{as})$:

$$DPR(\vec{x}_t, \mathbf{1}_{as}) = \frac{1}{|F_{a'}(\vec{x}_t + \mathbf{1}_{as})|} \sum_{w \in F_{a'}(\vec{x}_t + \mathbf{1}_{as})} DPR_w(\vec{x}_t, \mathbf{1}_{as}). \tag{5.3.4}$$

The same step has to be carried out for all $s \in F_a(\vec{x}_t)$, i.e., feasible slots when inserting order i in the tentative route in stage t. These terms, all together, form the vector $\overrightarrow{DPR} = (DPR(\vec{x}_t, \mathbf{1}_{as}))_{s \in F_a(\vec{x}_t)}$. Algorithm 5 provides a summary on how to compute the $\overrightarrow{DPR}$ at stage t for order i.

Upon execution of Algorithm 5, the computed reduction in delivery prices for all feasible slots is integrated into the opportunity cost calculation. An updated opportunity cost vector, $\overrightarrow{OC}^{updated} = (OC^{updated}(\vec{x}_t, \mathbf{1}_{as}))_{s \in F_a(\vec{x}_t)}$, is formulated according to

$$OC^{updated}(\vec{x}_t, \mathbf{1}_{as}) = OC(\vec{x}_t, \mathbf{1}_{as}) + DPR(\vec{x}_t, \mathbf{1}_{as}), \quad \forall s \in F_a(\vec{x}_t) \tag{5.3.5}$$

Here, OC represents the initial opportunity cost calculated by the employed method, which entails either the marginal insertion cost ($OC = IC$), or the combination of insertion cost and displacement cost ($OC = IC + RL$). The updated opportunity cost with DPR is instrumental in determining the optimal delivery prices ($\vec{d}^*$) as per Equation (3.3.9).

Algorithm 5: DPR calculation at stage t with simulation into stage $t+1$

Input : Feasible slots for order i at stage t: $F_a(\vec{x}_t)$,

Route plan $G_t(x_t)$,

MNL parameters: β_0, β_w, β_d

Output: Delivery price reduction at stage t: $\overrightarrow{DPR} = (DPR(\vec{x}_t, \mathbf{1}_{as}))_{s\in F_a(\vec{x}_t)}$

1 Create an artificial order $i+1$, denote its location by a' and revenue r';

2 Identify the feasible slots $F_{a'}(\vec{x}_t)$ for order $i+1$ considering $G_{t+1}(\vec{x}_t)$;

3 Determine $OC_{t+1}(\vec{x}_t, \mathbf{1}_{a'w}), \forall w \in F_{a'}(\vec{x}_t)$, considering $G_{t+1}(\vec{x}_t)$ and the employed method;

4 Calculate h with $F_{a'}(\vec{x}_t)$ using Eq. (5.3.1);

5 Initialise $DPR_s = 0, \forall s \in F_a(\vec{x}_t)$;

6 **for** *each slot* $s \in F_a(\vec{x}_t)$ **do**

7 Identify the feasible slots $F_{a'}(\vec{x}_t + \mathbf{1}_{as})$ for order $i+1$, considering $G_{t+1}(\vec{x}_t + \mathbf{1}_{as})$;

8 Determine $OC_{t+1}(\vec{x}_t + \mathbf{1}_{as}, \mathbf{1}_{a'w}), \forall w \in F_{a'}(\vec{x}_t + \mathbf{1}_{as})$ considering $G_{t+1}(\vec{x}_t + \mathbf{1}_{as})$, and the employed method;

9 Calculate h' with $F_{a'}(\vec{x}_t + \mathbf{1}_{as})$ using Eq. (5.3.2);

10 **for** *each slot* $w \in F_{a'}(\vec{x}_t + \mathbf{1}_{as})$ **do**

11 Compute $DPR_w(\vec{x}_t, \mathbf{1}_{as}) = d^*_w - d'^*_w$ using Eq. (5.3.3);

12 $DPR(\vec{x}_t, \mathbf{1}_{as}) \leftarrow$ average$(DPR_w(\vec{x}_t, \mathbf{1}_{as})_{w\in F_{a'}(\vec{x}_t+\mathbf{1}_{as})})$ according to Eq. (5.3.4);

5.4 Experimental Setting

To evaluate the impact of DPR on AHD solutions, DPR is integrated into several commonly used OC approximation approaches and run tests to explore the differences brought by it. These include some state-of-the-art approaches found in most recent AHD literature. The experiments are conducted in MATLAB, using a high-performance computing system equipped with an Intel Core i9-7940X processor running at 3.1GHz. To ensure the reliability of these findings, 30 separate experiment runs are performed for every parameter setting; results presented are based on the resulting averages.

5.4.1 OC approximation methods

In this section, the dynamic pricing methods introduced and employed in Chapters 3 and 4 to enhance opportunity cost (OC) calculations across various scenarios:

1. Dynamic Pricing with Insertion Cost without Forecast Orders (DP-IC)
2. Dynamic Pricing with Fixed Routing of Forecast Orders and Time Windows (DP-FR-F)
3. Dynamic Pricing with Dynamic Routing of Forecast Orders without Time Windows (DP-DR-F)
4. Dynamic Pricing with Dynamic Routing of Distributed Forecast Orders without Time Windows (DP-DR-DF)
5. Dynamic Pricing with Dynamic Routing of Distributed Forecast Orders without Time Windows and Augmented Time Windows (DP-DR-DF-ATW)

These methods are analysed to determine their effectiveness in real-world scenarios, focusing on their ability to adapt to changing conditions and optimise pricing strategies dynamically. The specifics of each method are detailed and annotated in Table 4.2 in chapter 4. It is important to note that the type of *OC* calculation employed in Algorithm 5 will be determined based on the method selected, as outlined in this table.

5.4.2 Characteristics of selected geographic areas for experiment

This study presents a comprehensive analysis of results obtained in seven different geographic areas (Area 1 to Area 7) for the integrated demand management and vehicle routing problem as explained in Section 3.2. The goal of this experiment was to evaluate the performance of various methods with and without the incorporation of Delivery Price Reduction (DPR). Key metrics are analysed to assess the impact of DPR on the efficiency and profitability of the logistics operations in each area. These areas were chosen to represent a spectrum of customer densities and distribution patterns. Actual customer locations and historical booking data are employed in these experiments with necessary anonymisations for data privacy. Table 5.1 furnishes a comprehensive summary of the selected areas, outlining their specific characteristics and the number of delivery vans deployed to fulfil the ultimately accepted delivery orders.

Table 5.1: Features of seven areas under study

Area	1	2	3	4	5	6	7
Geographical spread	Rural	Suburb	Semi-Rural	Semi-Rural	Suburb	City	City
Average number of customers per day	63	130	206	308	398	502	611
Average distance between customers	11.78	5.87	13.43	11.98	8.81	2.67	2.32
Number of delivery vehicles	2	4	6	9	11	14	16
Type of vehicle	Van	Van	Van	Van	Van	Van	Van

5.5 Computational results

This section delves deeper into the results of these experiments. Section 5.5.1 evaluates the performance of the DPR across various areas. Section 5.5.2 investigates DPR's impact on delivery prices, while Section 5.5.3 examines how DPR affects time slot availability and pricing dynamics.

5.5.1 Evaluating the impact of DPR across various methods and geographical areas

Table 5.2 shows detailed results from these experiments. Different methods across seven areas are tested and results on total profit, mileage per order, order size, number of orders accepted and mean price are reported. Each method is indicated by its name (e.g., DP-IC) as listed in Section 5.4, together with a tweaked version after adding DPR (e.g., DP-IC-DPR). This table provides a detailed comparison of how each method performs. It highlights the potential improvements achievable by adding the step of calculating and incorporating DPR into OC estimation, thereby enhancing our understanding of its impact on the overall methodology.

