Using stochastic and dynamic routing models to improve the performance of an e-grocer's delivery service

Master thesis

P. S. Bouwstra

TU Delft

Using stochastic and dynamic routing models to improve the performance of an e-grocer's delivery service

Master thesis

by

P. S. Bouwstra

In partial fulfilment of the requirements for the double degree consisting of the degrees of

Master of Science
in Mechanical Engineering

And

Master of Science
in Transport, Infrastructure and Logistics

Student number:	4350669	
MSc tracks:	Transport Engineering and Logistics	3mE
	Transport, Infrastructure and Logistics	CiTG
Report number:	2020.MME.8461	

Project duration:	Feb 17, 2020 – Nov 23, 2020	
Supervisors:	Prof. Dr. R.R. Negenborn,	Delft University of Technology
	Dr. B. Atasoy,	Delft University of Technology
	Dr. Ir. G. Correia,	Delft University of Technology
	Ir. P. M. Bijl	Picnic Technologies B.V.

An electronic version of this thesis is available at http://repository.tudelft.nl/.

TU Delft

Abstract

E-grocers are grocery markets which offer their assortment online and deliver a customer's groceries at home. For an e-grocer the quality of their delivery service is a crucial asset to create and maintain a loyal customer-base. Because the online share of the grocery market has grown rapidly over the last decade, there is an urgent need for e-grocer specific routing models.

In parallel, the research community has made rapid developments in the fields of stochastic and dynamic routing models during the past decade. Those studies suggest that the use of stochastic and dynamic elements improves the performance of routing models for a large variety of applications. However, it is yet unresearched how these stochastic and dynamic routing models can be used to improve the last-mile delivery for e-grocers specifically. The intended outcome of this research is the formulation of a promising modelling approach to use stochastic and dynamic routing model elements to improve the performance of an e-grocer's routing model.

In this thesis, first the objectives and requirements for an e-grocer's routing model are determined in more detail by means of a case-study at Dutch e-grocer Picnic. The strict requirements for an e-grocer's routing model are: (1) Offer customers a free one-hour delivery time window when they place an order. (2) Offer free communication of a 20-minute delivery time window at the morning of the delivery. (3) Calculation of the combinations of customers in a trip has to be completed within 45 minutes. (4) The trip planning has to be calculated within 5 hours. When these requirements are met, the objectives for the routing model are to maximize the on-time delivery performance and to minimize the operational costs. The on-time delivery performance is quantified by means of the on-time delivery rate and the rate of extreme lates (≥15 minutes late). The operational costs are quantified by means of the average time spent per delivery.

Next, scientific literature on the topic of dynamic and stochastic routing models is studied. Based on the findings in this literature study, it is suggested to make a distinction between "regular" deliveries and "flexible" deliveries. Regular deliveries have to be delivered within a 20-minute delivery time window. Flexible deliveries allow for larger delivery time windows. Once the trip has departed, the presence of these flexible deliveries allows for dynamic re-optimization of the customer sequence in that trip. This strategy has the potential to improve the on-time delivery performance. Moreover, this approach does not incur large additional operational costs because the number of vehicles used for delivery remains the same. At e-grocers, limited computation time is available for the calculation of the combinations of customers in a trip. However, plenty of time is available to determine the exact trip planning. Therefore, the traditional vehicle routing problem is split into two separate problems which are solved one after the other: The customer-trip assignment problem and the sequencing problem.

Then, different approaches for implementing the concept of flexible deliveries in an e-grocer's routing model are designed. Three different sequencing models are formulated to get insights in the benefits of a stochastic routing model: A deterministic sequencing model which resembles the current operations at Picnic (benchmark model), a stochastic sequencing model which uses a-priori simulation to choose the best solution out of a set of good solutions (simulation-based model), and a deterministic sequencing model which positions the flexible deliveries in the customer sequence of a trip in such a way that the effectiveness of re-optimization is maximized (heuristics-based model). In order to study the added value of a dynamic routing model, a deterministic re-optimization model is designed.

In order to quantify the performance of a routing model which makes use of flexible deliveries and compare the different model variants, four different routing model configurations are investigated by means of a computational experiment. A configuration consists of a sequencing model and, in some cases, a re-optimization model. The computational experiment uses real historic test instances from e-grocer Picnic to evaluate the performance of the configurations. In addition, historic data from Picnic is used to simulate realisations of planned trip times. Two different types of customer-trip assignments are investigated: customer-trip assignments for a small vehicle and for a large vehicle. Moreover, two different sizes of the flexible delivery windows are studied: 60 minutes and 75 minutes.

From the experimental results it can be concluded that the configurations including a re-optimization model outperform the static and deterministic benchmark configuration in terms of on-time delivery performance. However, re-optimization comes at the price of an increased average time spent per delivery. This can be explained by the fact that the re-optimization model optimizes the on-time delivery performance instead of the trip duration, resulting in larger average travel times. The effects of re-optimization become significant when 10% or more of all deliveries is flexible. For this percentage of flexible deliveries, the number of late deliveries can be reduced by up to 18% and the number of extreme lates can be reduced by up to 27%, depending on the type of vehicle and the size of the flexible delivery windows. The simulation-based sequencing model proves to be effective without re-optimization for the large vehicle type. For the small vehicle type it needs the re-optimization model to significantly outperform the static and deterministic benchmark configuration. When the re-optimization model is included in the configuration, the simulation-based and heuristics-based sequencing models show similar performance in terms of on-time delivery rate. The configuration with the simulation-based approach results in a lower average time spent per delivery. However, the heuristics-based approach results in fewer extreme lates.

In order to decide on the optimal routing model configuration for a specific e-grocer, the relative importance of the on-time delivery performance and the operational costs have to be specified. If an e-grocer is very cost-sensitive, the configuration consisting of the simulation-based sequencing model and the real-time re-optimization model would be optimal. However, if an e-grocer has more financial possibilities for an increase in operational costs the configuration including the heuristics-based sequencing model would be a better fit.

Extending the flexible delivery windows results in a significant improvement of the on-time delivery performance of the configurations. However, the increased number of successful re-optimizations comes at the expense of an increased average time spent per delivery. The added value of extended flexible delivery windows is largest for the small vehicle type. When comparing the two types of vehicles studied, it becomes evident that the use of stochastic and dynamic routing model elements is more effective for the large vehicle type than for the small vehicle type.

Preface

This thesis serves as the final project of my master's degree at the TU delft. It has been a great challenge to apply the knowledge that I have obtained in the courses I took during my studies. Managing a project of this size was a useful experience. I very much enjoyed the freedom that comes with a master thesis in terms of project topic and research approach.

However, my work would not have reached the same quality were it without the help of my project supervisors. I would like to thank Bilge for her assistance during my literature review and the debugging tips she provided whenever I struggled finding a fix for my broken code. Gonçalo helped me to improve the quality of my report by providing critical notes at each version I handed in. Moreover, I was supervised by Peter on behalf of Picnic Technologies B.V.. I very much valued our weekly meetings in which we brainstormed about the challenges I encountered. Thank you for your approachability and critical attitude towards my work which definitely helped me to perform at my best. Lastly, I would like to thank professor Negenborn for chairing my graduation committee and providing useful feedback during the official meetings.

Pieter Bouwstra
Delft, November 2020

Contents

1 Introduction

This chapter explores the context of e-grocer routing models and makes a case for the development of routing algorithms tailored to the needs of e-grocers. Some background information regarding e-grocers is provided after which the problem definition follows. The research questions are introduced along with the methodology that is employed to answer these research questions.

1.1 Context

In order to understand the context of this thesis, the concept of e-grocers and the different approaches to the home-delivery of groceries taken by various companies are introduced. Next, the relevance of the presented research is demonstrated by means of a discussion of the e-grocer market development, success-determining service qualities and e-grocer specific supply chain characteristics.

1.1.1. Concept of e-grocers

E-grocers offer their assortment of groceries online and deliver the customers' orders at home. They offer their customers a comfortable grocery shopping experience by allowing them to order from virtually any location and within moments. The groceries are delivered at the customer's front door or even into the kitchen. E-grocers use different pricing schemes to defray the costs of their delivery service; Some e-grocers accept orders above a certain threshold value, e.g. Picnic (2020), and offer free delivery. Others accept orders above a certain threshold value and charge an extra fee for delivery, e.g. Albert Heijn (2020) and Jumbo (2020). A third group of e-grocers only charges a delivery fee when the order value is below a certain threshold, e.g. Ocado (2020). Lastly, a group of e-grocers works with membership fees which unlock unlimited free delivery, e.g. Walmart (2020). Referring to the history of online grocery shopping, the first e-grocer businesses emerged in the late 1990s (Saunders and GlobalData, 2019a) and since the 2010s (Saunders and GlobalData, 2019b) e-grocers have started to seriously compete with traditional grocery stores. Because of the success of e-grocer start-ups, many traditional food retail market players have started to invest in an online grocery delivery service as well. Some of them completely separate the supply chains for the online and brick-and-mortar store customers while others see a potential in using their existing brick-and-mortar shops for the distribution of home-delivered groceries (Mkansi and Nsakanda, 2019).

1.1.2. E-grocer market analysis

The recent growth of the Dutch e-grocer market is illustrated in figure 1.1(a). The market share of online groceries in the total grocery market has increased by 325% between 2015 and the first half of 2019 (van der Weerd et al., 2019). Rabobank (2019) forecasts that the market share of the total Dutch grocery market occupied by e-grocers yields 15-20% in 2030. Figure 1.1(b) puts the development of the Dutch online grocery market in perspective with regards to other developed food retail markets. It can be concluded that the UK is ahead in terms of e-grocer market share. The Netherlands has seen the largest relative increase in e-grocer market share between 2016 and 2019.

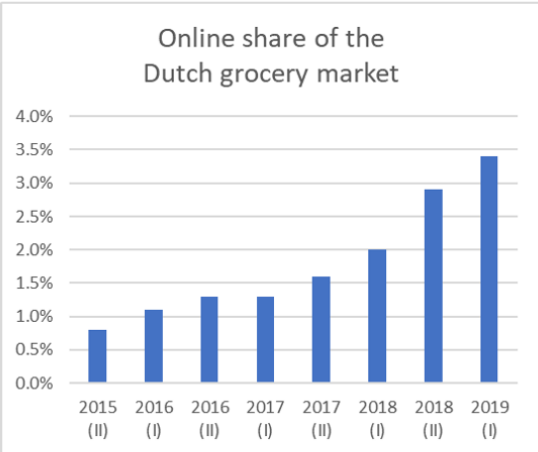

(a)

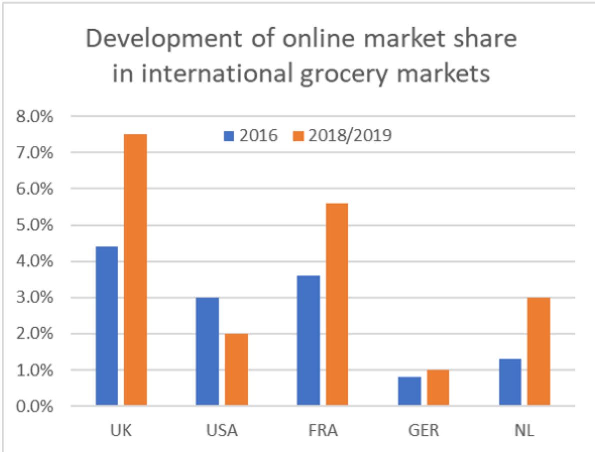

(b)

Figure 1.1: Development of the (a) Dutch e-grocer market and (b) the international e-grocer market (van der Weerd et al., 2019)

Joerss et al. (2016) stress that within the e-commerce industry, the quality of the delivery service is even more decisive for a company's success if their commodity is medicine or groceries. Wilson-Jeanselme and Reynolds (2006) studied the UK e-grocer market and point out that the reliability of the delivery service is one of the factors which determine whether online grocery stores successfully compete with traditional offline grocery stores. In their research on the European e-grocer market, Galante et al. (2013) discovered that for customers who have tried online grocery shopping, the most important reason to discontinue online shopping is the size of the delivery fee. They also point out that reducing these delivery costs would convince customers to retry online grocery shopping.

It can be concluded that the quality of the delivery service is a crucial asset for an e-grocer. Therefore, the research community has started to pick up on this (Warschun et al., 2012). For example, Mkansi et al. (2018) identified logistical elements inherent to e-grocers and revealed challenges related to each element. According to their study two of the areas in which challenges arise are communication with customers and last-mile delivery of orders.

1.1.3. E-grocer supply chain

The supply chains of most e-grocers consist of multiple fulfilment centres where the items for each order are picked. E-grocers mostly use a 2-stage distribution system in which the packed orders are shipped to a hub where they are transferred to smaller vehicles which complete the last-mile delivery (logistiek.nl, 2016). The scope of this study is limited to the part of the e-grocer supply-chain concerning the last-mile delivery to the customer.

The e-grocer market is relatively young and differs significantly from other more mature e-commerce markets. First of all, the commodity is highly perishable and requires strict cooling and careful handling. Moreover, not all products pose the same temperature constraints. A distinction has to be made between ambient, chilled and frozen products which have to be separated in conditioned compartments during transportation.

Although some e-grocers are experimenting with unattended home delivery (Paazl, 2018), most e-grocers depend on attended home delivery (AHD). AHD poses the challenge of alignment of the customer's and distributor's schedule. Because of the frequency of ordering, the size of the offered time windows is more important than for most other e-commerce. In addition, customers have to trust an e-grocer in terms of both product quality and delivery accuracy (Karahasanović et al., 2017). This is much more important than for the majority of other commodities, such as, for example, most parcels.

The consumer foods industry is a high volume low profit business (Creusen et al., 2008). This means that the average profit made by selling a product is small. For this reason the costs of the supply chain have to be kept at a minimum. In order to deliver groceries as cheaply as possible, e-grocers optimize their last-mile distribution system in such a way that the delivery costs per order are minimized.

1.2 Problem definition

E-grocers can achieve a competitive advantage by improving customer experience. Customer experience largely depends on the size of the delivery time windows communicated to the customers and the on-time delivery performance. As delivery time windows grow smaller, it becomes more challenging to maintain a good on-time delivery performance. Especially when the number of deliveries in a trip is large, the cumulative uncertainty of the travel- and service-times makes it challenging to complete the final deliveries in a trip on-time. At the same time, e-grocers have to keep the costs of their last-mile distribution system at a minimum in order to realise profit margins on the products they sell.

Because the e-grocer businesses have only recently started to gain a significant market share in the grocery market, not much research has been dedicated to the development of routing models for this industry specifically. Using an of-the-shelve routing model does not accurately take into account the peculiarities concerning the last-mile distribution of groceries, such as the perishability of the commodity and the importance of trust in the quality of an e-grocer's delivery service.

The research community has made rapid developments in the field of stochastic and dynamic routing models during the past decade (Oyola et al., 2018, Ulmer et al., 2017). The results of these studies suggest that the use of stochastic and dynamic routing models improves the performance of routing models for a large variety of applications. However, it is yet unresearched how these stochastic and dynamic routing models can be used to improve the last-mile delivery for e-grocers specifically.

The requirements and objectives for an e-grocer's routing model are determined based on a case-study at Dutch e-grocer Picnic. In the following chapters of this report, these requirements and objectives will be regarded as representative for the requirements which e-grocers in general set for their routing model. Picnic defines the general objective of their routing system as following: Further improve the **customer service** and minimize the **total travel time** while not sacrificing **safety** and **computation time**. Below the most important requirements for Picnic's routing model that follow from the case-study as presented in this chapter are summarized:

- Offer a free one-hour order time window for customers to choose when placing their order.
- Offer free communication of a 20-minute delivery time window to customers at the morning of the delivery.
- Calculation of the combinations of customers in a trip has to be completed within 45 minutes.
- The trip planning has to be calculated within 5 hours.

The most important objectives for the routing model are the on-time delivery performance, which is quantified by means of the on-time delivery rate and the rate of extreme lates, and the costs of the last-mile distribution system, which is quantified by means of the average time spent per delivery.

It is impossible to give a general definition for the relative importance of these two objectives. A routing model which results in a very good on-time performance but drastically increases the costs of operation is not a good routing model. Vice-versa, when the operational costs are very low but the on-time delivery performance is poor the routing model is no fit either. In other words, for each specific e-grocer case, a careful trade-off should be made between the performance of a routing model in terms of on-time delivery performance and operational costs. In this thesis multiple routing model variants are analysed based on these two objectives, using the outcomes of this research decision-makers should decide which routing model variant is best for their specific use-case.

1.3 Research objectives

The intended outcome of this research is the formulation of a promising modelling approach to use stochastic and dynamic routing model elements to improve the performance of an e-grocer's routing model. The effectiveness of the proposed modelling approach is assessed by means of an experiment that uses e-grocer specific test instances.

1.4 Research questions

The problem definition and research objectives lead to the formulation of the following main research question:

"How can stochastic and dynamic routing models improve the performance of an e-grocer's delivery service?"

Below, the sub-questions are presented. The answers to the sub-questions can be integrated to formulate an answer to the main research question.

1. What are the approaches for stochastic and dynamic routing models used in literature?
2. What is a promising approach to apply stochastic and dynamic routing models for e-grocers?
3. How does the proposed routing model perform?

1.5 Scientific gap

By means of the literature review conducted for this thesis (see chapter 2) an overview is created of the work done by other researchers in the fields of stochastic, dynamic and combined stochastic and dynamic routing models. The different approaches encountered in literature are tested against the specifications for an e-grocer's routing model. This literature review yields the promising solution approach of flexible deliveries, which is further investigated in this thesis. "Flexible" deliveries have a larger delivery time window than "regular" deliveries. These larger delivery time windows allow for dynamic re-optimization of the residual sequence of customers in a trip and thereby can mitigate the effects of running early or running late with respect to the trip planning. In this section a number of similar studies conducted by other researchers are introduced and the differences compared to the research presented in this thesis are explained.

Chen et al. (2018) research an approach to solve the capacitated vehicle routing problem with stochastic demands. They introduce the concept of premium customers. For those premium customers, the probability that their demand is met is larger compared to the other customers. This idea of making a distinction between customers is also investigated in this thesis. However, whereas Chen et al. (2018) make a distinction with regards of the probability of a customer's demand being met, the concept of flexible deliveries makes a distinction with regards of the delivery time window size.

Ng et al. (2017) investigate rerouting strategies which concern the exchange of certain remaining customers in a trip. They use real-time traffic data to make these adjustments to the trip planning. This thesis also investigates the added value of re-optimization of the sequence of residual customers in a trip. However, this problem involves time window constraints whereas the problem addressed by Ng et al. (2017) does not included time window constraints during the re-optimization stage.

Errico et al. (2016) investigate the vehicle routing problem with hard time windows and stochastic service times. The main difference between the problem they studied and the problem studied in this thesis is the complexity of the recourse action. Errico et al. (2016) investigate two simple recourse actions: skip the service at the current customer, or skip the service at the next customer. In the context of e-grocers such recourse actions are unviable. In this thesis, a more sophisticated re-optimization of the sequence of residual customers is investigated.

Table 1.1: Related literature

Author(s)	Problem type	Stochastic element	Dynamically re-optimized element(s)	Time constraints
Chen et al. (2018)	VRPSD[a]	Quantity of customer demand	None	None
Ng et al. (2017)	DVRP[b]	Travel times	Sequence of residual customers in a trip	None
Errico et al. (2016)	DVRPTW[c]	Service times	Skipping service at a customer	Yes
Taniguchi and Shimamoto (2004)	DVRPTW[c]	Travel times	1. Number of trips 2. Trip departure times 3. Sequence of customers in each trip	Yes
This thesis	DTSPTW[d]	1. Service times 2. Travel times	Sequence of residual customers in a trip	Yes

[a]Vehicle Routing Problem with Stochastic Demands
[b]Dynamic Vehicle Routing Problem
[c]Dynamic Vehicle Routing Problem with Time Windows
[d]Dynamic Travelling Salesman Problem with Time Windows

Taniguchi and Shimamoto (2004) look into the potential of using real-time traffic times for re-optimization of the allocation of customers to trucks and the sequence of customers in a trip. They found that incorporating real-time traffic information results in an improved reliability of the arrival times at customers. This problem differs from the problem studied in this thesis because it includes more degrees of freedom which can be optimized during re-optimization: the number of trucks, trip departure times and sequence of customers in each trip. In this thesis there is only one degree of freedom: the sequence of customers in each trip. This makes the problem more constrained and therefore it is more challenging to make a significant impact by means of re-optimization.

From the overview as presented in table 1.1 it can be concluded that the combination of using a re-optimization model to reconsider the optimal sequence of residual customers in a trip and the presence of time window constraints is precedented by Taniguchi and Shimamoto (2004). However, Taniguchi and Shimamoto (2004) only considers stochastic travel times while this thesis also considers stochastic service times. Moreover, this thesis focusses on trip instances specifically encountered by e-grocers. The other studies mentioned in table 1.1 address a general application of their problem type. Lastly, most researchers in the field of stochastic and dynamic routing models address a variant of the vehicle routing problem, whereas this thesis focusses on the travelling salesman problem. As a consequence the degree of freedom of the problem posed in this thesis is smaller compared to most studies found in the scientific literature.

1.6 Methodology

In order to answer the sub-questions and conclusively compose an answer to the main research question, the research steps as presented in figure 1.2 are followed.

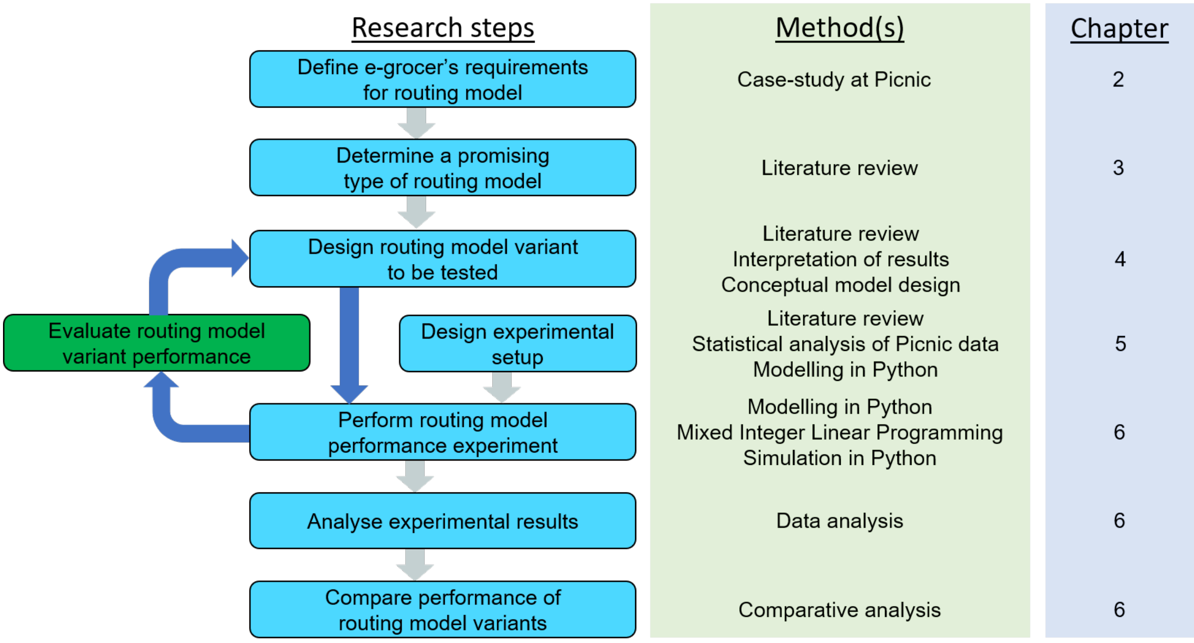

Figure 1.2: Research workflow

First of all, scientific literature on the topic of dynamic and stochastic routing models is reviewed. Different approaches to the use of dynamic and stochastic routing model features in different contexts are explored. This literature study leads to a solution approach which meets an e-grocer's requirements for its routing model.

Based on the solution approaches that result from the literature study, different types of routing models are designed via an iterative process; The starting point resembles the routing model as currently used by Picnic. The performance of this model is quantified by means of the experimental setup. The experimental results are analysed and evaluated in order to formulate a potentially improved version. The performance of that adapted routing model is then quantified again by means of the performance experiment. Via this iterative process promising model adaptations and/or additions are investigated and the impact of those changes on the performance of the model can be quantified effectively.

The experimental setup predicts how the tested routing models would perform when used in e-grocer operations. An important aspect of the experimental setup is the test instances. The test instances include the a-priori estimation of problem parameters and entail the expected travel time matrix and the expected service times for each delivery. Based on historic trip data from Picnic, realisations of planned trip times are simulated in the experimental setup.

Once all routing model variants are tested, their performances are compared. The experimental results provide an insight in the effectiveness of the concept of flexible deliveries in general. Moreover, it is discussed how the different routing model variants perform differently on the specified KPIs.

2

Theory on modelling approaches

This chapter provides an overview of the approaches to stochastic and dynamic modelling in the field of routing models. Sections 2.1, 2.2 and 2.3 together answer the following sub-question: "What are the approaches for stochastic and dynamic routing models used in literature?" Based on the overview of modelling approaches created in those sections, section 2.4 proposes a routing model tailored to the requirements and objectives of e-grocers as defined in chapter 1. Thereby sub-question 2 is answered: "What is a promising approach to apply stochastic and dynamic routing models for e-grocers?"

2.1 Stochastic modelling approaches

This section discusses the different approaches taken by researchers for modelling stochastic effects in routing models. First the different natures of the stochastic elements are discussed. Next, the different ways in which the stochastic elements can be taken into account are studied followed by a discussion of the types of probability distributions researchers use for these stochastic elements and how those distributions are obtained. The model performance evaluation is discussed after which the challenges that arise in the field of stochastic routing models and the possibilities that are yet to be explored are presented. The section is concluded by an analysis of the collected information in the light of its applicability for e-grocers. In this section papers are reviewed which were published after Gendreau et al. (2016) published their literature review on stochastic VRPs. For an overview of literature on this topic published before 2016, the work of Gendreau et al. (2016) can be consulted.

2.1.1. Nature of the stochastic element

In their literature review on stochastic VRPs Gendreau et al. (2016) distinguish three different natures of stochastic elements: stochastic demands, stochastic customers and stochastic times. Because the literature analysis performed in this report should ultimately provide insights for e-grocers specifically, only stochastic times are looked into. Stochastic demands are irrelevant in the context of e-grocers because the studied problem is delivery-only and the exact orders are known before departure. Moreover, stochastic customers are irrelevant because there is no uncertainty about the customers served in a trip. All customers whose groceries are placed in the vehicle are served. The presence/absence of customers is considered as integral to the service times. When a customer is not present, the driver will call the customer and try to find a solution. The time spent on these actions is included in the service time. For this reason the modelling approaches for two different stochastic elements are researched: stochastic travel times and stochastic service times.

Stochastic travel times

Different modelling approaches can be taken when the travel times are regarded as the stochastic element in a routing problem. The approach depends on the application of the routing problem studied. For example, the majority of the studied literature in this area focusses on routing problems with time-independent stochastic travel times. However, Varakantham et al. (2018) consider time-dependent travel times in the context of an orienteering problem. Their use-case is finding the optimal sequence of attractions in a theme park. They argue that the time-dependence plays a major role in terms of travelling and waiting time as the theme park is not equally crowded during the whole day. Similarly, one could argue that in other routing problems, such as parcel delivery problems, travel-times are time-dependent as well. However, the majority of the routing models does not take this time-dependency into account (Baradaran et al., 2019, Jaillet et al., 2016, Zhou et al., 2019). The main reason for a conservative attitude with respect to modelling time-dependent travel times is the computational challenges that arise when the problem is scaled (Chen et al., 2018). The use of time-dependent travel times drastically increases the demand for computer memory as well as calculation speed. For the modelling of stochastic travel times most researchers use an arc-based approach; A probability distribution for the travel time on each arc between delivery points is formulated. Then based on all possible arc travel times the optimal combination of arcs (trip) is computed (Adulyasak and Jaillet, 2016, Chen et al., 2018). Groß et al. (2016) take a different approach and divide the arcs between delivery points into smaller segments. They argue that logistics service providers can calculate these Interval Travel Times (ITTs) with relatively little effort compared to arc travel times by making use of the historic data they obtain.

Stochastic service times

Depending on the application, the service time at a customer comprises of different blocks. For example, in parcel-delivery the service time can be split into parking time, collecting the parcels from the vehicle, walking to the front door, waiting until the door is opened, handing-over the parcel, walking back to the vehicle and leaving the parking place. Researchers deal differently with the complexity of service times. Most researchers use aggregated service times including all phases of service. Other researchers, e.g. Han et al. (2017), take into account two different aspects of service times: the probability of customer show/no show and the customer's response time. In the field of stochastic service times, most researchers use the repairman use-case. In this use-case a repairman has to service a number of customers on a day without knowing the exact time required for each reparation. In most other use-cases it appears as if researchers of routing problems assume that the uncertainty of travel times is dominant over the uncertainty of service times. In some cases, a combination of stochastic travel times and stochastic service times is used (Ji et al., 2018). However, when the number of stochastic elements in a routing problem is increased, the maximum problem size is limited in order to keep the computational demand of the model within bounds.

2.1.2. How is the stochastic element taken into account?

There are different ways in which the stochastic element of a routing model can affect the solution of the routing problem. The most frequently used approaches are discussed in this subsection.

Risk minimization

In the risk minimization approach the risk is included in the objective function. A variety of methods can be adopted to achieve this. For example, Baradaran et al. (2019) use a multi-objective MILP in which they minimize the expected total costs and the variance of the total costs in separate objective functions. Other researchers choose an approach in which the variance of the objective metric is not taken into account explicitly. For example, Zhou et al. (2019) minimize the expected total delivery costs by means of predicting the recourse costs.

Han et al. (2017) include another interesting feature in the objective function of their Attended Home Delivery (AHD) problem: minimization of idle time, meaning the time that a driver has to wait until the customer answers the doorbell. It is mentioned that the model overestimates the time a driver has to wait for a customer. In the use-case studied by Han et al. (2017) it appeared taht the drivers rarely waste time waiting on customers. rather they leave before the customer has answered. Obviously the willingness of drivers to wait for a customer depends on the application and might be stronger in other use-cases.

Plenty of developments are going on when it comes to the definition of the objective function. Many researchers come up with their own performance indicator for risk which should be either minimized or maximized. For example, Groß et al. (2016) define a minmax regret criterion for their urban parcel delivery problem. This criterion measures the maximum regret that is obtained when a specific route is selected out of a set of routes. This maximum regret is calculated based on multiple simulations of each route in the set of routes. The route which results in the lowest maximum regret is selected as the optimal route. Jaillet et al. (2016) defined the “Requirements Violation (RV) index”. This index quantifies the total risk associated with the violation of requirements taking into account both the severity of the violations and their frequency. In the stochastic VRP with deadlines studied by Jaillet et al., these requirements translate into meeting delivery deadlines.

Chance-constrained modelling

Chance-constrained modelling was first introduced by Charnes and Cooper (1959) in the context of an inventory management problem. In most routing problems the ultimate objective of the distribution system is to minimize either the total costs or the total drive time, not the risk of arriving late. Therefore, in practice, most distributors set a risk-boundary. A maximum risk level can be captured in a routing model by means of a risk-constraint. This constraint ensures that the risk-level of a solution is maintained below a certain threshold which depends on the risk-attitude of the company. Alternatively, a customer-specific risk-attitude can be defined. For example, Chen et al. (2018) define “premium” customers who are exposed to lower risks of late deliveries. Varakantham et al. (2018) take into account two different risk-attitudes of the visitors of a theme park. They consider risk-seeking and risk-averting visitors. The objective is to visit as many attractions as possible whilst respecting a maximum risk of not being able to enjoy all planned attractions before the theme park closes.