Methods	Total mileage	Mileage/Order	Order size	#Order	Mean price	Total profit
Area 1						
DP-IC	585.52	8.52	248.11	68.23	-3.14	2034.34
DP-IC-DPR	574.71	8.28	251.73	69.27	-2.93	2078.09
DP-FR	571.82	7.99	261.14	71.07	-2.89	2161.34
DP-FR-DPR	564.85	7.79	266.42	71.97	-2.74	2217.05
DP-DR-F	504.12	7.08	265.60	70.50	0.61	2447.53
DP-DR-F-DPR	506.75	6.98	269.88	71.97	0.72	2494.65
DP-DR-DF	507.52	6.93	272.32	72.47	0.76	2521.52
DP-DR-DF-DPR	506.41	6.90	273.20	72.67	1.05	2550.45
DP-DR-DF-ATW	515.71	6.85	278.35	74.60	1.62	2641.51
DP-DR-DF-ATW-DPR	519.36	6.79	284.49	75.77	1.89	2717.52
Area 2						
DP-IC	443.55	3.17	522.84	139.90	-2.74	4352.13
DP-IC-DPR	440.85	3.09	532.10	142.23	-2.66	4441.04
DP-FR	446.74	3.18	524.38	139.83	-2.54	4393.27
DP-FR-DPR	436.93	3.04	540.47	143.43	-2.47	4541.15

Methods	Total mileage	Mileage/Order	Order size	#Order	Mean price	Total profit
DP-DR-F	384.89	2.57	564.12	148.67	1.23	5289.97
DP-DR-F-DPR	392.12	2.59	570.45	150.73	1.36	5369.52
DP-DR-DF	390.49	2.61	565.97	149.03	1.25	5309.40
DP-DR-DF-DPR	383.67	2.50	581.66	152.80	1.35	5472.42
DP-DR-DF-ATW	389.27	2.54	574.05	152.30	2.07	5512.43
DP-DR-DF-ATW-DPR	388.54	2.50	583.63	154.30	2.15	5616.60
Area 3						
DP-IC	620.59	2.71	867.14	228.10	-2.80	7213.05
DP-IC-DPR	619.39	2.67	878.79	231.57	-2.71	7331.38
DP-FR	621.42	2.75	862.11	225.70	-2.44	7257.77
DP-FR-DPR	620.05	2.67	885.68	231.63	-2.42	7460.52
DP-DR-F	545.83	2.34	894.59	232.70	1.21	8381.70
DP-DR-F-DPR	538.81	2.28	903.69	235.47	1.29	8486.19
DP-DR-DF	541.27	2.31	897.45	233.20	1.14	8390.63
DP-DR-DF-DPR	534.56	2.28	895.22	233.73	1.40	8430.50
DP-DR-DF-ATW	550.10	2.31	908.29	237.70	2.11	8726.34
DP-DR-DF-ATW-DPR	547.64	2.28	912.74	239.50	2.37	8832.35
Area 4						
DP-IC	1184.54	3.72	1196.65	318.73	-3.07	9856.20
DP-IC-DPR	1185.25	3.70	1198.82	319.97	-3.01	9891.16
DP-FR	1178.86	3.69	1203.20	318.87	-2.82	9994.52
DP-FR-DPR	1182.08	3.69	1207.82	319.57	-2.76	10052.80
DP-DR-F	1041.46	3.06	1288.48	339.37	0.96	11990.57
DP-DR-F-DPR	1024.28	2.99	1296.07	342.07	1.09	12104.20
DP-DR-DF	1010.66	2.95	1291.70	342.03	1.06	12057.28
DP-DR-DF-DPR	1017.38	2.96	1299.13	343.60	1.08	12132.72
DP-DR-DF-ATW	1036.77	2.98	1310.21	346.77	1.75	12470.28
DP-DR-DF-ATW-DPR	1033.98	2.96	1318.81	349.07	1.77	12557.99

Methods	Total mileage	Mileage/Order	Order size	#Order	Mean price	Total profit
Area 5						
DP-IC	1544.51	3.66	1579.50	421.23	-2.56	13222.20
DP-IC-DPR	1528.50	3.56	1603.76	429.20	-2.54	13428.10
DP-FR-F	1532.32	3.62	1578.68	423.00	-2.23	13348.19
DP-FR-F-DPR	1526.05	3.57	1595.32	426.47	-2.22	13497.81
DP-DR-F	1328.67	2.95	1677.45	450.07	2.42	16275.05
DP-DR-F-DPR	1327.24	2.95	1682.87	449.70	2.45	16335.30
DP-DR-DF	1325.52	2.93	1690.64	451.57	2.26	16324.28
DP-DR-DF-DPR	1324.28	2.93	1692.01	452.03	2.27	16341.63
DP-DR-DF-ATW	1334.95	2.93	1696.09	455.27	2.93	16687.71
DP-DR-DF-ATW-DPR	1320.54	2.89	1700.60	455.57	3.06	16789.40
Area 6						
DP-IC	447.23	0.91	1916.15	488.80	-2.97	15900.52
DP-IC-DPR	443.11	0.90	1922.14	490.87	-2.84	16012.96
DP-FR-F	442.68	0.91	1920.54	488.57	-2.56	16143.92
DP-FR-F-DPR	441.76	0.90	1928.44	490.93	-2.54	16218.92
DP-DR-F	391.83	0.73	2114.95	533.20	0.14	19228.54
DP-DR-F-DPR	388.69	0.73	2118.96	534.53	0.21	19301.92
DP-DR-DF	394.01	0.74	2117.49	534.80	0.19	19274.42
DP-DR-DF-DPR	391.43	0.73	2117.04	533.93	0.23	19296.35
DP-DR-DF-ATW	397.22	0.74	2122.64	538.00	1.06	19790.04
DP-DR-DF-ATW-DPR	390.97	0.72	2131.05	543.00	1.17	19933.24
Area 7						
DP-IC	475.53	0.86	2167.89	551.00	-3.45	17733.63
DP-IC-DPR	476.54	0.86	2175.70	553.13	-3.40	17819.05
DP-FR-F	476.38	0.87	2160.45	548.83	-3.04	17897.48
DP-FR-F-DPR	475.85	0.85	2186.28	556.70	-3.02	18119.03
DP-DR-F	418.22	0.66	2467.85	630.27	0.82	22866.20

Methods	Total mileage	Mileage/Order	Order size	#Order	Mean price	Total profit
DP-DR-F-DPR	417.08	0.66	2479.96	629.57	0.86	22997.64
DP-DR-DF	411.34	0.65	2471.75	627.93	0.56	22732.69
DP-DR-DF-DPR	415.25	0.66	2469.69	629.57	0.81	22870.58
DP-DR-DF-ATW	418.59	0.66	2472.38	629.71	1.30	23206.18
DP-DR-DF-ATW-DPR	419.05	0.66	2465.34	632.29	1.53	23290.49

Table 5.2: Obtained results for different methods in each area.

Table 5.2 demonstrates that incorporating DPR into OC calculation leads to modest improvements in key performance metrics across all tested areas. This is most notably reflected in the increased total profit and increased number of orders. The operational benefits of DPR are also evident in its indirect improvement of route efficiency in different approaches. By considering the marginal insertion cost and for some OC approaches the potential revenue loss associated with the artificial order in the subsequent time step, DPR calculations contribute to more efficient routing, as seen in the reduced mileage per order. This is a clear indication of enhanced resource utilisation and route optimisation.

To clarify the results shown in Table 5.2, where larger orders may correlate with shorter delivery distances, the improved routing optimisations enabled by DPR are key. By recalculating opportunity costs with updated delivery prices, DPR allows for strategic routing decisions that accommodate increased order volumes without proportionally extending delivery routes. This method ensures promoted slots are more likely to be booked, resulting in more efficient routing. This is particularly effective in densely populated or compact areas, where routing decisions have a greater impact. For instance, in Area 3, the DP-FR method reduced average mileage per order from 2.75 to 2.67 after implementing DPR, while the number of orders increased from 225.70 to 231.63. This demonstrates how DPR optimises routing to manage more orders efficiently within compact regions, enhancing operational efficiency and customer satisfaction.

Graphical comparison of profit and number of accepted orders are also presented in Figure 5.2 and Figure 5.3, with paired bars depicting the results both with and without the inclusion of DPR across different operational terrains. Percentage labels are also provided to highlight the tangible profit/number-of-orders gains attributable to DPR.

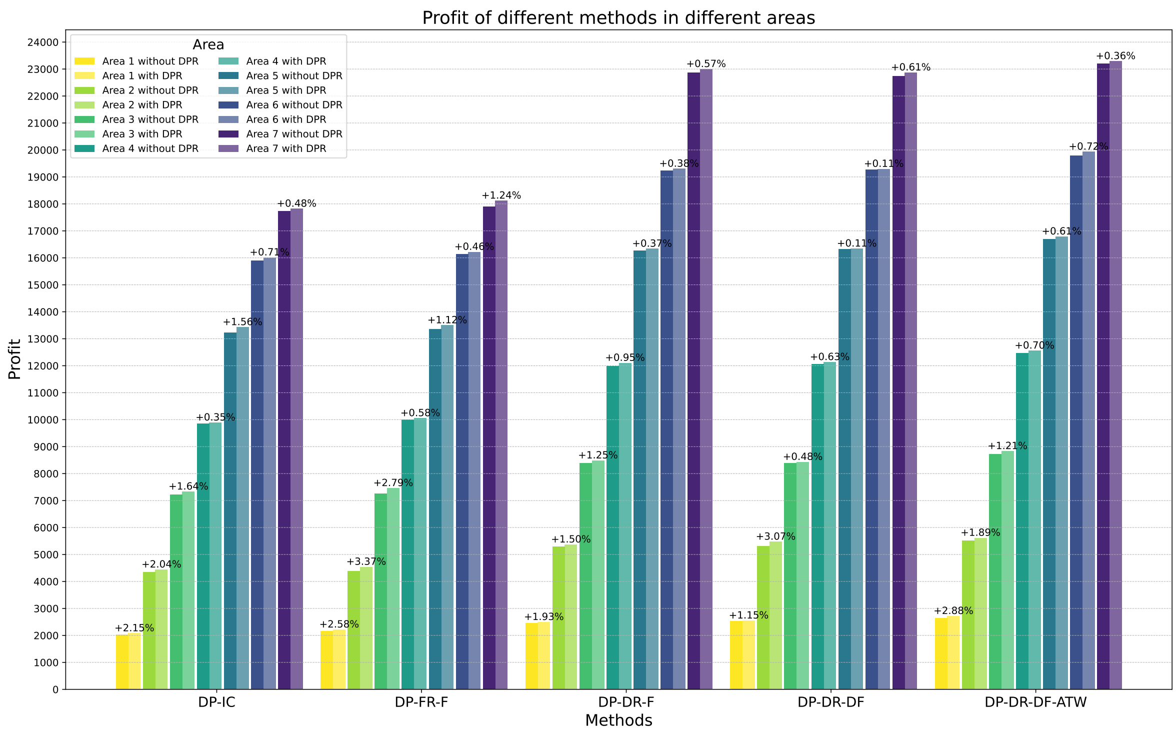

Figure 5.2: Average Profit Improvement Comparison Across All 7 Areas with and without DPR

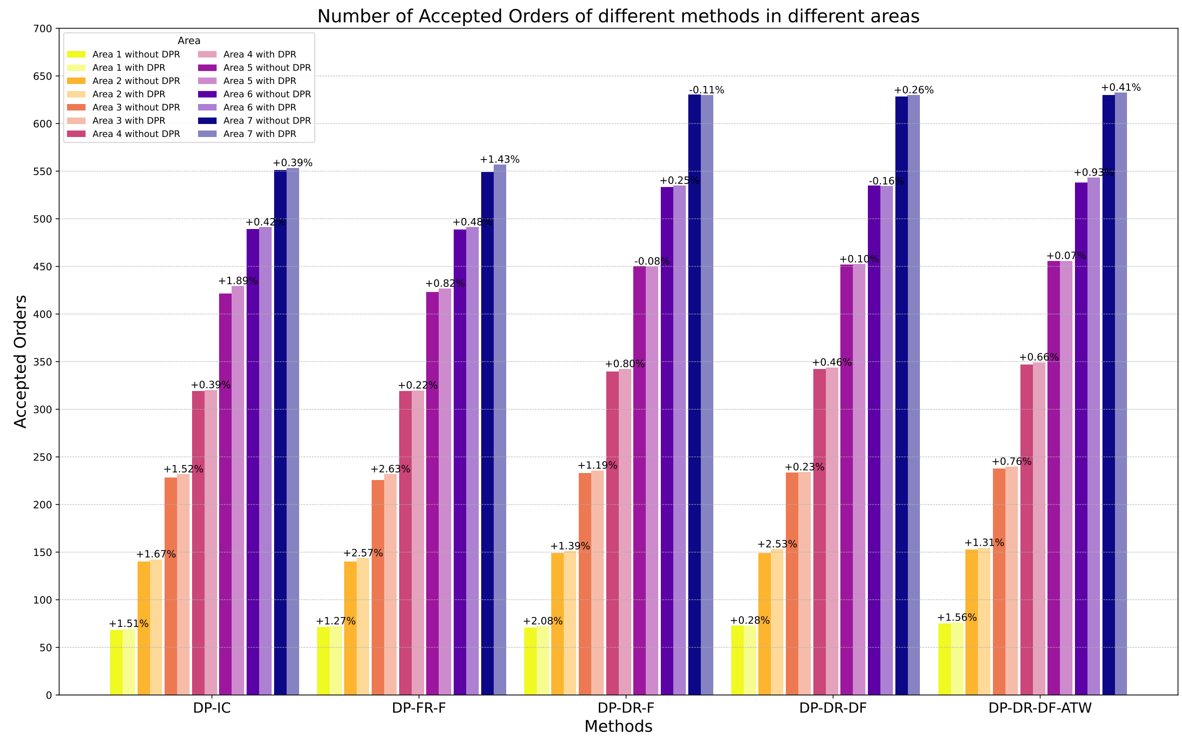

Figure 5.3: Average Request Improvement Comparison Across All 7 Areas with and without DPR

In examining the effectiveness of DPR methods across different geographic settings, a general pattern emerges: regardless of the area's density, from rural or suburban to distinctly urban environments, there is a consistent trend of increased number of acceptances, reduced mileage per order and increased profits. This trend is particularly evident in less dense areas like Area 1, Area 2 and Area 4.

Figure 5.4 displays the percentage improvement introduced by adding DPR to every approach. Each line in the plot corresponds to a particular geographical area.

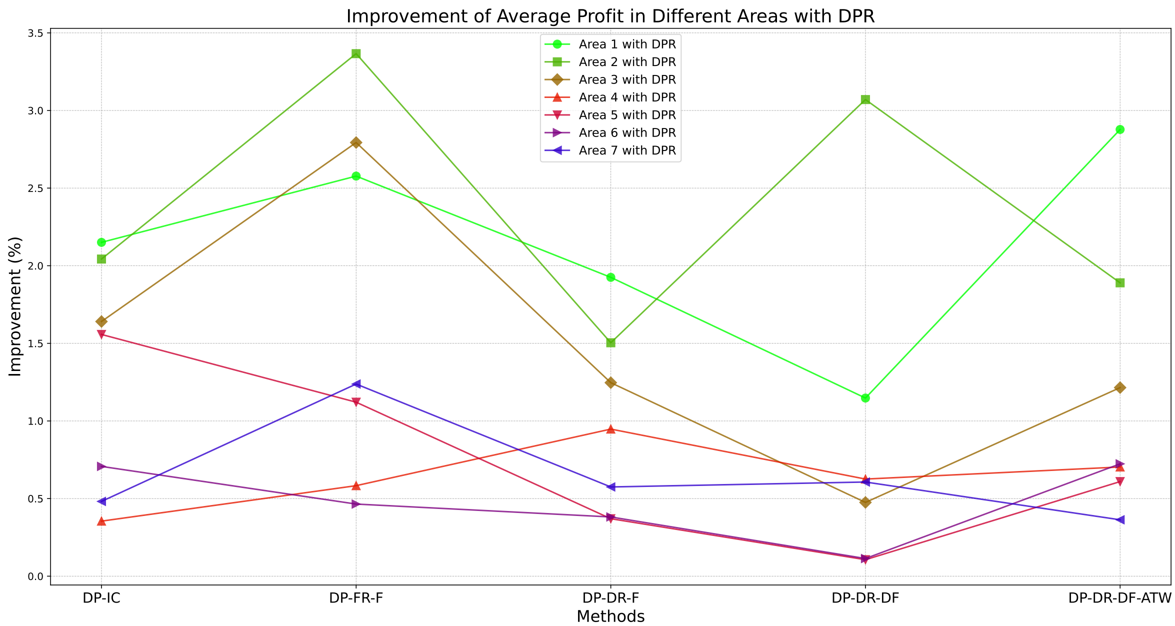

Figure 5.4: Improvement of Average Profit of the Studied Methods in Different Areas with DPR