Errico et al. (2018) use a chance-constrained approach because they argue that while “a stochastic program with recourse can indirectly decrease the failure probability of the a-priori plan by increasing the penalty cost associated with recourse actions, a chance-constrained model is more suited for a direct control of the failure probability”. In other words, they argue that chance-constrained models are more effective in limiting the failure probability. Ji et al. (2018) combine both methods, chance-constrained modelling and risk minimization, in a repairman problem. They use chance-constrained modelling on the working time of the repairman and minimize the sum of the following two route attributes: 1) the maximum risk of arriving late at a customer and 2) the summated risk of arriving late at each customer.

Multi-stage dynamic re-optimization

This type of stochastic modelling can also be referred to as stochastic and dynamic models. These models consider much more complicated recourse actions and are extensively discussed in chapter 2.3.

2.1.3. Probability distribution

In order to benefit from the use of stochastic elements in a routing model, the approximation of the probability distribution of the stochastic element is important. Researchers take different approaches to describe those probability distributions and obtain their data from different sources.

The majority of the travel time distributions is assumed to follow a normal distribution (Adulyasak and Jaillet, 2016, Baradaran et al., 2019, Ji et al., 2018). Baradaran et al. (2019) observed a normal distribution of travel times from historic data obtained from their use-case application: trucking distribution in Iran. Varakantham et al. (2018) assume a gamma distribution for the travel time. They realise that the distribution of the total trip time of a route is equivalent to the sum of the travel and waiting time distributions. A normal and gamma distribution are assumed for, respectively, travel and waiting times. They note that when a complex distribution is used, e.g. a gamma distribution with different thetas, sampling techniques can be used to obtain the cumulative travel-waiting times. Groß et al. (2016) use the 5^{th} and 95^{th} percentiles of a gamma distribution that is observed from artificial historic operation to set the boundaries of a uniform distributions of travel arcs which is used in the routing model. The boundaries of the uniform distribution can be adjusted depending on the risk-attitude of the distributor.

In general, most of the studies assume that distributions of individual travel times are mutually independent. However, as argued by Jaillet et al. (2016) this assumption is weak. When traffic congestion occurs in a neighbourhood, chances are that not just one travel-arc in a trip is affected. For this reason, traffic would be modelled more accurately when correlations between arc travel times are introduced. However, as mentioned by Jaillet et al., the independence assumption significantly simplifies the computation of a cumulative distribution. For example, Chen (2018) developed a solution technique that will solve a stochastic routing problem reasonably quickly when the distribution is log-concave.

The use of log-concave distributions simplifies models because a linear combination (convolution) of multiple log-concave distributions is log-concave as well. This means that a multiple integral can be reduced to an iterated integral. An iterated integral is a series of singular integrals which have only one stochastic variable. These singular integrals are easy to calculate and therefore calculation of the iterated integral requires modest computational power. When the distribution is not log-concave, most researchers use a sampling technique to evaluate the sum of the probability distributions (Chen et al., 2018, Ji et al., 2018). Monte-Carlo simulation is one of the most used techniques because it can be used for all probability distributions. In monte-carlo simulation random scenarios are created by means of a random sample from the probability distribution of each variable. When the number of evaluated scenarios is large enough, the cumulative probability distribution can be approximated. Depending on the number of stochastic parameters, many scenarios have to be evaluated.

The probability distribution of service times depends on the type of service that is offered. Chen et al. (2018) express the preference for a normal distribution in large datasets. The calculation will become easier and faster because a normal distribution belongs to the class of log-concave distributions. As Oyola et al. (2018) conclude in their literature survey, a wide range of probability distributions is used to describe stochastic service times. However, the vast majority of these distributions have in common that they are log-concave distributions. Ji et al. (2018) remark that the starting time of a service for their case-study will not be normally distributed because the driver always has to wait until the customer's time window has begun. According to this line of reasoning one could argue that service times do not follow a normal distribution because the extremely short service times under the normal distribution are never realised. For this reason, a truncated normal distribution would more accurately model service times. Alternatively, Errico et al. (2016) assume that service time probability distributions are of the type discrete triangular.

2.1.4. Model performance

In order to put the performance of their models in perspective, most researchers refer to Solomon's instances (Solomon, 1987). In his paper Solomon describes a method to generate test instances for deterministic VRP's with time windows. Consequently, Solomon's instances require an adaptation before they can be used to generate instances to test stochastic routing models. For example, Nguyen et al. (2016) test their stochastic VRP model with time windows using an adaptation of Solomon's instances. Gamma distributions are used to model the stochastic travel times. Both the shape parameter (α) and scale parameter (θ) are calculated as a function of the deterministic Solomon's travel resistance. Although different researchers use different adaptations of Solomon's instances, replication of a test instance is simplified when that test instance is based on a Solomon instance. Many researchers who propose a stochastic routing model compare the results of their stochastic model with those obtained by a deterministic version. For example, Nguyen et al. (2016) reports a decrease in lateness of 72% for their pick-up problem when stochastic instead of deterministic travel times are used. However, the number of vehicles required increases by 13% and the calculation time increased by 188%. Errico et al. (2018) studied the VRP with hard time windows and stochastic service times and conclude that the added value of a stochastic model compared to a deterministic model depends on the size of the delivery time windows. For smaller time windows (Solomon type 1) the fraction of on-time deliveries increases with 22%, whereas for larger time windows (Solomon type 2) the fraction of on-time deliveries remains the same.

2.1.5. Challenges & possibilities

One of the biggest challenges for the field of stochastic routing problems is the use of realistic probability distributions for the stochastic element(s). Until now the majority of the researchers assumed independent log-concave distributions. This assumption drastically reduces the computational power required to solve the problem. As Gendreau et al. (2016) mentioned as well, this simplification widens the gap between the studied problems and the real-life applications of routing problems. In order to increase the impact stochastic routing models can have on real-life operations, the probability distributions should better represent real applications.

It is debatable whether the mutual independence assumption is flawed for all stochastic elements. For example, as also argued by Gendreau et al. (2016), the assumption might be reasonable for stochastic service-times. In order to reduce the gap between science and practice the nature of the stochastic element should be investigated when formulating a stochastic routing problem and based on that investigation the validity of the independence assumption can be assessed.

Recent developments in the availability of ICTs and the possibilities to collect and structure data provide opportunities for a more realistic representation of the probability distribution for a specific use-case. The most accurate distributions are obtained when based on the historic data of that exact same operation. These historic data, for example, can also demonstrate whether the mutual independence assumption for a specific stochastic element holds; The correlation between travel time distributions in an area can be investigated by means of comparing trip data on different arcs in that area and during the same time period. Moreover, depending on the application a customer-specific service time probability distribution could be deducted from historic data, resulting in a more realistic model. Another example is a driver-specific travel-time distribution.

Dropping the independence assumption of travel times provides possibilities for a more accurate prediction of the second half of the driving arcs in a trip after the first half of the trip is completed. In a dynamic routing problem this correlation between arc travel times in the first and second half of a trip could be used in order to more accurately compute the optimal route for the second half of the trip.

The evaluation of cumulative probability functions poses challenges. The challenge of the exact computation of a cumulative probability function is the primary reason for the use of simplified probability functions by most researchers. Some advanced techniques for the evaluation of complicated cumulative distributions have emerged, but it appears as if the most used solutions are simulation tools. For example, monte-carlo simulation is widely used to compute complex distributions. These simulation tools are less accurate than analytic evaluation but allow for a more realistic problem description.

Many stochastic routing problems minimize the total operating costs whilst taking into account the probability for the need of a recourse action and its associated costs. The complexity of the recourse actions encountered in literature is limited. Mostly, a simple penalty for arriving late is applied. Alternatively, in a pick-up problem the costs of a return to the depot are taken into account in case of a shortage of capacity to meet the demand of a customer

As predicted by Oyola et al. (2018) the interest in multiple-objective routing problems is increasing (Baradaran et al., 2019). This trend can be further explored in order to balance the variety of objectives inherent to a real-life routing problem. An example of a combination of objectives could be on-time delivery rate and transportation costs. In many cases, specifying a constraint for one of those is not effective as the maximum on-time delivery rate that can be realised at a reasonable cost is different for each trip.

As Gendreau et al. (2016) also mention in their literature review, only a limited variety of recourse actions has been considered so far. For example, the skipping of customers is only rarely considered. Moreover, to the best of our knowledge, there is no literature on stochastic routing models that consider changing the sequence of customers in a trip as the recourse action while calculating the optimal a-priori route. Lastly, it is very rare that researchers consider recourse actions that involve multiple vehicles. An example of such recourse action could be the following; Given that a parcel delivery vehicle is running late, the driver could meet with a nearby vehicle of the same fleet and exchange a few parcels in order to increase the on-time delivery rate.

Lastly, it is interesting to note that most stochastic routing models only consider one stochastic element. This was also noted by Braekers et al. (2016). The only combination of two stochastic elements that was found in the studied papers was in the paper by Ji et al. (2018) who considered both stochastic travel times and service times. As mentioned by Braekers et al. (2016), the use of multiple stochastic elements further burdens the tractability of the problem. Consequently, this limits the maximum problem size.

2.1.6. Applicability for e-grocers

In order to assess the applicability of the discussed stochastic modelling approaches they are examined based on the requirements for an e-grocer's routing model as specified in chapter 1. The purpose of this section is to discuss how stochastic routing models can be used to improve an e-grocer's routing model.

E-grocers encounter uncertainty in both travel times and service times. In a typical e-grocer trip, travel distances between customers are short resulting in a significant fraction of the total trip time spent at customer service. Therefore it is relevant to take into account stochastic travel times as well as stochastic service times. According to the studied papers, it is reasonable to assume that service times at different customers are uncorrelated. This offers opportunities to use a log-concave function for the probability distribution of the service time which would relieve the computational burden of the summation of stochastic elements significantly. On the contrary, the independence assumption for travel times does not hold and that raises computational challenges. It makes one question whether the performance improvement due to the use of stochastic travel times outweighs the increased computation time. Looking at e-grocer trips, the stem times in particular depend on the traffic conditions. Therefore it might be interesting to only take into account stochastic travel times for the modelling of stem times. Even more so because the independence assumption for stem-times is stronger.

For e-grocers one of the ways to increase the quality of their service is an improvement of the on-time delivery rate. Through the use of stochastic elements the chance of late arrival at a customer can be taken into account when calculating the optimal trips, see section 2.1.1. Inclusion of lateness in the model by means of a chance-constrained approach is complicated because the lowest feasible risk-level depends on characteristics of the area in which the trip is performed. For example, considering stochastic service times, the uncertainty in areas with many apartment buildings might be larger compared to areas with many bungalows. Considering stochastic travel times, in densely populated areas travel times will be more uncertain than in rural areas because of the unpredictability of traffic. Because the determination of a feasible risk constraint for each specific delivery area is a challenging task, risk minimization by means of including the probability of arriving late in the objective function seems to be advantageous. For example, this could take shape as a lateness-index for each trip which is assigned a weight in the objective function.

Most e-grocers have access to historic trip data. This trip data can be used in order to formulate the probability distribution for a stochastic element. The large majority of an e-grocer's customers orders repetitively. This phenomenon could be used to obtain a customer-specific service time distribution. The probability distribution for travel times could also be obtained based on historic data. For example by means of aggregating travel times on area level.

The determination of an effective classical recourse action in case of a delay is challenging. Skipping customers or a shortened service time is not an option from a customer experience perspective. This leaves a more complicated re-optimization approach as the only satisfactory strategy. How re-optimization can be used by e-grocers is discussed in section 2.2.7.

2.2 Dynamic modelling approaches

In dynamic routing models, new information becomes available during the execution of a trip. This information is used to re-optimize the residual trip. Because this calculation takes place during the trip, dynamic routing models are also referred to as real-time or online routing models. The number of papers on dynamic routing models has increased rapidly over the past decade because of technological developments that allow for efficient communication of the dynamic element in those models (Pillac et al., 2013). In this section a variety of dynamic routing models is discussed based on the nature of their dynamic element, moment of reconsideration of the residual route, rerouting criteria, dynamically optimized element and model performance. The section again concludes with an analysis of the current research status, challenges and possibilities, and applicability in routing models for e-grocers.

2.2.1. Nature of the dynamic element

The dynamic element in a dynamic routing model is the element which is revealed over time. The category of dynamic routing models which has received most attention in literature is the category of dynamic customer requests (Pillac et al., 2013, Psaraftis et al., 2015). These could be either demands for goods or services. For example, Sarasola et al. (2016) consider the dynamic VRP with dynamic service requests which come in while the vehicle has already departed on its trip. These new requests have to be implemented in the current trip or in a new trip executed by another vehicle while it is the objective to minimize both total travel time and late returns at the depot. Ulmer et al. (2017) suggest to solve the dynamic VRP with stochastic service requests by means of a special type of sequential decision model. They introduce a route-oriented approach as a better alternative to the commonly used customer-by-customer oriented approach. The customer-by-customer approach entails that during each moment of re-optimization only the next customer to serve is computed. They argue that the route-oriented approach is more accurate than the customer-by-customer approach. The larger computation times inherent to a route-oriented approach are mitigated by means of combining an offline value function approximation with a roll-out algorithm.

Han et al. (2017) study an Attended Home Delivery (AHD) problem with soft time window constraints. Their focus is on the customer behaviour and they distinguish two types of behaviour: no-show and random response time. They investigate a model in which the driver adjusts the service start time and maximum waiting time for the next customer based on the departure time at the current customer; If the response time of the current customer is longer than planned, the maximal waiting time at the next customer will be shortened in order to arrive at the second next customer on time.

Vodopivec and Miller-Hooks (2017) consider a special routing problem in which multiple customers are picked up and dropped off at various locations by a single vehicle. Their customers are elderly people who rely on taxi-services for transportation. In case the vehicle is running late, there is a possibility to call a third-party taxi service to bring the customers to their destination in time. The authors equip optimal stopping theory to determine the need and moment of stopping based on the progress of a trip.

A topic that is inherently related to dynamic customer requests is the taxi-routing problem. Amongst others, this problem is studied by Bertsimas et al. (2019). They study the assignment of a new request to a vehicle based on the real-time location of the taxis in the fleet and the nature of the request. For example, when a taxi is on its way to pick up customer A, a new request from customer B might come in. The developed routing model decides whether the taxi still picks up customer A or prioritizes customer B. This decision depends on the nature of the request (trip length and revenue) and the positions of other taxis in the fleet (possibility for another taxi to pick up customer A or B).

One of the simplest forms of the dynamic routing problem is the choice of paths to take between a predefined pair of customers or service locations. For example, this problem is studied by Ng et al. (2017). The model they develop monitors the real-time traffic situation and uses that information to determine the optimal path to the next customer upon departure. The model is expanded as following: the possibility to dynamically exchange customers in the sequence of remaining customers is considered in order to reduce the uncertainty due to road traffic.

2.2.2. Dynamically optimized element

The dynamically optimized element in a dynamic routing model is the element of the trip that is reconsidered during re-optimization. The mostly used dynamically optimized element is accepting a new request or not. This trip element can be studied in different settings depending on the application of the routing model. It is encountered in applications such as taxi companies (Bertsimas et al., 2019), maintenance services (Borenstein et al., 2010) and pick-up services (Wu et al., 2019). In some of these problems, it is only possible to add extra customers to a trip. In others, new requests replace previously scheduled customers. Lastly, in some models there is the possibility of using other vehicles in addition to the currently deployed vehicle(s) to serve new requests. This poses the challenge to decide when it becomes profitable to deploy an extra vehicle. When an extra vehicle is deployed real-time re-optimization is required to redistribute all customers over the new set of vehicles in an optimal way.

Another category of dynamically optimized elements concerns changing the path between customers. In this type of problem the sequence of customers in a trip remains mostly unchanged. However, the paths between the customers are optimized upon departure based on the real-time traffic situation. The effectiveness of such a re-optimization strategy depends on the travel distances between the customers. If the real-time delays on alternative paths are uncorrelated this re-optimization strategy might be effective. However, if the delays are correlated this approach is not likely to yield significant positive effects. This type of dynamically optimized elements could be more effective when the sequence of customers can be re-optimized as well (Ng et al., 2017). However, the possibility to do so depends on the customers' time window constraints.

Lastly, a broadly studied family of models considers the decision to return to the depot as the dynamically optimized element. For example, Pandelis et al. (2013) study a pick-up problem in which the customers' exact demand is revealed upon arrival at the customer but a probability distribution of the demand is known a-priori. Considering the maximum capacity of the vehicle, the dynamic routing model re-optimizes the moment to return to the depot to unload the vehicle. Obviously, a trip back to the depot is considered as a penalty. However, arriving at a customer and not being able to meet the demand is penalized as well. Two distinct variants of this problem exist: with and without the possibility to split the demand of a single customer over multiple vehicles or runs (Lee et al., 2006).

2.2.3. Moment of reconsideration

Many dynamic routing problems make use of periodic re-optimization. In this case, an optimal route is computed before departure of the trip. At specific moments in time during the trip, the residual of the trip is re-optimized. These moments in time are referred to as decision epochs (Chen and Xu, 2006) or time slices (Kilby et al., 1998). Depending on the application of the routing model the re-optimization frequency is determined. For example, Liang et al. (2020) developed a taxi routing model in which re-optimization is initiated after set intervals of time. They adopt the rolling horizon approach to limit the computational power required to find the optimal routing solution for the near future in the context of a dial-a-ride problem with automated taxis. In this case the appearance of new requests does not affect the re-optimization frequency directly.

Alternatively, for the majority of the delivery or maintenance service problems, the residual route is re-optimized when a specific event happens. For example, in the context of Pick-up and Delivery Problems (PDPs), re-optimizations are mostly initiated by new requests. Bertsimas et al. (2019) explain that the trip planning in their taxi routing model is reconsidered every time a new customer requests appears. This type of dynamic models is also referred to as continuous re-optimization models as opposed to periodic re-optimization models.

However, some authors, for example Ghiani et al. (2009), observe inefficiency in the fact that the total residual route is recalculated at each moment of reconsideration. Therefore, they propose solution algorithms which only determine the optimal action on a limited horizon. In the PDP studied by Ghiani et al. (2009) this means that at each re-optimization the next X customers out of a set of more than X to-be-served customers are determined. An approach like this reduces the calculation time which is a crucial factor in dynamic routing models because of their real-time nature. However, this type of model can not benefit from the experience in static routing models as much as models in which the total residual trip is re-optimized at each moment of reconsideration.

As opposed to running a routing model in real-time, some authors prefer the definition of a set of states in which a recourse action has to be taken. This set of states can be determined a-priori. For example, Vodopivec and Miller-Hooks (2017) define a time-distance function which forms the boundary between either continuing to follow the current trip planning, or calling a taxi to take over a customer. This type of dynamic routing model does not take into account real-time inputs other than the progress of a trip. Therefore this approach does not benefit from the more reliable prediction of, for example, travel times when the moment of departure gets closer.

Ulmer (2020) proposes a concept in which a state function similar to the one used by Vodopivec and Miller-Hooks (2017) is combined with real-time re-optimization and test it on a problem with stochastic service requests. A model in which the optimal next X customers in a trip are determined based on a re-optimization algorithm is proposed. The other customers in the residual trip are determined according to the a-priori defined state function. By taking this approach the re-optimization time is drastically decreased whilst still the output is a whole residual trip instead of a mere appointment of the next (couple of) customer(s) to visit.

2.2.4. Rerouting criteria

At the moment of reconsideration the routing model determines whether or not the current route has to be changed. In other words, should the previous route be replaced by the re-optimized route? In literature not much attention is given to this topic. Most dynamic routing models do not penalize a change of route and therefore simply adopt the optimal route whenever it is calculated. However, there might be cases were changing the route appears disadvantageous. For example, when a customer is notified about the planned arrival time each time when it changes, it might be confusing when the planned arrival time is changed multiple times. In this case, the customer experience might benefit from fewer re-optimizations (and therefore perhaps a suboptimal route) resulting in less confusion about the arrival time.

2.2.5. Model performance

For dynamic routing models there is no benchmark set of test instances which is used by many researchers. Researchers rather define their own test instances (Jaillet et al., 2016, Nakagawa et al., 2017, Ng et al., 2017, Ulmer, 2020, Wu et al., 2019). For example, Ng et al. (2017) adopt the test instances as defined by Augerat (1995) for their online VRP with deadlines. They compare the performance of a static and dynamic routing model using traffic conditions as measured in the city of Hong Kong during 6 consecutive days. It is found that the reduction of average travel time due to re-optimization equals 18%. In this problem heavy congestion is assumed; The free flow travel times are approximately half of the shortest simulated travel times. Similarly Nakagawa et al. (2017) use multiple types of test instances for their numerical experiments. They investigate a request acceptance model for a large capacity network service which is aimed at increasing the fairness of blocking a request. An example of such a network is the demand of electricity from an electric grid or the demand of transportation from a road network. For their numerical experiments a 4x4 polygrid network, Japan's national road network and the COST266 road network (standardized European road network) are used. They experiment with a static and dynamic version of their algorithm. The dynamic version redirects flows through the network in order to redistribute the network load over its arcs. It can be concluded that the version without re-optimization always exhibits a larger request blocking probability. However, when the network load increases the benefits from using a dynamic model decrease.

2.2.6. Challenges & possiblities

As already recognised by Pillac et al. (2013), dynamic routing models could be improved when stochasticity is taken into account during re-optimization. In the last couple of years only few papers that feature this combination of routing model characteristics appeared. Many researchers take stochastic elements into account when calculating the optimal a-priori route. However, for the re-optimization inherent to dynamic models, mostly deterministic variables are used. The most promising explanation for this approach is the time-pressure on re-optimization in dynamic routing models (Sarasola et al., 2016) while the use of stochastic variables increases the calculation time drastically (Oyola et al., 2017). Pillac et al. (2013) state that when designing a dynamic routing model a trade-off has to be made between computation time and decision quality. Nonetheless, the use of stochastic variables for the re-optimization steps would be a promising step forward to increase the performance of dynamic routing models. Examples of approaches to make this combination work are the shortening of the decision horizon of dynamic models, or changing the moment of reconsideration. For example, in a delivery problem, the residual of a trip can be re-optimized upon departure from a customer or upon arrival at a customer. When the latter is chosen, the computation time of the model would be less critical. Section 2.3 presents a more elaborate overview of the combination of stochastic and dynamic features in routing models.

Currently it is common practice to build in some extra time (slack) in routing schedules to compensate for unforeseen delays. This slack reduces the chances of arriving late but also increases the costs of the system because unnecessary waiting times are costly. Amongst others, Vodopivec and Miller-Hooks (2017) recognise that ride-sharing service providers might be able to reduce unnecessary slack in their schedules through intelligent re-optimization. Obviously, this improvement of a schedule is not limited to ride-sharing service providers only, but might yield benefits to logistic service providers or maintenance companies in general.

Alternatively, the computation time of a dynamic routing model can be reduced by means of the intelligent use of offline optimization. This means that a-priori a state function can be defined which determines the type of trip adaptation (if any) depending on the state of the system; Different scenarios are defined together with their re-optimization strategy. Online and offline re-optimization can be combined as following: When using both, the online re-optimization can be employed for the short term horizon whilst the effects for the other customers in the trip can be calculated by means of state functions. Amongst others, Yildiz and Savelsbergh (2020) studied this approach.

2.2.7. Applicability for e-grocers

E-grocers have a complex supply chain of which the last-mile delivery is the final phase. During the last-mile delivery, uncertain parameters are revealed over time. The most important parameters are the travel times and the service times because together those describe the progress of a trip. The progress in a trip can be the dynamic element in an e-grocer's routing model and can be used as an input for re-optimization. During that re-optimization other dynamically revealed parameters might be used, for example, the real-time traffic conditions and the positions of other delivery vehicles.

Based on the studied literature there is a variety of aspects of a trip that can be re-optimized in real-time. Nearby vehicles could redistribute orders if one vehicle is running early and another vehicle is running late. Moreover, there could be a possibility to use a back-up vehicle and driver who are on stand-by and take over orders from vehicles which are running late. However, having drivers and vehicles on stand-by is not feasible for e-grocers due to the increased operational costs incurred by that extra capacity. Moreover, the 20-minute time window constraint challenges the redistribution of orders over nearby vehicles as this redistribution would consume valuable time. This suggests that the exact re-optimization strategies as found in literature do not fit e-grocers perfectly. However, other suitable strategies can be deduced from them.

The preferred moment of re-optimization depends on the nature of the dynamically re-optimized trip element and on the time it takes to complete re-optimization. The routing model is inefficient if idle time is added to a trip unintentionally because the driver has to wait until re-optimization has finished. Therefore, depending on the calculation time the re-optimization could be initiated a couple of minutes before the solution will be used.

Most dynamic routing models place no burden on changing the planned trip. In most applications changing the trip does not affect the customers. However, e-grocers communicate a 20-minute delivery time window to customers. If a customer receives multiple notifications with mutations of their delivery time window this would negatively affect the customer experience. A possible solution is to limit the number of trip schedule changes. Alternatively, not after each change a notification is sent to the customers. Instead, shortly before the delivery time the customer receives a notification.

2.3 Stochastic and Dynamic modelling approaches

This section intends to provide insights in how the research community approaches the combination of stochastic elements and dynamic re-optimization in a routing model. Sarasola et al. (2016) mention that combined stochastic an dynamic VRPs have not received much attention in literature until recently. Most research on this combination has been performed in the context of Dial-A-Ride Problems (DARP) because of the inevitably dynamic nature of the incoming requests. This section first introduces a few examples of different types of optimization studies along with the natures of the stochastic and dynamic elements they researched. Secondly, the authors' motivation for combining these model features is analysed. Then the interaction between the dynamic and the stochastic element is explained. Because of the complexity of including both stochastic elements and dynamic re-optimization, computation time often poses a challenge for this type of models. The solution approaches used by researchers to overcome this challenge are explored in subsection 2.3.4 followed by the challenges and possibilities in this field of research as recognised by a variety of researchers. The section is rounded off by an analysis of stochastic and dynamic modelling approaches in the light of application in an e-grocer's routing model.

2.3.1. Natures of stochastic and dynamic element

As mentioned above, the most frequently studied models featuring a combination of stochastic elements and dynamic re-optimization are used to solve a DARP. Most of those models feature stochastic travel times and/or requests. Li et al. (2019), for example, consider both stochastic time-dependent travel times and stochastic requests in the context of a van-pool service. They study the trade-off between intelligent prepositioning of vehicles in areas where future demand is expected and serving existing requests. Therefore, the dynamically optimized element in their routing model is the next customer to serve, or the area in which to preposition a vehicle. Another example of a DARP is the problem studied by Schilde et al. (2014). They use a network-consistent stochastic time-dependent travel time layer as a method to generate time-dependent travel times that are guaranteed to be network-consistent. They focus primarily on travel times affected by the occurrence of accidents. The probability of an accident is based on historic data and aggregated on a level of city districts. A combination of dynamic and static requests is used. It is assumed that rejection of requests is not allowed and that the objectives are threefold and ranked based on importance: (1) Minimize the sum of lateness, earliness and ride time violations. (2) Minimize the number of vehicles used. (3) Minimize total route duration. Pandi et al. (2019) study a special model to solve the DARP in which vehicle breakdowns are modelled as stochastic events. The occurrence of a vehicle breakdown is modelled as a new request with its pick-up position at the location where the vehicle broke down. The new request is considered as the dynamic element and has to be served within the constraints of the previously solved problem (before the vehicle broke down).

Alternatively, Toriello et al. (2014) looked into the TSP with stochastic arc costs following a distribution which is known a-priori. The salesman is able to observe the outgoing arc costs at each city before deciding on which city to visit next. Therefore the dynamically revealed element is the arcs costs and the dynamically optimized element is the next city to visit. This problem is similar to the problem studied by Yu and Yang (2019) in which real-time traffic information is used to select the next customer to visit. Grippa et al. (2019) study a problem in which a number of drones delivers goods from a set of depots to customers. According to their problem definition, the customer requests come in dynamically over space and time. In essence they solve a version of the job assignment problem as the drones serve only one customer in each run. The dynamic element is the assignment of a vehicle to a customer once the customer request comes in.

Güner et al. (2017) study the modelling approaches for stochastic and dynamic elements in the context of milk-run tours; In this type of problem similar trips are completed on a regular basis. This approach can be taken when the customers prefer regularity concerning the time of their delivery. Güner et al. (2017) study this problem in the context of just-in-time production for which regularity and predictability of the delivery of components is important. The considered stochastic element is the time-dependent travel times. The dynamically optimized element is the set of arcs traversed from one customer to the next. This means that the sequence of customers is fixed from the start of the trip. However, the optimal path between predetermined customers is computed in real-time. The residual of a route between two customers is reconsidered at the end of each traversed arc. So the route between a pair of customers is re-optimized several times.

2.3.2. Why is this combination useful?

The distinction between recourse actions and re-optimization might be ambiguous for some readers. Recourse actions are predetermined actions to undertake when the realisation of a stochastic element takes a certain value. For example, a recourse action might be to skip a customer or the service at a customer when the progress in a trip lacks behind schedule (Errico et al., 2016). In a nutshell, recourse actions are uncomplicated response actions and are completely defined before departure. When X happens, undertake action Y. Dynamic re-optimization in the essence is a very complicated recourse action. When X happens perform re-optimization according to Y. Often but not always re-optimization uses more real-time parameters than the determination of a recourse action. Due to the complexity of re-optimization and the use of more real-time system parameters, dynamic re-optimization is more computationally challenging (Vangipurapu et al., 2019). Even more so given the fact that the computation has to be executed in real-time resulting in limited available calculation time.

Re-optimization is an intelligent response to the realisation of uncertain parameters. In order to respond as intelligently as possible, the uncertainty about the future should be taken into account as well. For this reason, the model becomes more realistic when the uncertainty of future parameters is taken into account during re-optimization. This allows for a re-optimized solution that anticipates on events which might occur in the future. Models in which stochastic elements are combined with re-optimization have proven to be superior to both only stochastic and only dynamic models in the fields of parcel delivery (Archetti et al., 2020), general VRP (Yu and Yang, 2019) and general TSP (Toriello et al., 2014).

2.3.3. Interaction between dynamic and stochastic elements

There are two ways to combine stochastic elements and dynamic re-optimization. Firstly, one can use stochastic elements for computing the optimal a-priori route and use deterministic parameters for re-optimization. One can question the validity of such an approach. However, it can be argued that at the moment of re-optimization the uncertainty of the stochastic parameters has reduced to a point where a deterministic modelling approach is sufficient. For example, travel times are uncertain a couple of hours ahead. Upon departure from a customer, the real-time travel times are a good approximation of the travel times to be encountered. Alternatively, stochastic parameters can be taken into account during re-optimization in order to capture the uncertainty that is still left. The decision to use stochastic or deterministic parameters for re-optimization depends on the constraints on calculation speed, availability of data and length of the re-optimization horizon.

Most of the studied DARPs use stochastic elements in the re-optimization stage. Depending on the variant of the problem the stochastic element is either the emergence of additional requests or the travel times. In most DARPs a short horizon is considered (Ho et al., 2018). This means that in the re-optimization step only the current requests and in some cases the near future expected requests are taken into account. The assignment of requests to vehicles depends on the real-time traffic conditions (Li et al., 2019). Based on historic data of the city of Vienna, Schilde et al. (2014) argue that a realistic DARP in an urban environment requires time-dependent stochastic travel times.