This visualisation underscores the prevailing trend that implementing DPR generally boosts profit metrics, although the extent of enhancement varies depending on both the specific method and geographical context. In general, DPR was able to improve DP-IC and DP-FR-F methods more obviously compared to other methods. This is understandable as these methods employ a myopic and simple OC calculation mechanisms so that by including a foresight effect of delivery price variance in their OC calculation, the methods' performance witness a considerable increase. While the extent of average profit improvement is contingent upon the methods employed, this analysis generally indicates a minimum profit gain of 0.2%, with the potential to achieve increases up to a significant 3.4%. Notably, even with advanced methods for OC estimation such as DP-DR-DF-ATW, the inclusion of DPR still consistently demonstrates improvement. This justifies the necessity of DPR and its distinct focus compared to traditional opportunity cost terms such as insertion cost (marginal delivery cost) and revenue loss (displacement cost).

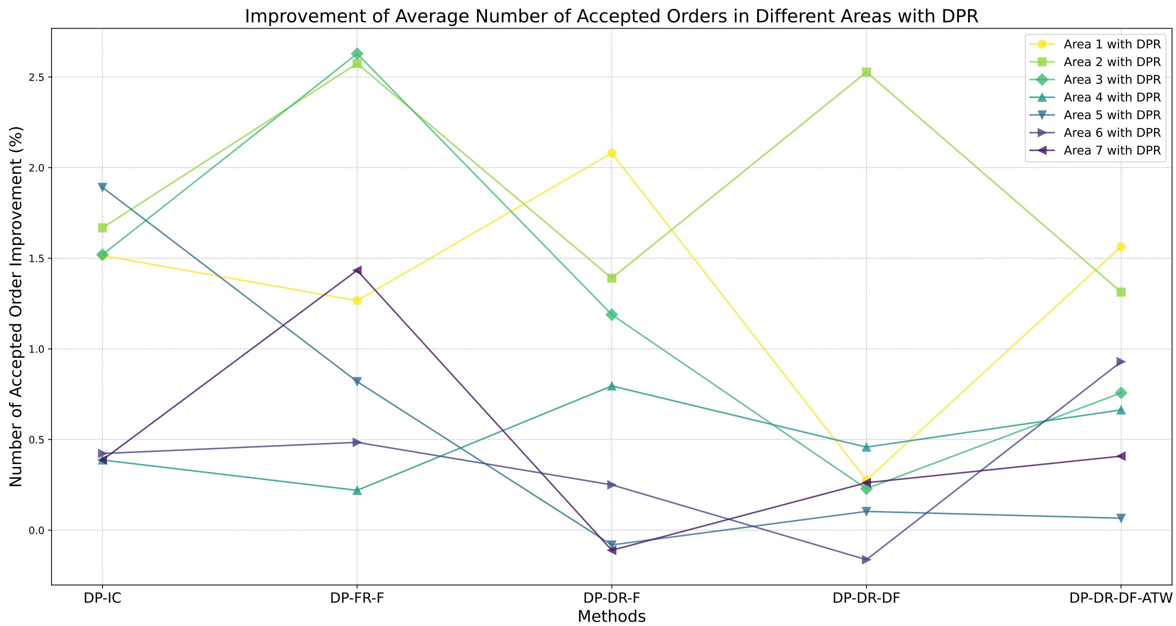

Figure 5.5: Improvement of Average Number of Accepted Orders in Different Areas with DPR

Similar to its impact on profitability, Figure 5.5 depicts a discernible percentage increase in order volumes with the adoption of DPR in the majority of areas and across various methods. Specifically, Areas 1, 2, and 3 experienced a more pronounced boost in order acceptance rates due to the incorporation of DPR. Meanwhile, other areas demonstrated a smaller yet steadier rise in the number of accepted orders. It is noteworthy, however, that in a few instances, methods such as DP-DR-F and DP-DR-DF with DPR resulted in a marginal decrease in order volume, ranging from 0.1% to 0.2%, compared to similar methods without DPR. Despite this, their overall profitability was higher. This uptick in order volumes highlights DPR's ability to enhance dynamic programming solutions for AHD, reinforcing the case for its broader implementation.

5.5.2 Influences of DPR on prices

There is an increase in the average price with the inclusion of DPR in OC calculation. The increase in average price is linked to the integration of DPR into the dynamic pricing model, which considers both slot availability and popularity. Figure 5.6 illustrates a clear relationship between slot pricing and popularity using bars for different slots (Areas 2, 3, and 6) and a line chart of scaled β_s values for DP-IC method.

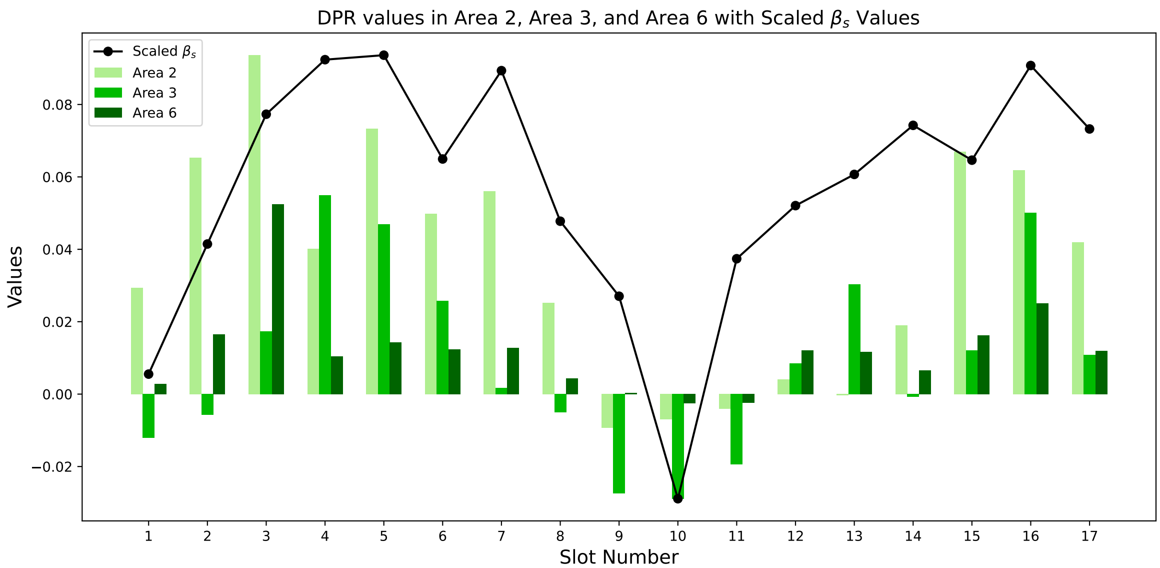

Figure 5.6: Comparative Analysis of DPR Values Across Area 2, Area 3, and Area 6 with Adjustments for Scaled β_s Parameters using DP-IC Method

As shown in Figure 5.6, dynamic pricing model, enhanced by DPR, effectively distinguishes between high-demand (popular) and low-demand (unpopular) slots. For popular slots, characterised by higher demand and limited availability, DPR leads to a greater price increase. In contrast, for slots that are typically less popular and more likely to remain un-booked, the price increment is smaller, making these slots more attractive to customers. Consider a popular slot s: if a new order is inserted, it often leads to s becoming unavailable for subsequent orders. Consequently, in such cases, Equation (5.3.2) contains one less term than Equation (5.3.1), resulting in $h > h'$. This scenario typically yields a positive DPR for popular slots. In contrast, for an unpopular slot s, it is more probable that the slot remains available post-insertion. Therefore, the difference between Equations (5.3.1) and (5.3.2) primarily lies in the OC values. With a more compact route in (5.3.2), the OC tends to be lower, making $h < h'$ and thus contributing to a negative DPR for these slots. Such a pricing approach encourages customers to opt for less popular slots which are priced more affordably, thereby achieving a more balanced and effective distribution of bookings across all available slots. Such a pricing strategy revision not only addresses immediate operational expenses but also considers the wider impact on future order scheduling and the availability of time slots, aligning operational efficiency with long-term strategic planning.

The distribution of average slot prices across different areas and methods is illuminated in Figure 5.7. The 'swarm' of points distinctly highlights the data's dispersion and concentration in various areas, providing a detailed understanding of the intricate price dynamics. Each method's and area's uniqueness is further accentuated by the dual-tone colour palette. The light blue tones represent methods without DPR integration, while light coral signifies those with DPR, facilitating an immediate visual differentiation and insight into the distinct impacts of DPR on pricing.

The implementation of DPR generally leads to higher slot prices. IC-based methods, specifically DP-IC and DP-FR-F, demonstrate negative average slot prices, contrasting with others that present positive values. This discrepancy suggests that marginal delivery cost alone is not sufficient to capture the whole opportunity cost. Nevertheless, the forward-looking feature brought by the DPR drives the prices slightly higher, leading to a more balanced pricing strategy that positively affects profit in the long run.

The methods that incorporate both IC and RL in opportunity cost estimation, such as DP-DR-F, DP-DR-DF, and DP-DR-DF-ATW, perform better to respond to area-

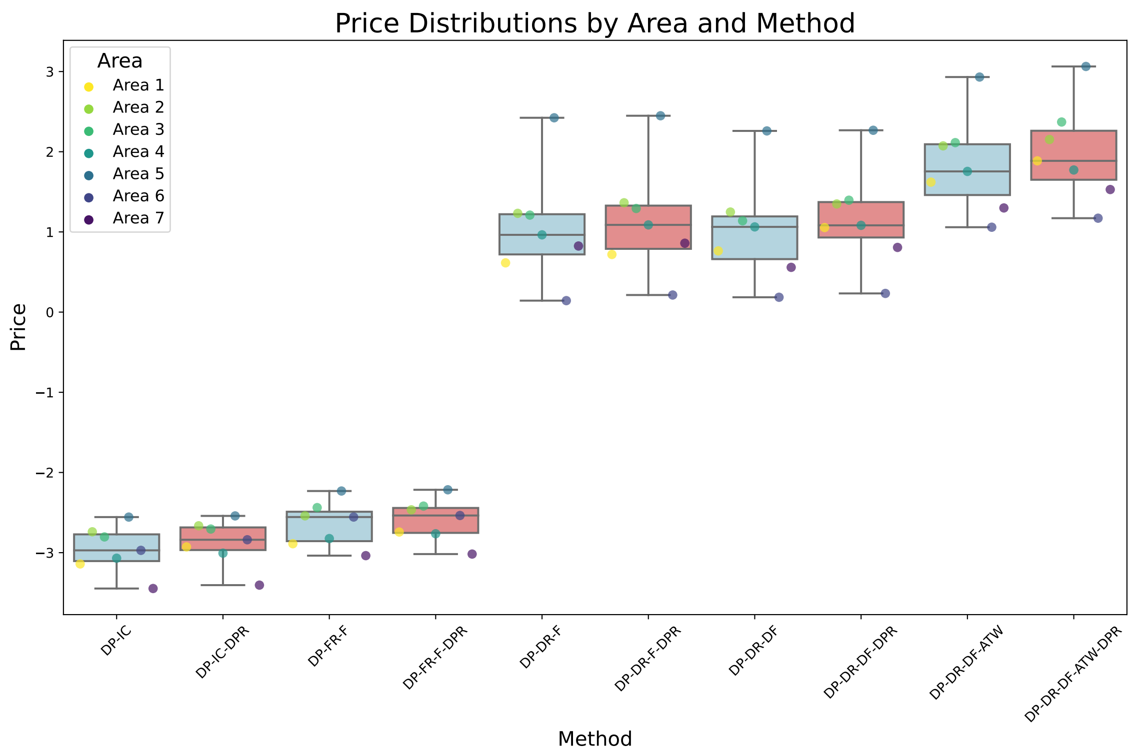

Figure 5.7: Price Distributions by Area and Method

specific details, resulting in notable variations in slot prices across different areas. These methods adeptly account for local demand patterns and travel requirements unique to each geographical area and customer distribution. For instance, Area 5 consistently has the highest prices among all areas tested. This is attributed to its limited delivery capacity relative to a high volume of requests and longer average distances between customers. In scenarios where an anticipatory pricing system predicts a potential capacity shortfall before the end of the planning horizon, it automatically raises the prices shown to customers to manage demand. Conversely, Areas 6 and 7, despite experiencing higher traffic, benefit from having customers located closer to each other. This proximity reduces travel requirements, enabling delivery vans to serve more customers efficiently, which in turn leads to generally lower prices in these areas compared to Area 5.

5.5.3 Exploring the influence of DPR on time slot availability and pricing dynamics

The following analysis delves deeper into the influence of DPR on the prices offered to customers. Dynamic pricing strategies take into account several factors, including the number of available slots offered for each customer arrival, order revenue, and the current estimated OC. These components collectively contribute to price adjustments. Visualising and examining how DPR affects the number of available slots and the average slot price offered provides a better understanding of the sources of profit improvement. Figure 5.8 illustrates the time slot availability and slot price over time for Area 4, as an example.

As shown in Figure 5.8, the integration of DPR leads to consistent increases in time slot availability for new arrivals. This translates into more slot options displayed for customers to choose from, which increases the chance for them to find a desired one to receive their orders and so as to boost the overall probability. It is worth emphasising that a key objective of a well-structured OC approximation is to provide a realistic estimate of future order values, thereby enhancing present-time pricing strategies. Including anticipated DPR makes the overall dynamic pricing approach more forward-looking. It leads to better recognition of which slot is the best one to offer and to promote at the current stage, as opposed to myopic adjustments that might deplete attractive time slots early in the booking horizon, leaving fewer options for high-volume, profitable orders in the future. This is considered the key reason for why higher availability and more efficient routing could be observed after adding DPR.

Another observation is the higher delivery prices offered throughout the entire booking horizon. Notably, this boost in average slot prices did not lead to a drop in the number of accepted orders. Instead, the number of acceptances increases as well as the