When studying the delivery of parcels by means of drones, Grippa et al. (2019) recognised that the timing of the job assignments (re-optimization) is challenging. On the one hand, the re-optimization should be performed as soon as a new request comes in because this enables the company to serve the new customer as soon as possible. On the other hand, from a system's perspective, it might be beneficial to postpone the re-optimization until more new requests have arrived. In their specific case this means that re-optimization would be performed when a drone returns to a depot from where goods are shipped to customers. In the general dynamic and stochastic routing problem there is a trade-off between early re-optimization and postponement until more information has become available (for example, more new requests in the case of Grippa et al. (2019)). In road-oriented routing problems postponing the moment of re-optimization results in a more accurate prediction of the travel times (Yu and Yang, 2019).

Li et al. (2019) take a different approach to determine the moment of re-optimization. In their van-pool routing problem they suggest that re-optimization can be initiated by two events. Re-optimization is initiated when a new request has come in or after a given time interval as elapsed. Periodic re-optimization prevents a van from waiting for a request in the same area for a long period of time. During re-optimization the potential demand in another area might be higher causing the optimal waiting area for the van to change.

A variety of authors (Güner et al., 2017) recognise the potential that Intelligent Transport Systems (ITS) provide for the improvement of routing models. They advocate for the exploitation of real-time traffic data from ITS for the purpose of dynamic re-optimization.

2.3.4. Computational power

The combination of online re-optimization and stochastic model inputs poses challenges in terms of a the computational power required to run the model. Especially when stochastic parameters are considered during re-optimization, the computation time of the model can be troublesome. Authors of different papers deal with this challenge in different ways.

In their milk-run tour problem, Güner et al. (2017) take stochastic travel times into account during re-optimization. The moment of re-optimization is at the end of one of a set of arcs which has to be traversed to drive from one customer to the next. In order to keep the computational demand of the model manageable Güner et al. (2017) implement a limited look ahead approach. This means that instead of observing the traffic situation on all residual arcs to get to the next customer, they only observe the traffic situation on the next couple of arcs. The number of future arcs to be observed depends on the arc length, value of real-time information with respect to the predetermined travel times and congestion state. The other arcs on the route are approximated by means of the a-priori approximated travel times. Limitation of the problem size is an alternative approach to keep a model's computation time manageable. For example, Grippa et al. (2019) limit their problem to 24 vehicles and a maximum of 16 depots from where goods are transported to customers.

Toriello et al. (2014) consider deterministic arc costs in the re-optimization steps. They assume that the outgoing arc costs can be approximated as known values. Approximate Linear Pogramming (ALP) is applied to reduce the size of the solution space. Grippa et al. (2019) preserve the reactivity of their model by means of assuming that the inter-arrival times of customer requests follow a Poisson distribution. The customer request location is assumed to be independently and uniformly distributed. Both a Poisson distribution and uniform distribution are log-concave distributions and therefore do not require sampling for evaluation (Gendreau et al., 2016) (see subsection 2.1.3).

In the literature on routing problems, heuristics have started to occupy a prominent place. For recent literature surveys on heuristics for routing problems see the work by Oyola et al. (2017) and Cao and Yang (2017). For routing problems in general, the computation time becomes troublesome when the problem size grows. Even more is this the case for stochastic models because the evaluation of probability density functions requires additional computational power. When the solution is re-optimized in real-time, minimization of the model's calculation time is even more important. For that reason, some researchers adopt special heuristics to limit the computation time and arrive at a good but mostly not optimal solution. For example, Li et al. (2019) recognise that the use of exact methods to solve their van-pool problem requires a long computation time. For that reason they look into the use of heuristics and meta-heuristics such as tabu search, variable neighbourhood search and hybrid approaches to solve their problem.

Schilde et al. (2014) realise that recalculating the shortest path between any pair of nodes in a network every time the traffic situation changes is demanding when the network is large. For this reason, they choose to fix the shortest path between any pair of nodes. However, the length of that path changes as a consequence of the traffic situation. This simplification might lead to unrealistic paths in cases of severe congestion but preserves the tractability of the problem.

2.3.5. Challenges & possibilities

When it comes to dynamic models Psaraftis et al. (2015) suggest to research differentiated weights assigned to future events. Near future events would have a larger weight than far future events. They argue that the further away an event is, the less influence it has on the current decision because of the uncertain events that will occur beforehand. This approach can be applied stochastic and dynamic models as well. It would mean that in the stage of re-optimization the possibility of, for example, arriving late at the next customer weighs heavier than a late arrival at the 5th next customer.

Gendreau et al. (2016) studied the literature published before 2016 related to stochastic routing models. They encounter that recourse actions in stochastic models often lack a degree of flexibility that is necessary in order to accurately model a real logistics system. Implicitly Gendreau et al. (2016) advocate for the development of stochastic models with sophisticated re-optimization strategies. Amongst others these re-optimization strategies could expand their focus from a single vehicle to a fleet of vehicles (such as often encountered in literature focusing on the DARP). Braekers et al. (2016) arrive at a similar conclusion; The recourse actions do not capture enough complexities to model a logistics operation realistically. They conclude that most stochastic an dynamic models consider only one type of uncertain and/or dynamic parameter. However, in reality a system comprises of many more uncertain and/or dynamic parameters.

Güner et al. (2017) suggest that, similar to static stochastic routing models, large improvements can be achieved by modelling the correlation between the stochastic parameters more accurately. However, they realise that this is challenging especially for stochastic re-optimization models. Assuming dependence between model parameters rapidly increases the effort required to evaluate the cumulative probability distribution. This is especially the case for problems which do not allow for pre-optimization, e.g. the a-priori generation of subsets from a large number of links. For example, pre-optimization might be effective when it is possible to define different sets of arcs for the different routes between specific pairs of customers. This drastically reduces the solution space in which the re-optimization model has to search for the optimal set of arcs between a pair of customers.

2.3.6. Applicability for e-grocers

As discussed in section 2.1 two stochastic elements are relevant for e-grocers: travel times and service times. However, for the calculation of the a-priori route other stochastic elements might be relevant than for re-optimization. An e-grocer has to finish calculation of the a-priori trips before the morning of the delivery. So at the time of that calculation, the traffic conditions are more uncertain than when the trip has actually started. For this reason, it might be interesting to consider stochastic travel times only for the a-priori route calculation and use deterministic real-time travel times for re-optimization. However, in case the total trip length is long, the uncertainty regarding travel times at the end of the trip is large during the first re-optimization moments. This would be an argument to use stochastic travel times for re-optimization as well. For e-grocers the uncertainty of service times does not decrease when the moment of service gets closer. Therefore, it is suggested to include stochastic service times in both the a-priori optimization and re-optimization.

A well-thought use of stochastic and dynamic elements in a routing model can release the burden of long computation times to a large extent. During re-optimization there is no time to wait for excessive computation times. However, for e-grocers the problem size at each re-optimization step is limited; Only the residual part of a single trip has to be re-optimized. This means that during re-optimization the problem only entails a maximum of 30 customers. A stochastic routing problem of this size can be solved using simulation techniques within a limited amount of time (Ehmke et al., 2015).

For the computation of the a-priori route the size of the problem is much larger. For a large problem, tractability becomes an issue when stochastic elements are taken into account. Especially when considering that the trip-customer assignments have to be computed within 45 minutes. A promising strategy to work around this tractability issue is to use deterministic travel and service times for this assignment problem. Once the assignment is completed, stochastic travel and service times can be used to compute the optimal sequence of customers in each trip. In this phase the size of the problem is much smaller and much more time is available for its calculation.

2.4 Conclusion

The use of stochastic and dynamic elements improves the performance of a routing model. However, the degree of improvement depends on problem parameters such as delivery window size, length of trips and road network characteristics. In many cases Solomon's instances are used to evaluate the performance of a routing model. However, multiple researchers test their model on both a specific test instance (case-study) and a benchmark instance from literature. Depending on the nature of the probability distribution of the stochastic elements involved, the computational power required for stochastic models poses challenges. In order to maintain the computational demand at a manageable level, in most cases the stochastic elements are assumed to be independent and the probability distributions are defined as log-concave functions. Regarding service times this is a realistic approach. However, regarding travel times the independence assumption is inaccurate in most cases. The computation time for a dynamic model is influenced not only by the use of stochastic elements but also by the moment of re-optimization. When the time required for re-optimization is known beforehand, the re-optimization can be scheduled before the expected point in time at which the solution is going to be used.

The conclusions from the literature study lead to the recommendation for a routing model specifically tailored to the needs of e-grocers. This recommendation answers the following sub-question: "What is a promising approach to apply stochastic and dynamic routing models for e-grocers?" Inspired by the premium customer approach as used by Chen et al. (2018), it is suggested to make a distinction between "regular" deliveries and "flexible" deliveries. Regular deliveries have to be delivered in the communicated 20-minute delivery time window. Flexible deliveries allow for larger delivery time windows. This means that the time window constraint remains, however, the size of a time window depends on the type of delivery.

The concept of flexible deliveries could work as following: When customers order, they opt for a regular or flexible delivery. Once the trip has departed, the presence of flexible deliveries allows for dynamic re-optimization of the customer sequence in that trip. This strategy has the potential to improve the on-time delivery performance and reduce the waiting time caused by early arrivals at customers. Moreover, this re-optimization strategy does not incur large additional operational costs because the number of vehicles required to complete the trips will not increase. For e-grocers changing the sequence of customers in a trip does not result in a large increase in total travel times because the travel time between any pair of customers served in the same trip is short. The concept of flexible deliveries is explained in more detail in section 3.1.

The VRP solved by e-grocers is large in size and the resulting customer-trip assignments have to be computed within 45 minutes. However, calculation of the planned arrival times at each customer is allowed to take up to five hours. Therefore, it is interesting to split the traditional VRP into two separate problems: the customer-trip assignment problem and the sequencing problem. Because of the short computation time available to solve the large assignment problem, a deterministic approach is suggested. In the sequencing problem it is feasible to use stochastic travel and service times in order to determine the optimal sequence of customers because the size of this problem is smaller and the available computation time is long. The a-priori sequence of customers is computed in such a way that online re-optimization of the customers' indices is effective at optimizing the on-time delivery performance. The travel times used for this a-priori model are modelled as inter-dependent, whereas the service times are modelled as independent. Both distributions can be deduced from historic data. When the trip has begun, re-optimization of the sequence of residual customers takes place shortly before the driver departs from a customer. Stochastic service times and deterministic real-time travel times are used for re-optimization. The structure of the routing model is summarized in figure 2.1.

Customers who opt for a regular delivery receive a notification of the planned 20-minute delivery time window on the morning of the delivery. Customers who opt for a flexible delivery accept a larger delivery time window when placing their order and do not receive a smaller delivery time window on the morning of the delivery.

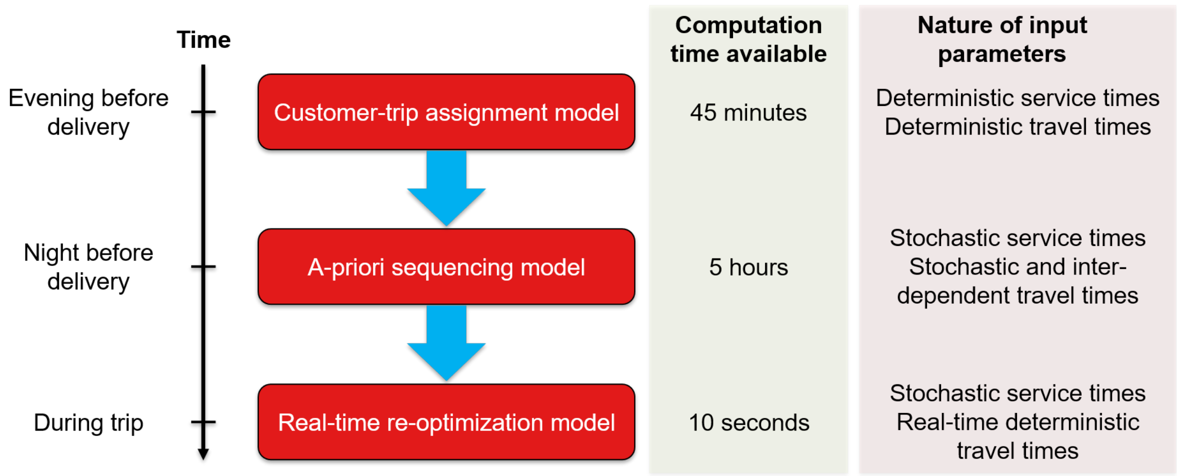

Figure 2.1: Structure of the routing model as proposed based on the literature review

The suggested routing model has to be tested in order to evaluate its performance. Testing can be performed on benchmark test instances or on case-study instances. The following chapters of this thesis present an approach to evaluate the performance the concept of flexible deliveries. Different routing model variants are compared, each implementing the concept of flexible deliveries in a different way. By means of a computational experiment conclusions can be drawn on the relative performances of the different routing model variants.

3
Methodology

The literature review as presented in chapter 2 yields a proposed solution approach to improve an e-grocer's routing model. The investigated routing model consists of three sub-models, see figure 3.1. The customer-trip assignment model is left out of the scope of this research. In this thesis a specific combination of an a-priori sequencing model and a re-optimization model (if any) is referred to as a "configuration".

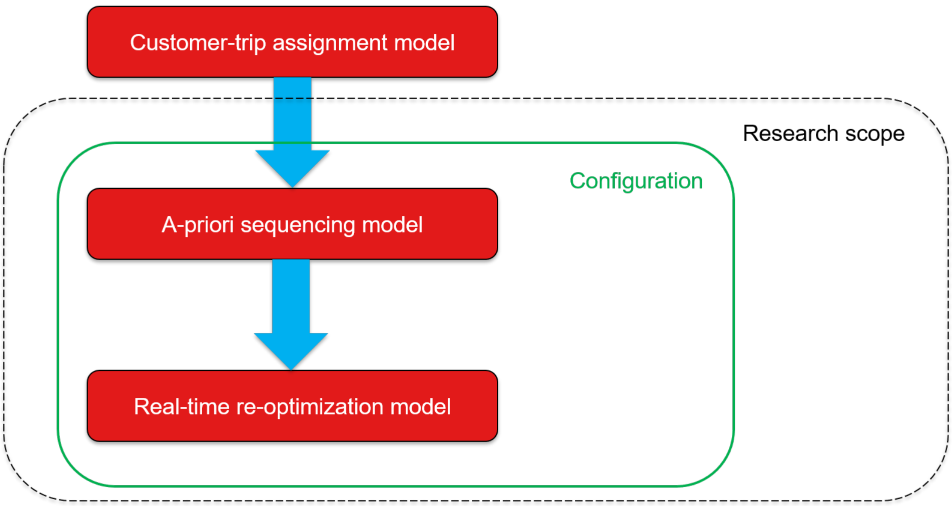

Figure 3.1: Components of the routing model

Given a set of customers who are served in the same trip, the optimal sequence of customers is calculated by the a-priori sequencing model for each trip separately. The real-time re-optimization model revisits the residual sequence of customers in real-time. This chapter starts with a more detailed description of the concept of flexible deliveries in section 3.1. Next, in sections 3.2, 3.3 and 3.4 the different a-priori sequencing models are discussed. Lastly the real-time re-optimization model is presented in section 3.5.

3.1 Flexible deliveries

Figure 3.2 illustrates how the time windows are designed around the scheduled arrival times ($T_{arrival}$). The majority of the time windows is determined as following: $[T_{arrival} - 10\,min, T_{arrival} + 10\,min]$. The customer index is defined as the delivery position within a trip; A customer with customer index 2 is the second customer served in a trip.

The flexible delivery approach suggests to change the strict time window constraint. The idea is to allow customers to opt for "flexibility" when they place their order. This means that they accept a larger time window compared to a "regular" delivery, see figure 3.2.