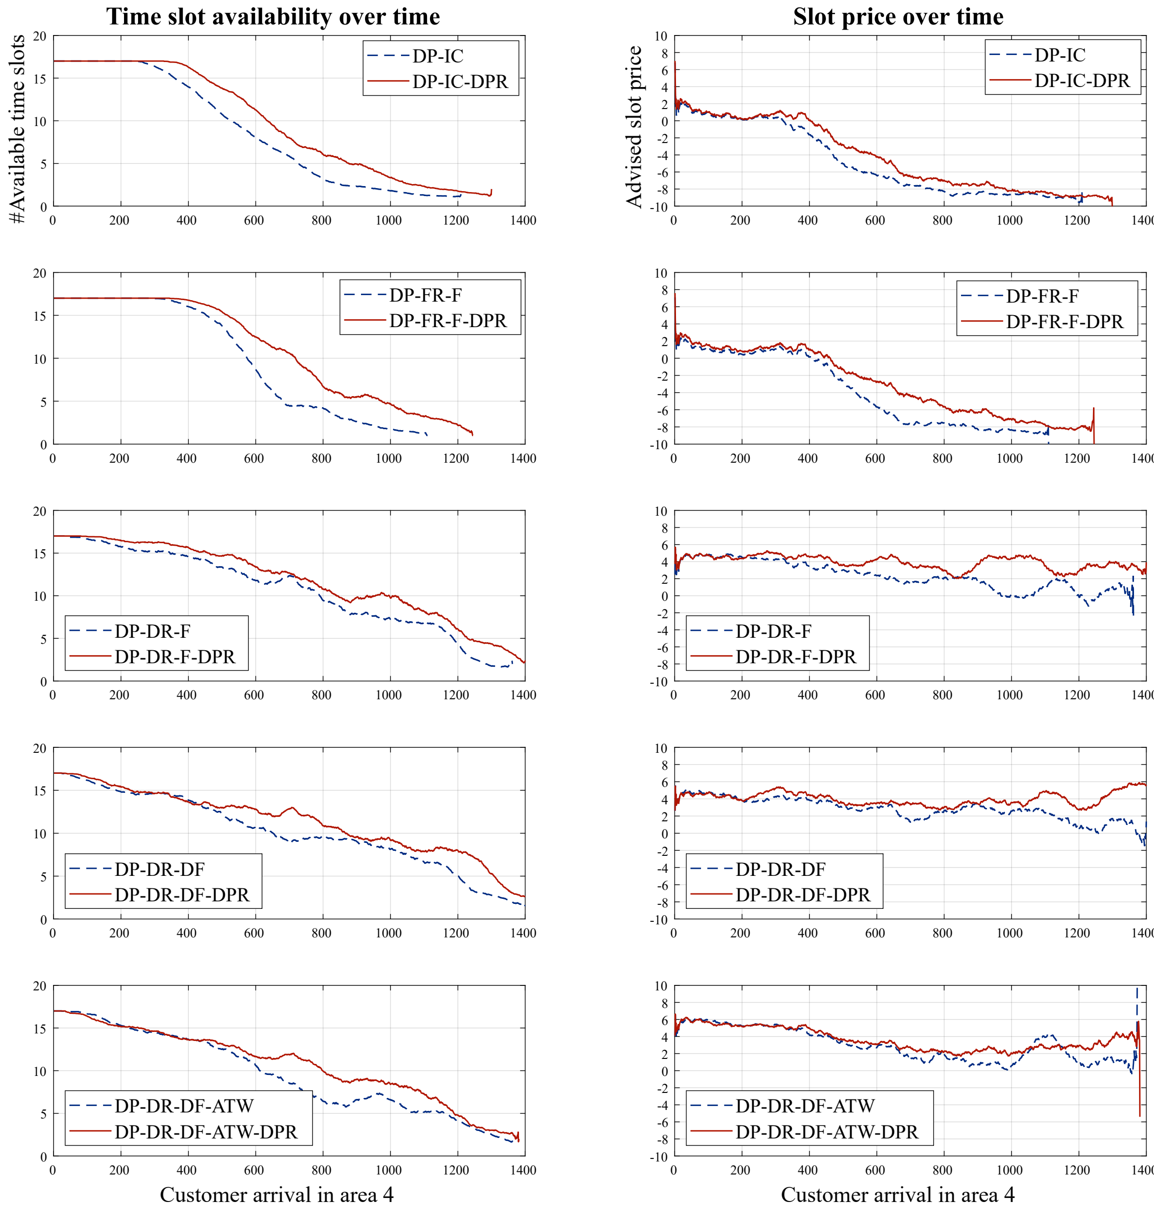

Figure 5.8: Effect of DPR on the number of available slots and slots' price over time for different methods in area 4.

delivery prices as a result of increased availability of slots. This justifies DPR's ability of helping identify suitable slots to accommodate the order and generating differentiated costs accordingly. Additionally, it can be observed that offered slot prices over time exhibit greater resistance to price drops when DPR is integrated. This suggests that the inclusion of DPR facilitates fairer slot pricing, aligning with the objectives of ideal AHD demand management.

In summary, the analysis illuminates the nuanced dynamics of slot pricing across various areas and methods, emphasising both the significant impact of DPR and the imperative for context-specific strategies. This tailored approach is essential to optimise efficiency and maximise profitability effectively.

5.6 Conclusion and Future Work

The analysis of the Delivery Price Reduction (DPR) and their impacts on attended home delivery (AHD) reveals considerable insights and advancements in optimising e-commerce delivery systems. These results across various geographic contexts, from rural to urban areas, consistently indicate an enhancement in total profit and operational efficiency due to the incorporation of DPR. Moreover, the integration of DPR into the opportunity cost estimation has proven to be a decent enhancement, supporting better informed and more optimal pricing decisions.

Each geographic area presented distinct patterns in how DPR affected profitability, influenced by factors such as customer density and geographical layout. This highlights the flexibility of DPR in adapting to various operational environments.

Among the various methods employed, those incorporating dynamic and distributed forecast orders, especially when amalgamated with augmented time windows, showcased optimal performance with profit increase from 0.2% up to 3.4%. This convergence of strategies not only maximised total profit but also improved the mean price, establishing a holistic optimisation of the AHD system.

Graphical illustrations further reinforced the tangible gains attributable to DPR, offering a visual testament to its efficacy. The visible increments in profit and order acceptance across all areas and methods validate DPR's applicability and effectiveness. The effectiveness of DPR remained consistent, even in complex urban areas with challenging logistics, proving its capability to optimise prices to maximise the revenue.

In conclusion, DPR is effective in the complex and changing field of attended home delivery. It is an important factor for better profit, operational efficiency, and improved customer satisfaction. The diverse advantages it reveals in different geographic and operational environments support its widespread integration and adaptation, marking it as a foundational element for the evolution of e-commerce delivery frameworks. Expected upcoming research could explore the scalable and adaptable features of DPR within a wider range of contexts. These in-depth analyses are anticipated to reveal crucial insights, contributing to the enhancement and expansion of DPR's utility in the global optimisation of e-commerce delivery landscapes.

5.6.1 Future Work

Building upon these findings, several promising avenues for further research emerge that not only refine the theoretical aspects of the DPR system but also pave the way for practical applications. Firstly, the integration of machine learning such as reinforcement learning offers a good opportunity to introduce greater sophistication into the DPR system. By incorporating more stages in the decision-making process, machine learning can lead to more precise predictions and enable a more dynamic response to market changes, thereby enhancing the effectiveness of DPR.

Secondly, acknowledging the limitations of using a single artificial order set, future research should diversify this approach. Introducing multiple artificial orders from various segments of the operational area, or developing a more sophisticated random technique for generating these orders, will provide a richer dataset. This enhancement is crucial for uncovering intricate patterns and gaining deeper insights into the adaptability and effectiveness of DPR across different contexts.

Thirdly, it is essential to extend the scope of research to include real-time implementations of DPR in various dynamic pricing structures. Exploring alternative formulations and testing these in real-world scenarios will not only provide a detailed understanding of how pricing decisions can be optimised but also validate the operational changes proposed in the thesis. Such studies could involve pilot implementations in selected e-commerce delivery contexts to observe the practical impacts of different DPR strategies and to refine them based on empirical evidence.

Lastly, Another area for future research involves examining the contribution of the DPR to opportunity costs. Quantifying this impact, potentially in terms of a percentage, will enable a more comprehensive understanding of DPR's importance in strategic decision-making within dynamic pricing systems. This line of inquiry promises to reveal the financial implications of different pricing strategies and their effects on revenue management

By addressing these facets, future studies will not only reinforce the theoretical understanding of DPR's benefits but also demonstrate its practical applications, ensuring that the research contributes directly to operational improvements and innovations in the e-commerce delivery domain.

6

Conclusion

6.1 Academic Contributions

This thesis contributes to the field of AHD in the e-commerce sector through the development of novel strategies aimed at enhancing delivery systems. The key academic contributions include:

- **Dynamic Pricing and Route Optimisation:** The introduction of a new dynamic pricing method using forecast orders without time windows has led to considerable improvements in order efficiency, cost reduction, and profitability. This method has been validated through real-world data and experiments across diverse geographic areas, demonstrating robustness and adaptability.

- **Opportunity Cost Approximation:** A simple-to-implement dynamic opportunity-cost approximation for marginal delivery cost and potential revenue loss was proposed, based on a dynamically managed routing system that includes both actually accepted orders and forecast orders without time windows. This approximation was integrated into an order-replacement and routing re-optimisation framework to capture the influence of new order commitments.

- **Forecast Orders Integration:** Incorporating forecast orders without time windows into the vehicle-routing system allowed for more appropriate time slot suggestions, guiding the choice of incoming orders accordingly. This approach's superiority was demonstrated over four benchmark approaches using real data-sets from various geographical and demographic settings.

- **Augmented Time Windows:** The introduction of augmented time windows expanded customer time slot choices and improved the dynamic routing system's efficiency and effectiveness. This resulted in higher slot availability and improved overall system performance.

- **Enhanced Opportunity Cost Approximation in AHD:** This research introduced an novel approach to refine opportunity cost estimation by analysing the effects of accepting one order now on potential delivery price reductions for future orders. This integration with existing demand management strategies enhanced the accuracy of delivery price determinations.

While the methods developed in this thesis offer numerous advantages, they also present certain limitations and trade-offs that merit discussion:

- **Computational Demands:** The dynamic pricing and route optimisation methods introduced require substantial computational resources, particularly as they scale to larger datasets and more complex scenarios and at the same time all forecast orders are expected to be evaluated by selecting a large radius of neighbourhood. The computational time can be excessive, which may not be practical for all operational contexts. To alleviated this issue selecting a smaller radius would be beneficial to handle excessive run-time.

- **Assumptions and Real-world Applicability:** The presumption that dynamic pricing can adequately balance customer demands across all time slots is an oversimplification. In reality, customer availability and preferences play a significant role and can lead to discrepancies between predicted and actual system performance.
- **Incremental Improvements vs. Computational Cost:** Some of the enhancements, such as the opportunity cost approximation and forecast orders integration, offer relatively small improvements at the cost of notable increased computational demands. This trade-off must be carefully considered when implementing these methods in practice, as the marginal gains may not always justify the additional resource expenditure.
- **Complexity in Implementation:** The integration of forecast orders and augmented time windows, while beneficial, adds complexity to the dynamic routing systems. This complexity could hinder adoption by firms lacking the technological infrastructure or expertise to implement such systems effectively.

6.2 Implications for Practice/Business

The findings of this research have several practical implications for the e-commerce and logistics industry, specifically in the domain of Attended Home Delivery (AHD):

- **Commercial Interest in AHD Operations:** The dynamic pricing method developed can be integrated into existing dynamic-routing packages without requiring additional learning or complex modifications, making it a commercially viable option for enhancing delivery efficiency and profitability.
- **Informed Pricing Decisions:** The integration of potential delivery price reductions into opportunity cost estimation supports more informed and optimal pricing decisions, thereby enhancing total profit and operational efficiency across various contexts.
- **Enhanced Customer Experience:** The introduction of augmented time windows not only improves the availability of delivery slots but also provides practical guidance for implementing such systems, potentially leading to increased customer satisfaction and loyalty.
- **Financial Gains and Operational Efficiency:** Leveraging forecast orders in dynamic routing decisions can result in considerable financial gains and improved system performance. Businesses can achieve better profitability and operational efficiency by incorporating the innovative strategies proposed in this research.

The integration of these methods into practice can lead to more efficient and profitable AHD systems, addressing both customer needs and operational challenges effectively.

6.3 Future Work

Building upon the findings of this thesis, several promising avenues for further research are identified:

1. Integration of Machine Learning in Dynamic Pricing: The current system as described in Chapter5 primarily employs a one-step look-ahead forecasting approach to assess and compute the DPR. This method operates under the assumption that the simplest yet effective strategy in MDP scenarios—like the one under discussion—is to predict the immediate future. However, integrating a multi-step forecasting strategy into the DPR calculations could notably improve the system's ability to anticipate future conditions. By incorporating advanced machine learning techniques e.g. reinforcement learning, the DPR system's complexity and effectiveness could be elevated. This enhancement might include multiple stages in the decision-making process, enabling a better understanding of customer behaviour patterns in slot selection. Such an understanding could then be leveraged to refine the dynamic pricing algorithm, allowing it to adjust prices more accurately for each time step.

2. Diversification of Artificial Order Sets: Exploring the diversification of artificial order sets drawn from various customer segments within a targeted operational area represents a promising avenue for future research. This approach aims to enrich the forecasting models by better mirroring the true customer characteristics of the area under analysis. Accurately reflecting these characteristics is crucial as it contribute to reducing the risk of underestimating or overestimating DPR calculations, thereby ensuring more reliable pricing strategies.

   One innovative method to achieve this involves developing a sophisticated technique for generating these artificial orders. This could be implemented by applying clustering algorithms to categorise the operational area based on order density. Subsequently, artificial orders can be randomly generated from these segments in proportions that correspond to the volume of historic orders observed within each segment. Such a method would not only yield a set of data that is representative of actual customer behaviour but also enhance the understanding of the DPR system's effectiveness and its adaptability to different market contexts.

3. Exploration of Alternative Dynamic Pricing Structures: The current study utilises a dynamic pricing system grounded in the MNL choice model, which adeptly predicts customer choices by evaluating the utility values linked to various pricing options. Although this method provides a solid foundation, the exploration of alternative dynamic pricing strategies, such as segmented pricing and time-based pricing, promises substantial optimisation opportunities within the swiftly evolving e-commerce delivery industry. Segmented pricing allows for tailored prices based on distinct customer demographics and purchasing behaviours, while time-based pricing adjusts rates according to fluctuations in demand during different times of the day or week, aligning price with anticipated customer activity.

4. Expansion to Varied Geographical and Operational Contexts: To thoroughly assess the effectiveness and scalability of the methodologies discussed in this thesis, it is essential to test them across a diverse array of geographical and operational contexts. Expanding the study to include various environments—from densely populated urban areas to more sparsely populated rural locations—will provide critical insights into how these methods adapt to different market dynamics and customer behaviours. Furthermore, exploring a variety of operational scenarios, such as different delivery infrastructures and local regulations, will help determine the robustness of the proposed methods. This broadened scope will ensure that the developed methods are not only versatile but also reliably effective in varied settings, thus affirming their practical utility in a global marketplace.

5. Longitudinal Studies and Real-Time Implementation: Conducting longitudinal studies and implementing real-time operations are important steps for assessing the long-term effects and sustainability of proposed methods. Longitudinal studies allow researchers to observe how these strategies perform over extended

periods, providing valuable insights into their efficacy, adaptability, and impact on consumer behaviour. Additionally, real-time implementation of these methods offers a direct insight into their practical viability, providing immediate feedback on how they perform under live market conditions. This approach not only helps in assessing the adaptability of the strategies as market dynamics evolve but also allows for iterative improvements based on continuous data collection and analysis. Such comprehensive evaluation techniques are essential for developing robust, sustainable pricing models that can withstand the test of time and fluctuating market conditions.

By addressing these facets, future research can not only reinforce the current understanding of the proposed methods' benefits but also extend their potential applications and refinements. This continued exploration is vital for keeping pace with the rapidly evolving field of e-commerce and the ever-changing demands of the global market.

A

Publication List

- Mohammad Abdollahi, Xinan Yang, Moncef Ilies Nasri, and Michael Fairbank. Demand management in time-slotted last-mile delivery via dynamic routing with forecast orders. *European Journal of Operational Research*, vol 309, Issue 2, pages 704-718, 2023. https://doi.org/10.1016/j.ejor.2023.01.023.

- Mohammad Abdollahi, Xinan Yang, and Michael Fairbank. Efficient forecast-based routing and dynamic time window management for attended home deliveries. Submitted to *Annals of Operations Research*, June 2024.

- Mohammad Abdollahi, Xinan Yang, and Michael Fairbank. Enhancing Opportunity Cost Approximation through Incorporation of Delivery Price Reduction.