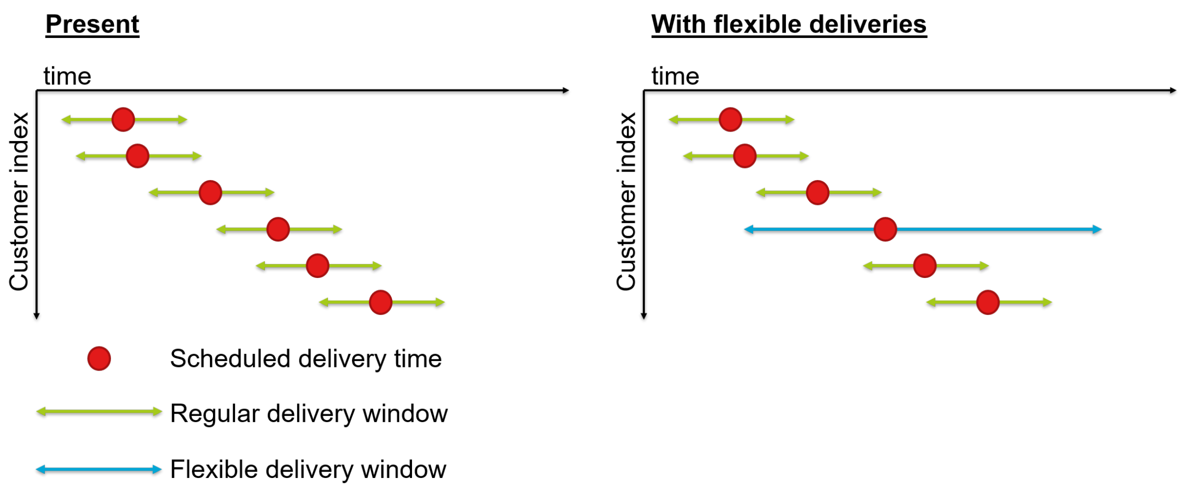

Figure 3.2: Illustration of the communicated delivery time windows in the present operations (left) and the suggested solution approach (right)

Making use of the flexible deliveries, the sequence of residual customers in the trip is reconsidered upon departure from a customer. This moment of reconsideration allows for taking the best estimate of the trip progress as an input to the re-optimization model. Either moving the flexible delivery backwards or forwards in the customer sequence can respectively mitigate the following problems: 1) A driver running late and delivering groceries after the customer's time window. 2) a driver running early and having to wait until the customer's time window starts.

The illustration in figure 9 shows how real-time re-optimization can compensate for running late. The left route is the shortest route as calculated by the a-priori routing algorithm. After serving customers A and B there is a severe delay because the sum of the realised travel times and service times at customers A and B was larger than estimated a-priori. Due to re-optimization the flexible customer is moved to the end of the trip, securing on-time arrival at customers D and E.

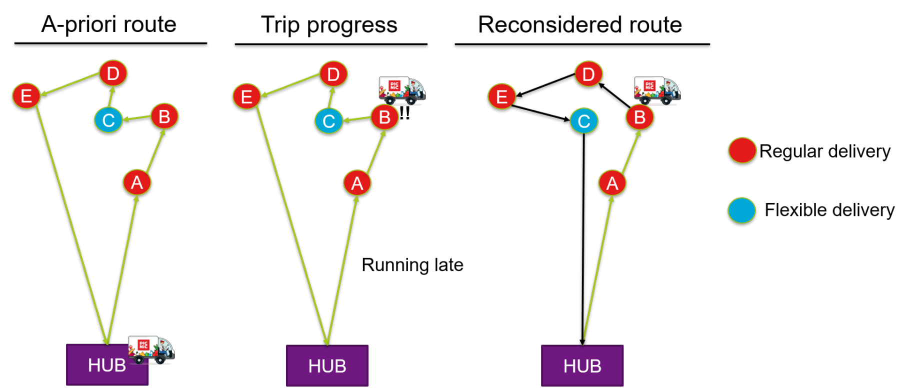

Figure 3.3: A flexible delivery can mitigate the effects of running late.

The illustration in figure 3.4 shows how real-time re-optimization can compensate for running early. Again the left route is the shortest route as calculated by the a-priori routing algorithm. After serving customer A the driver is running early. In order to prevent waiting time due to early arrival at other customers, customer C is served first.

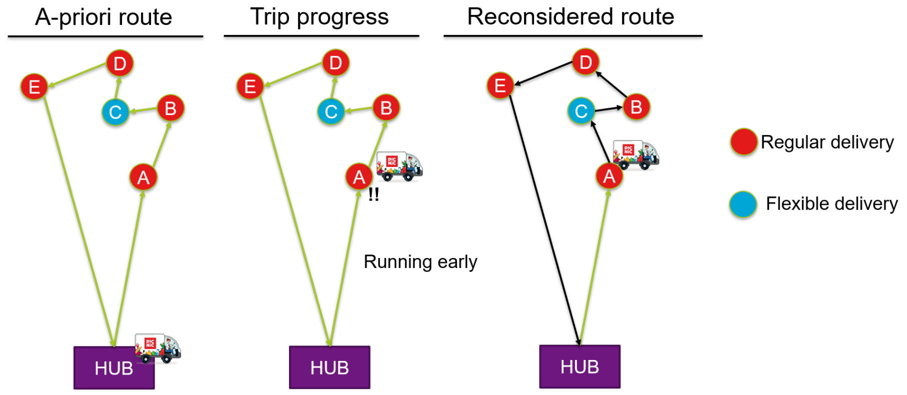

Figure 3.4: A flexible delivery can mitigate the effects of running early.

3.2 Benchmark sequencing model

This routing model is referred to as the benchmark sequencing model because of its simplicity and deterministic natur and uses the Gurobi MILP solver to find the optimal solution to the formulated problem.

3.2.1. Mathematical model

The notation of symbols in this mathematical model is adopted from Dumans et al. (1995).

Choice sets

N	Set of all nodes. The hub is represented as two separate nodes: the $start\_hub$ node and the end_hub node.
V	Subset of N containing only customer nodes.

Decision variables

$X_{i,j}$	$\forall i \in N, \forall j \in N$	Takes value 1 if the arc from node i to node j is included in the trip. Takes value 0 otherwise.
T_j	$\forall j \in N$	Arrival time at node j as a unix timestamp

The optimal sequence of customers is determined based on the minimization of the total trip duration. However, this objective does not define the optimal departure time from the hub. A trip can depart on several moments in time resulting in the same value for the objective function. In this the case, the departure time is determined based on the optimization of time-window overlap which results in a larger chance of on-time arrival.

SEA_j	$\forall j \in V$	Seconds early arrival with respect to the start of the customer order time window + 10 minutes.
SLA_j	$\forall j \in V$	Seconds late arrival with respect to the end of the customer order time window - 10 minutes.

Objective function

The dominant objective of this sequencing model is to minimize the duration of the trip. Subordinately it minimizes the number of deliveries with a planned arrival time within 10 minutes from the bounds of the customer order window. The latter results in a higher probability of arriving on-time at the customers.

$$Minimize \quad M * (T_{end\_hub} - T_{start\_hub}) + \sum_{j \in V} (SEA_j^2 + SLA_j^2) \tag{3.1}$$

Continuity constraints

All customers have one outgoing arc.

$$\sum_{j \in N} X_{i,j} = 1 \quad \forall i \in V \tag{3.2}$$

All customers have one incoming arc.

$$\sum_{i \in N} X_{i,j} = 1 \quad \forall j \in V \tag{3.3}$$

$start\_hub$ has no incoming arc and one outgoing arc.

$$\sum_{i \in N} X_{i,start\_hub} = 0 \tag{3.4}$$

$$\sum_{j \in N} X_{start\_hub,j} = 1 \tag{3.5}$$

end_hub has one incoming arc and no outgoing arc.

$$\sum_{i \in N} X_{i,end\_hub} = 1 \tag{3.6}$$

$$\sum_{j \in N} X_{end\_hub,j} = 0 \tag{3.7}$$

The same node can not be at both the start and end of the same arc.

$$X_{i,i} = 0 \quad \forall i \in N \tag{3.8}$$

Ensure time continuity during the trip. Where ST_i is the planned service time at customer i and $TT_{i,j}$ is the planned travel time from node i to node j.

$$T_j \geq \sum_{i \in N} X_{i,j} * (T_i + ST_i + TT_{i,j}) \quad \forall j \in V \tag{3.9}$$

Ensure time continuity at the $start\_hub$. T_{start_hub} represents both the imaginary arrival and departure time at $start\_hub$ because the service time at $start\_hub$ is set at zero seconds.

$$T_{start\_hub} \leq \sum_{\forall j \in V} X_{start\_hub,j} * (T_j - TT_{start\_hub,j}) \tag{3.10}$$

Ensure time continuity at the end_hub.

$$T_{end\_hub} \geq \sum_{\forall j \in V} X_{i,end\_hub} * (T_i + ST_i + TT_{i,end\_hub}) \tag{3.11}$$

Time-window constraints
Arrive at customers within their planning windows. $PTWS_j$ and $PTWE_j$ respectively represent the start and end times of the planning window of customer j.

$$PTWS_j \leq T_j \leq PTWE_j \quad \forall j \in V \tag{3.12}$$

Constraints 3.13 and 3.14 calculate the seconds early (SEA_j) and seconds late (SLA_j) for the arrival time at each customer with respect to the optimal arrival window. Here $OTWS_j$ and $OTWE_j$ are, respectively, the start and end times of customer j's order time window. CWW is the width of the communicated delivery time window and equals 20 minutes.

$$T_j - OTWS_j + SEA_j \geq 0.5 * CWW \quad \forall j \in V \tag{3.13}$$

$$OTWE_j - T_j + SLA_j \geq 0.5 * CWW \quad \forall j \in V \tag{3.14}$$

3.2.2. Model validation

The benchmark sequencing model follows the same constraints and objective functions as Picnic's sequencing model. However, it uses a different approach to solve the LP, namely an exact method instead of heuristics. For this reason it can be expected that the trip duration of Picnic planned trips is equal to or longer than trips planned by the benchmark sequencing model. This is verified by investigating trips where Picnic's sequencing model and the benchmark sequencing model yield different customer sequences. Trips completed in the period between 18/Feb/2020 and 15/March/2020 are used for this analysis. The results are presented in table 3.1. Note that the percentage of planned trips with a different customer sequence is much larger than the percentage of trips with a different duration. This can be explained by the fact that two customers living in the same building are ordered arbitrarily (because the travel time between them is zero seconds). Amongst the trips with a difference in trip duration, it was found that the benchmark sequencing model outperforms Picnic's sequencing model in terms of trip duration by an average of 26 seconds.

Table 3.1: Comparison between Picnic's sequencing model and the benchmark sequencing model

Percentage of trips with a different sequence	Percentage of trips with a different duration	Average deviation [a]	Maximum deviation
9.6%	4.2%	-26 seconds	-230 seconds

[a] Amongst trips for which the trip durations are different

An example of a planned trip along with the 1-hour order windows (green double-sided arrows) and the communicated 20-minute delivery time windows (green double-sided arrows) is presented in figure 3.5. Note that the communicated delivery time windows are mostly designed around the planned arrival time at a customer. Except for the cases where the planned arrival time is close to the edges of the one-hour order window.

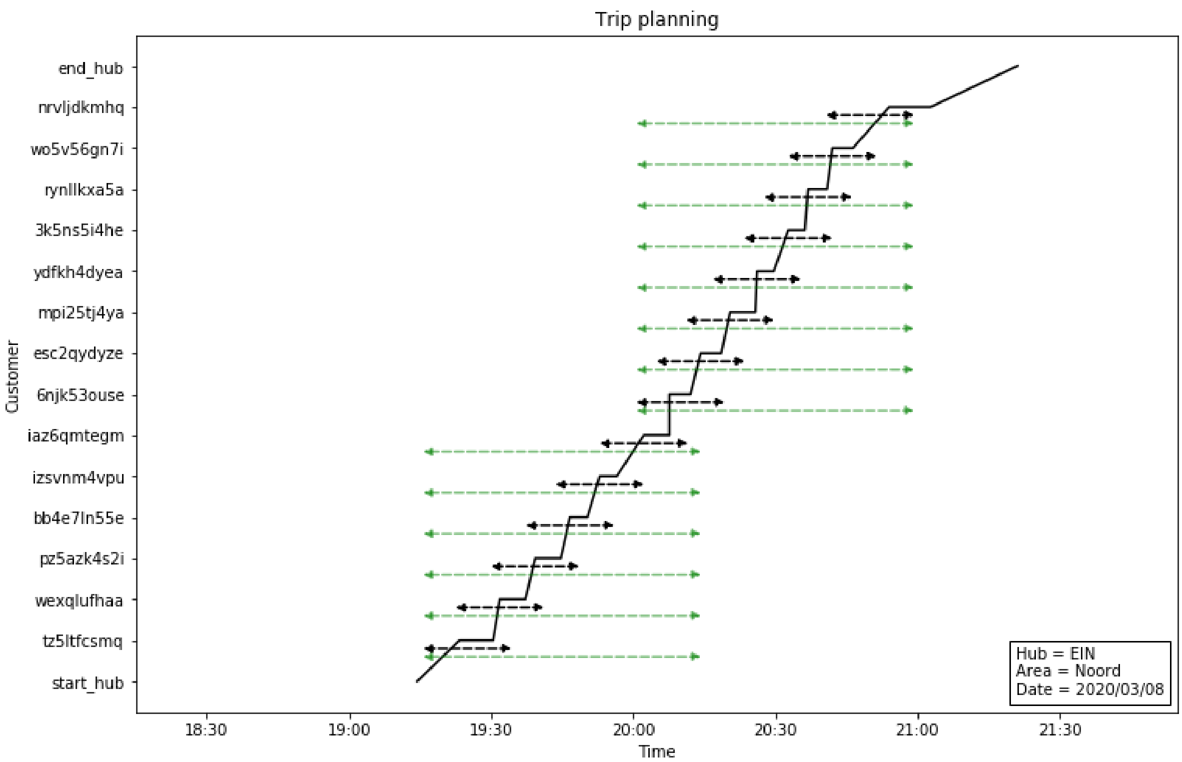

Figure 3.5: Example of a planned trip

3.3 Simulation-based sequencing model

The simulation-based sequencing model is a stochastic a-priori sequencing model. This model uses the benchmark sequencing model as its basis. It concerns a stochastic sequencing model because the optimal sequence of customers is selected based on a comparison of performance predictions of multiple solutions. A similar approach in the context of the Vehicle Routing Problem with Stochastic Demands (VRPSD) is taken by Juan et al. (2011). They search for the optimal number of trips to deliver uncertain amounts of goods to a set of customers. For different numbers of trips, they solve numerous iterations with customer demands generated according to Monte-Carlo simulation. The results for the different numbers of trips are compared, and this comparison yields an optimal number of trips. This sequencing model uses a similar approach to search for the optimal sequence of customers. An overview of the simulation-based sequencing model is provided in figure 3.6.

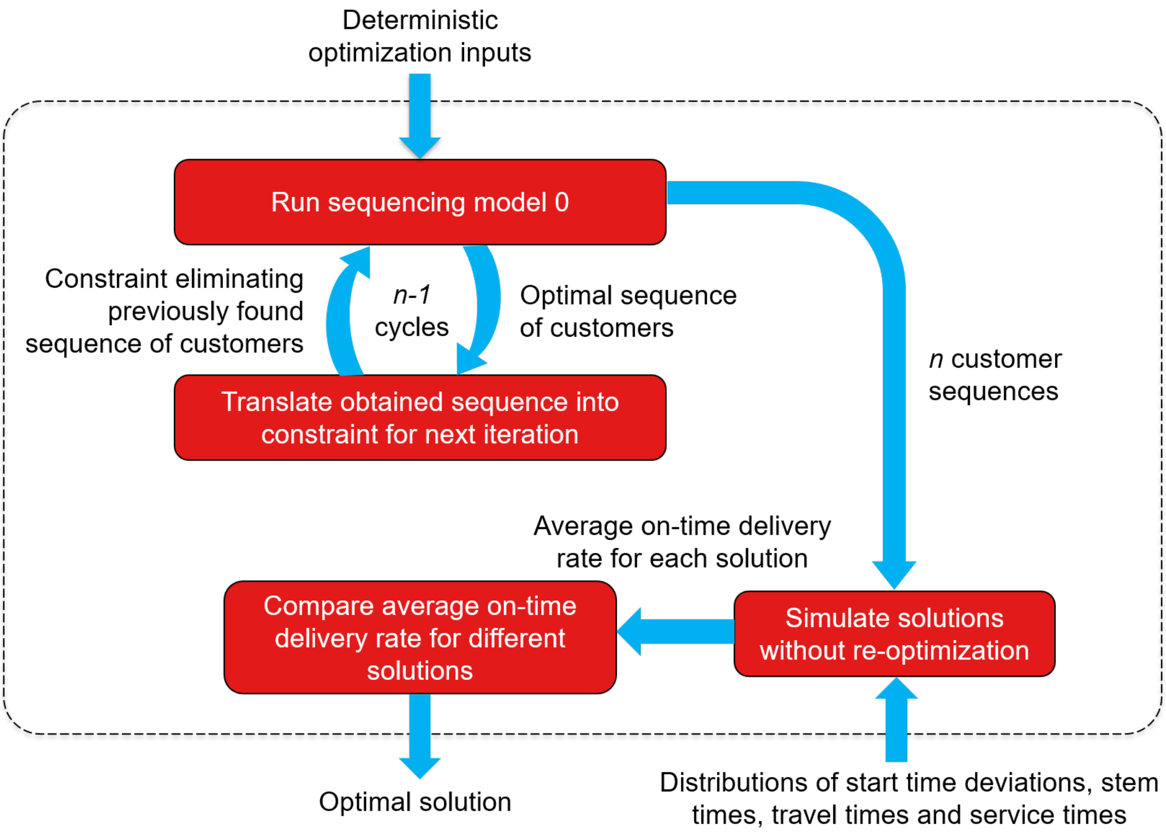

Figure 3.6: Overview of the structure of the simulation-based sequencing model

It is common practice to use the same distributions of stochastic parameters for both obtaining the optimal solution and evaluating it (Juan et al., 2011, Lei et al., 2012). In the context of this research this means that all n customer sequences are simulated using the simulation environment as presented in section 4.3. Based on the outcome of these simulations, the optimal solution is selected. The performance of this sequencing model is evaluated by means of the same simulation environment (see chapter 4).

3.3.1. Mathematical model

This sequencing model completes n runs of the benchmark sequencing model. In each next run, the solution(s) found in the previous run(s) are eliminated. This results in a different solution for each run. The corresponding mathematical model is nearly identical to the mathematical model defined for the benchmark sequencing model, see section 3.2.1. In order to calculate not only the optimal solution, but also a set of good solutions, the following choice sets are added to the model. A different solution is obtained for each iteration because of one additional constraint (constraint 3.15).

S		Set of all previously computed solutions
AS_s	$\forall s \in S$	Set of arcs used in solution s. The arc from node i to node j is denoted as (i,j).

In order to arrive at different solutions constraint 3.15 is added to the model. The *n* best solutions are looked for. In case less than *n* feasible solutions are available a smaller set of solutions is compared.

$$\sum_{(i,j)\in AS_s} X_{(i,j)} \leq len(N) - 1 \quad \forall s \in S \tag{3.15}$$

3.3.2. Determining the optimal solution

The *n* different solutions are evaluated by means of running an a-priori simulation. The simulation outputs for the *n* different solutions are compared based on the on-time delivery rate and the average time spent per delivery. In order to determine the optimal solution out of a set of good solutions, the on-time delivery rate should be decisive because the differences in average trip duration between the different solutions are small ($\leq\pm 0.4\%$).

Other researchers take different approaches to determining the optimal solution out of a set of solutions. For example, Groß et al. (2016) select an optimal solution out of a set of solutions based on the maximum regret of each solution. The regret of a solution relative to other solutions in the set is calculated for each iteration of the a-priori simulation. Once all iterations of the a-priori simulation are completed, the solution which has the lowest maximum regret is regarded as optimal. The approach taken by Groß et al. (2016) is more conservative compared to the approach taken in this sequencing model because it picks the best solution based on the worst case performance of each solution, instead of the average performance of each solution.

During the a-priori simulation, re-optimization is not considered. Although considering re-optimization during a-priori simulation could improve the performance of this sequencing model, the computational burden of simulating trips including re-optimizations is large and therefore not included in this research.

The number of iterations of this a-priori simulation is determined based on a comparison of the average values for the on-time delivery rate obtained through different numbers of iterations. An analysis has shown that the number of iterations has a significant effect on the average on-time delivery rate. The best decision regarding the optimal solution would be based on a large number of iterations. However, taking computation times into account, performing in the order of 1000 iterations per trip per solution is infeasible. Therefore, a compromise is made and only 25 iterations are performed. This will result in the selection of a sub-optimal solution in some occasions.

3.3.3. Number of solutions to compare

In order to determine the optimal number of customer sequences to compare (*n*), two aspects of the model have to be taken into account: the computation time which scales linear with the number of solutions compared, and the quality of the optimal sequence of customers returned by the model. The performance of the model is compared for three numbers of compared solutions (*n* = 1, *n* = 3 and *n* = 5). Figure 3.7 illustrates which solution was selected out of the set of good solutions for two different values of *n*.

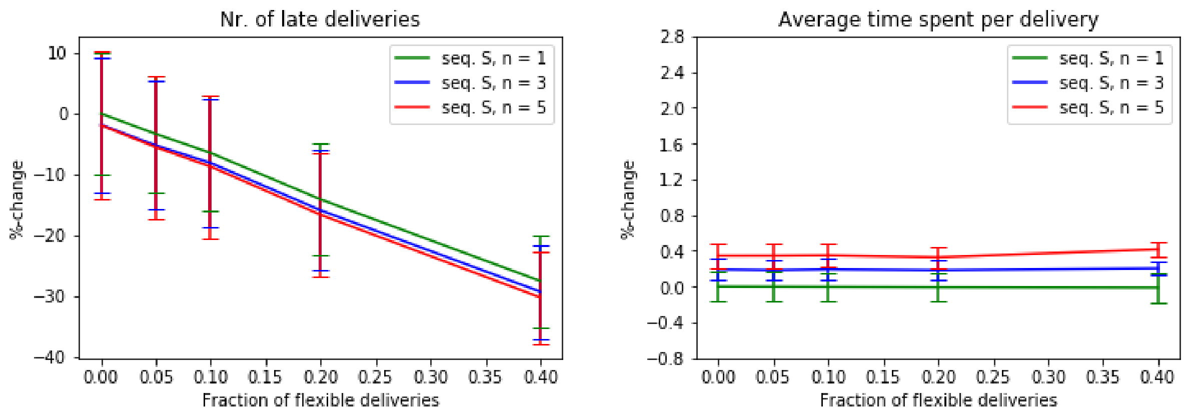

Figure 3.8: Comparison of the simulation-based sequencing model variants with different numbers of solutions considered. The results are presented as relative to the version with n = 1 and a fraction of flexible deliveries of 0.0.

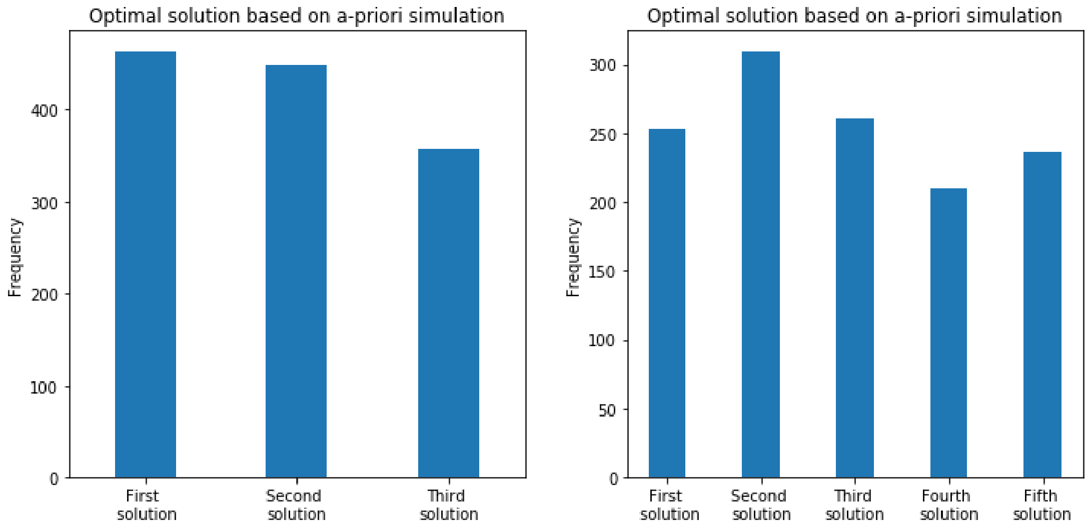

(a) $n = 3$

(b) $n = 5$

Figure 3.7: Optimal solution based on a-priori simulation for (a) $n = 3$ and (b) $n = 5$

The results show that the average simulated on-time delivery rate does not necessarily decrease when additional constraints are added to the mathematical model. The relative frequency of a certain solution being the optimal one depends on the number of solutions considered. The trend for $n = 3$ differs from the trend for $n = 5$. In order to gain an insight in the increased performance of the sequencing model when the number of considered solutions is increased, an experiment is performed with three different values of n, see figure 3.8. These results are obtained when running the experiment as described in chapter 4 without the re-optimization model.

From the results as presented in 3.8 it can be concluded that when no re-optimization is used, the number of solutions considered does not significantly affect the on-time delivery rate. However, the average time spent per delivery increases when the number of considered solutions increases. This can be explained by the fact that the total trip duration is the dominant factor in the objective function of this sequencing model. As a consequence, the total trip duration never decreases when a constraint is added to the mathematical problem.

It is decided that $n = 3$ yields the most promising sequencing model. Comparison of three good solutions still tests the concept of this sequencing model while limiting the computation time required to run the experiment.

3.3.4. Model validation

Figure 3.9 illustrates three distinct solutions for one specific trip. The objective function values are 85450996, 85451067 and 85481097 for solution 0, solution 1 and solution 2 respectively. Based on the values for the objective function one would expect solutions 0 and 1 to be close in terms of performance, while solution 2 would be significantly inferior. However, the simulated on-time delivery rate of solution 0 appears to be higher, and therefore solution 0 is selected as the optimal solution. This indicates the fact that the benchmark sequencing model is not optimal for achieving the highest on-time delivery rate.

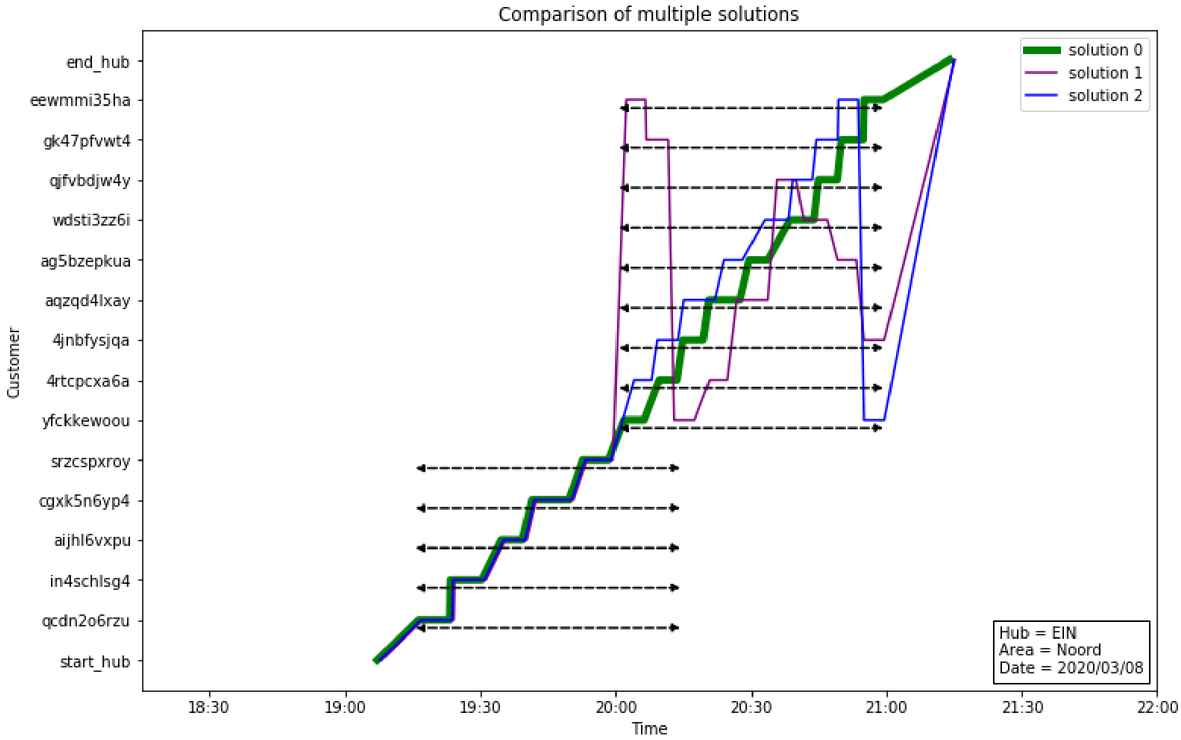

Figure 3.9: Visualisation of a set of good solutions. The double-sided arrows represent the one-hour planning window for each customer.

3.4 Heuristics-based sequencing model

In order to maximize the potential of real-time re-optimization it is interesting to investigate how the a-priori sequencing model can contribute to creating possibilities for successful re-optimization. The design approach for the heuristics-based sequencing model is presented in figure 3.10. In order to gain an insight in the correlation between the index of a flexible delivery in a trip and the effectiveness of re-optimization, an exploratory research is conducted.

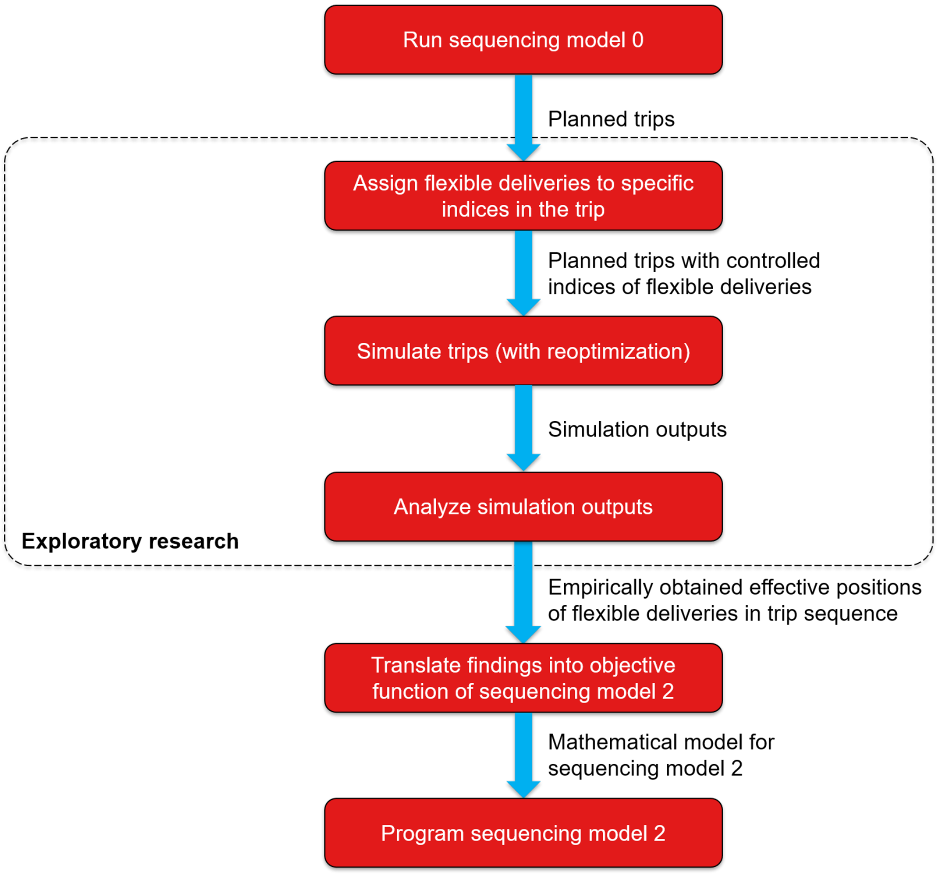

Figure 3.10: Design approach for the heuristics-based sequencing model

3.4.1. Exploratory research

This exploratory research isolates the index of a flexible delivery in a trip from other model parameters. It uses the simulation of one day of trips from a single hub to obtain an insight in the effect of certain rules regarding the indices of flexible deliveries in a trip. The simulation is executed using the simulation environment as presented in section 4.3. The re-optimization model as presented in section 3.5 is used. For each tested model variant 100 iterations are performed. Through this exploratory research a heuristic is developed that improves the effectiveness of re-optimization. This heuristic is not guaranteed to be the optimal way to use flexible deliveries. However, it should contribute to an enhanced effect of flexible deliveries. The different rules regarding the indices of flexible deliveries investigated in this exploratory research are graphically illustrated in figure 3.11.