Submitted to *IMA Journal of Management Mathematics*, May 2024.

Bibliography

Abdollahi, M., Yang, X., Nasri, M. I., and Fairbank, M. (2023). Demand management in time-slotted last-mile delivery via dynamic routing with forecast orders. European Journal of Operational Research.

Agatz, N., Campbell, A., Fleischmann, M., and Savelsbergh, M. (2011). Time slot management in attended home delivery. Transportation Science, 45(3):435–449.

Agatz, N., Campbell, A. M., Fleischmann, M., Nunen, J. V., Savelsbergh, M., van Nunen, J., and Savelsbergh, M. (2013). Revenue management opportunities for internet retailers. Journal of Revenue and Pricing Management, 12(2):1–19.

Agatz, N., Fan, Y., and Stam, D. (2021). The impact of green labels on time slot choice and operational sustainability. Production and Operations Management, 30(7):2285–2303.

Agatz, N. A. H., Fleischmann, M., and van Nunen, J. A. E. E. (2008). E-fulfillment and multi-channel distribution: A review. European Journal of Operational Research, 187(2):339–356.

Asdemir, K., Jacob, V. S., and Krishnan, R. (2009). Dynamic pricing of multiple home delivery options. European Journal of Operational Research, 196(1):246–257.

Bent, R. W. and Van Hentenryck, P. (2004). Scenario-based planning for partially dynamic vehicle routing with stochastic customers. Operations Research, 52(6):977–987.

Bühler, D., Klein, R., and Neugebauer, M. (2016). Model-based delivery cost approximation in attended home services. Computers & Industrial Engineering, 98:78–90.

Campbell, A. M. and Savelsbergh, M. (2006). Incentive schemes for attended home delivery services. Transportation Science, 40(3):327–341.

Campbell, A. M. and Savelsbergh, M. W. P. (2005). Decision support for consumer direct grocery initiatives. Transportation Science, 39(3):313–327.

Chiang, Wen-Chyuan Russell, R. A. (1996). Simulated annealing metaheuristics for the vehicle routing problem with time windows. Annals of Operations Research, 63(2):3–27.

Cleophas, C. and Ehmke, J. F. (2014). When are deliveries profitable? Considering order value and transport capacity in demand fulfillment for last-mile deliveries in metropolitan areas. Business and Information Systems Engineering, 6(3):153–163.

Dayarian, I. and Savelsbergh, M. (2020). Crowdshipping and same-day delivery: Employing in-store customers to deliver online orders. Production and Operations Management, 29(9):2153–2174.

Dong, L., Kouvelis, P., and Tian, Z. (2009). Dynamic pricing and inventory control of substitute products. Manufacturing & Service Operations Management, 11(2):317–339.

Ehmke, J. F. and Campbell, A. M. (2014). Customer acceptance mechanisms for home deliveries in metropolitan areas. European Journal of Operational Research, 233(1):193–207.

Figliozzi, M., Mahmassani, H., and Jaillet, P. (2007). Pricing in dynamic vehicle routing problems. Transportation Science, 41:302–318.

Fleckenstein, D., Klein, R., and Steinhardt, C. (2022). Recent advances in integrating demand management and vehicle routing: A methodological review. European Journal of Operational Research.

Gevaers, R., Voorde, E., and Vanelslander, T. (2014). Cost modelling and simulation of last-mile characteristics in an innovative b2c supply chain environment with implications on urban areas and cities. Procedia - Social and Behavioral Sciences, 125.

Ichoua, S., Gendreau, M., and Potvin, J.-Y. (2006). Exploiting knowledge about future demands for real-time vehicle dispatching. Transportation Science, 40(2):211–225.

Klein, R., Koch, S., Steinhardt, C., and Strauss, A. K. (2020). A review of revenue management: Recent generalizations and advances in industry applications. European Journal of Operational Research, 284(2):397–412.

Klein, R., Mackert, J., Neugebauer, M., and Steinhardt, C. (2018). A model-based approximation of opportunity cost for dynamic pricing in attended home delivery. OR Spectrum: Quantitative Approaches in Management, 40(4):969–996.

Klein, R., Neugebauer, M., Ratkovitch, D., and Steinhardt, C. (2017). Differentiated time slot pricing under routing considerations in attended home delivery. Transportation Science, 53.

Klein, V. and Steinhardt, C. (2022). Dynamic demand management and online tour planning for same-day delivery. European Journal of Operational Research.

Klein, V. and Steinhardt, C. (2023). Dynamic demand management and online tour planning for same-day delivery. European Journal of Operational Research, 307(2):860–886.

Koch, S. and Klein, R. (2020). Route-based approximate dynamic programming for dynamic pricing in attended home delivery. European Journal of Operational Research, 287(2):633–652.

Köhler, C., Ehmke, J. F., and Campbell, A. M. (2020a). Flexible time window management for attended home deliveries. Omega, 91:102023.

Köhler, C., Ehmke, J. F., and Campbell, A. M. (2020b). Flexible time window management for attended home deliveries. Omega, 91:102023.

Kumar, S. N. and Panneerselvam, R. (2012). A survey on the vehicle routing problem and its variants. Intelligent Information Management.

Lin, I. and Mahmassani, H. (2002). Can online grocers deliver?: Some logistics considerations. Transportation Research Record Journal of the Transportation Research Board, 1817:17–24.

Liu, F., Lu, C., Gui, L., Zhang, Q., Tong, X., and Yuan, M. (2023). Heuristics for vehicle routing problem: A survey and recent advances. arXiv preprint arXiv:2303.04147.

Mackert, J., Steinhardt, C., and Klein, R. (2019). Integrating customer choice in differentiated slotting for last-mile logistics. Logistics Research, 12.

McFadden, D. et al. (1973). Conditional logit analysis of qualitative choice behavior. Frontiers in Econometrics.

Mintel (2022). Online Grocery Retailing - UK - 2022. Technical report, Mintel Group Ltd, London.

Mintel (2023). Consumer Trends in Online Grocery Retail. Technical report, Mintel Group Ltd, London.

Ojeda Rios, B. H., Xavier, E. C., Miyazawa, F. K., Amorim, P., Curcio, E., and Santos, M. J. (2021). Recent dynamic vehicle routing problems: A survey. Computers & Industrial Engineering, 160:107604.

Punakivi, M. and Saranen, J. (2001). Identifying the success factors in e-grocery home delivery. International Journal of Retail & Distribution Management, 29:156–163.

Snoeck, A., Merchán, D., and Winkenbach, M. (2020). Revenue management in last-mile delivery: state-of-the-art and future research directions. Transportation Research Procedia, 46:109–116. The 11th International Conference on City Logistics, Dubrovnik, Croatia, 12th - 14th June 2019.

Soeffker, N., Ulmer, M. W., and Mattfeld, D. C. (2022). Stochastic dynamic vehicle routing in the light of prescriptive analytics: A review. European Journal of Operational Research, 298(3):801–820.

Strauss, A., Gulpinar, N., and Zheng, Y. (2020). Dynamic pricing of flexible time slots for attended home delivery. European Journal of Operational Research.

Ulmer, M. W. (2020). Dynamic pricing and routing for same-day delivery. Transportation Science, 54(4):1016–1033.

Voccia, S. A., Campbell, A. M., and Thomas, B. W. (2019). The same-day delivery problem for online purchases. Transportation Science, 53(1):167–184.

Wang, X., Zhan, L., Ruan, J., and Zhang, J. (2014). How to choose "last mile" delivery modes for e-fulfillment. Mathematical Problems in Engineering, 2014.

Waßmuth, K., Köhler, C., Agatz, N., and Fleischmann, M. (2022). Demand management for attended home delivery–a literature review. ERIM Report Series Reference Forthcoming.

Yang, X. and Strauss, A. (2017). An approximate dynamic programming approach to attended home delivery management. European Journal of Operational Research, 263:935–945.

Yang, X., Strauss, A. K., Currie, C. S. M., and Eglese, R. (2016). Choice-based demand management and vehicle routing in e-fulfillment. Transportation Science, 50(2):473–488.