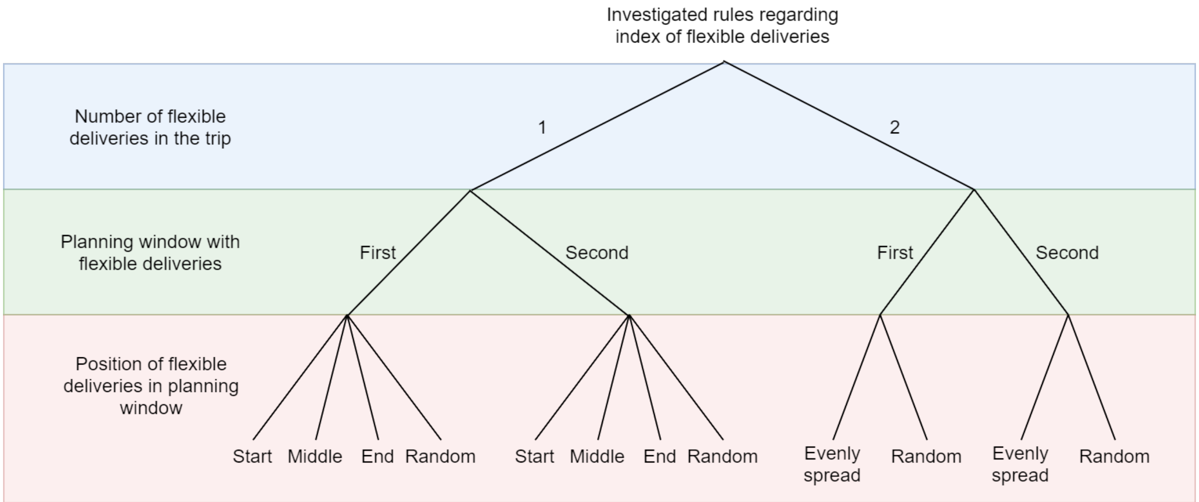

Figure 3.11: Model variants tested in the exploratory research

The results from the exploratory research suggest the following:

- In case there is only one flexible delivery in a planning window, the position of the flexible delivery within that window correlates with the performance on the KPIs.
- For the first planning window in a trip, the possibility for re-optimization is most effective when the flexible delivery is positioned at the middle of that window.
- Regarding the second planning window in a trip, the possibility for re-optimization is most effective when the flexible delivery is positioned at the start of that window.
- In case there are multiple flexible deliveries in one planning window, their positions within that window do not significantly affect the performance of the model variant on the KPIs.

The results from the exploratory research point out that the positions of flexible deliveries within the trip sequence significantly affect the performance of the routing model. Therefore, the findings from the exploratory research are translated into the mathematical model of the heuristics-based sequencing model in order to maximize the re-optimization possibilities offered by flexible deliveries.

3.4.2. Mathematical model

The majority of this mathematical model is identical to the mathematical model formulated for the benchmark sequencing model (see section 3.2.1). Therefore only the adapted objective function, added choice set, decision variables and constraints are presented here.

Choice sets

OD	Set of all nodes of which the index within the trip has to be optimized. This amounts to a set consisting of a maximum of two nodes; one customer node in the first planning window and one customer node in the second planning window. According to the findings of the exploratory research (see section 3.4.1) the indices of flexible deliveries do not impact the overall performance of the routing model when the number of flexible deliveries in the same planning window exceeds one.

Decision variables

I_j	$\forall j \in N$	Index of node j in the trip
OID_j	$\forall j \in OD$	Deviation from the optimal index for node j

Objective function

In the first place, the objective function minimizes the flexible deliveries' deviation from their optimal delivery index. Subordinately it minimizes the total trip duration. Lastly the delivery window dummy variables are minimized, resulting in planned delivery times closer to the center of the corresponding communicated delivery time windows.

$$Minimize \quad M * \sum_{j \in OD} OID_j + M/1000 * (T_{end\_hub} - T_{start\_hub}) + \sum_{j \in V} (DWS_j^2 + DWE_j^2) \tag{3.16}$$

Optimal delivery index constraints

The below constraints serve to calculate the absolute value of the difference (OID_j) between the optimal index of node j (OI_j) and the index node j is assigned to in the solution (I_j). The optimal index of each node is determined based on the heuristics that were obtained by means of the exploratory research (see section 3.4.1).

$$\sum_{j \in N} X_{i,j} * (I_j - I_i) = 1 \quad \forall i \in N \tag{3.17}$$

$$I_{start\_hub} = 0 \tag{3.18}$$

$$OID_j \leq OI_j - I_j \quad \forall j \in OD \tag{3.19}$$

$$OID_j \leq -OI_j + I_j \quad \forall j \in OD \tag{3.20}$$

3.4.3. Model validation

In order to validate that the heuristics-based sequencing model steers the customer sequence towards a solution that aligns with the exploratory research findings as presented in section 3.4.1, the position of flexible deliveries in a trip is compared for the benchmark and heuristics-based sequencing models. Figure 3.12 presents the position of a flexible delivery within the planning window for both planning windows. This plot only includes trips were there is exactly one flexible delivery in either the first, second or both planning windows. This is the case for 36% of all trips simulated in the experiment as described in chapter 4. Therefore, this sequencing model only affects a subset of all simulated trips. It can be concluded that for this subset, the heuristics-based sequencing model succeeds at positioning flexible deliveries in the middle of the first planning window and at the start of the second planning window.

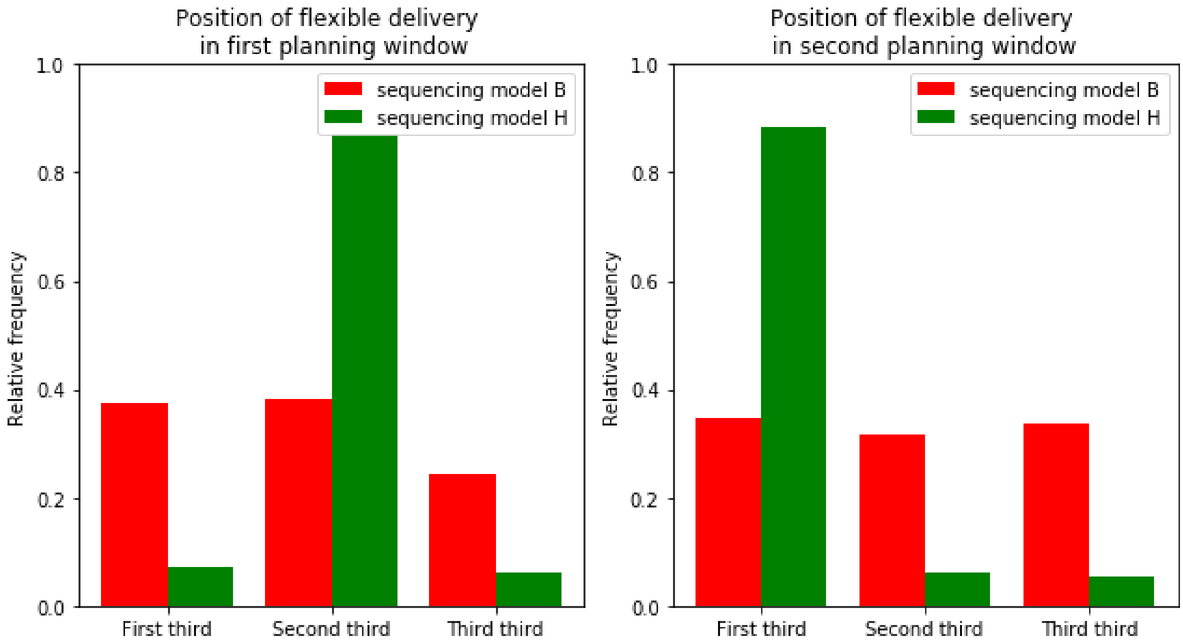

Figure 3.12: Position of flexible delivery in each planning window, based on trips with one flexible delivery in the first and/or second planning window

The additional term in the objective function of the heuristics-based sequencing model negatively influences the computation time of the model. In order to limit the time required to run the experiment with this sequencing model, the time limit for the a-priori optimization was set at 30 seconds per trip. When the time limit is reached before the optimal sequence of customers is found, the best solution at that point is taken. When no feasible solution is found within the time limit, the benchmark sequencing model is used to compute the optimal sequence of customers. The effects of this time limit are displayed in figure 3.13. Figure 3.13 demonstrates that an optimal solution is found within the time limit in 86% of the trips simulated in the experiment. For 10% of the trips a feasible solution is found, but this solution is not the optimal solution. In 4% of the cases no feasible solution was found within the time limit.

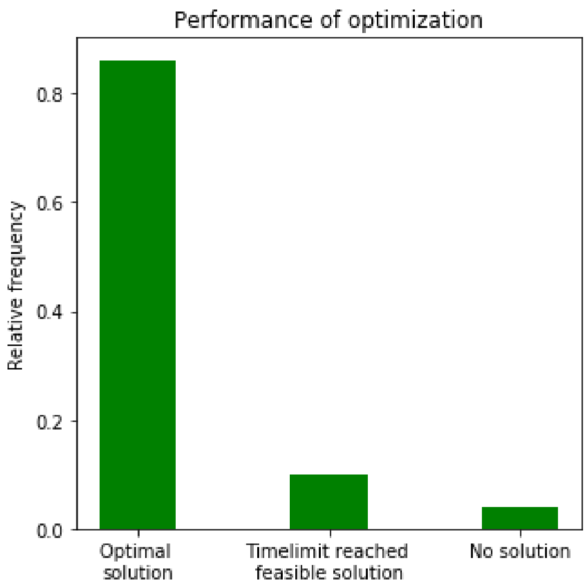

Figure 3.13: Performance of the gurobi solver

3.5 Re-optimization model

This re-optimization model is ran upon departure from each customer in a trip, except for the last two customers. Inputs to this model are the currently followed trip planning (planned stem times, travel times and service times), communicated delivery time windows and the trip progress. The output is the re-optimized customer sequence. Similar to Ulmer (2020), a route-oriented approach is used instead of a customer-by-customer approach. This means that the complete residual route is calculated during re-optimization instead of merely the next customer in the trip.

3.5.1. Mathematical model

The notation of symbols in this mathematical model is adopted from Dumans et al. (1995).

Choice sets

R	Set containing the current node, all residual customer nodes and the end_hub node
V	Set of residual customer nodes

Decision variables

$X_{i,j}$	$\forall i \in N, \forall j \in N$	Takes value 1 if the arc from node i to node j is included in the residual trip. Takes value 0 if not included in the residual trip.
T_j	$\forall i \in N$	Arrival time at node j as a unix timestamp

This re-optimization model has to take into account 20 minute delivery time windows (versus one-hour planning time windows for the sequencing models). When the progress of a trip falls behind the trip planning calculated a-priori, the use of hard time window constraints might result in an infeasible problem. For this reason, this re-optimization model makes use of soft time window constraints. Opposed to other researchers who penalize lateness or earliness in a linear fashion (Balakrishnan, 1993, Calvete et al., 2007, Taş et al., 2013), this model penalizes earliness and lateness in an exponential fashion. This approach ensures that a solution including two planned arrival times which are five minutes late is selected instead of a solution including one planned arrival time which is ten minutes late. In order to optimize the chance that a driver arrives on time, the planned arrival time should be close to the middle of the delivery time window. The optimal arrival window is defined as following: [delivery time window start + 10 min, delivery time window end - 10 min]. In order to include the deviation of the planned arrival time from the optimal arrival window in the objective function the below decision variables are introduced.

L_j	$\forall j \in V$	Seconds later than optimal at customer j
E_j	$\forall j \in V$	Seconds earlier than optimal at customer j
D_j	$\forall j \in V$	Absolute deviation of the planned arrival time from the optimal arrival window at customer j

Objective function

The objective function minimizes the cumulative penalty due to planned arrival times outside of the optimal arrival windows at customers. Secondly, the duration of the residual trip is minimized. dev_opt describes the weight of the penalty of arriving outside of the optimal arrival window and is set at 100. The value of the parameter dev_opt depends on the application-specific trade-off between the two elements in the objective function. In Picnic's case, on-time delivery is more important than total trip duration, therefore a weight of 100 is assigned to the planned arrival time penalties. For a different e-grocer this dev_opt could take another value. The correlation between the planned arrival time and the penalty in the objective function ($D_j * D_j$) is illustrated in figure 3.14.

$$Minimize \quad dev\_opt * \sum_{j \in V} D_j * D_j + T_{end\_hub} - T_{current\_node} \tag{3.21}$$

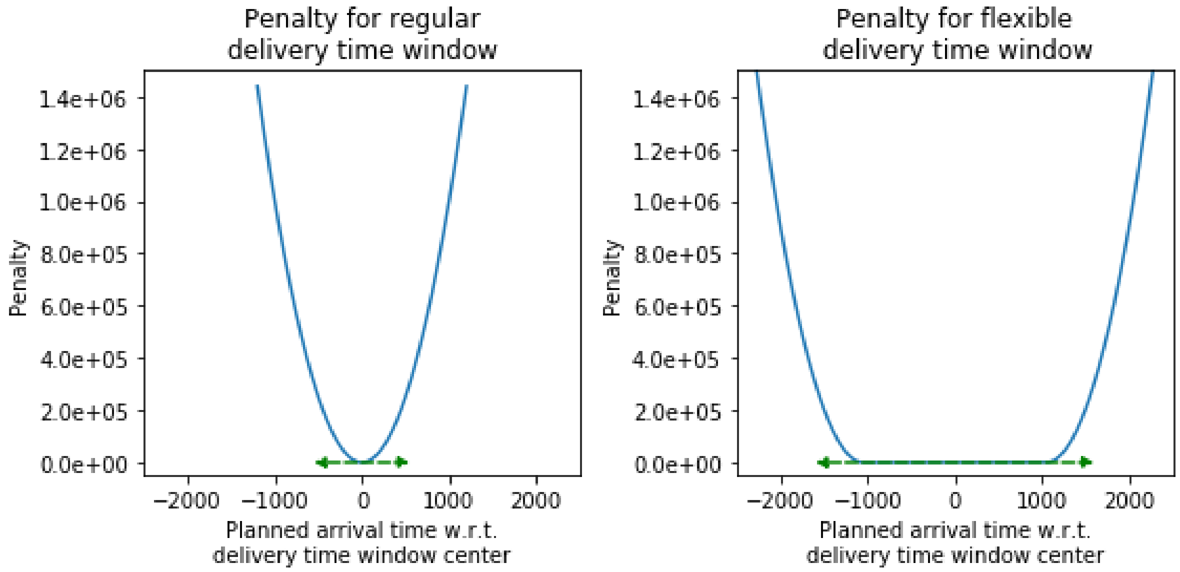

Figure 3.14: Correlation between the penalty and the arrival time

Continuity constraints
All residual customers have one outgoing arc.

$$\sum_{j \in R} X_{i,j} = 1 \quad \forall i \in V \tag{3.22}$$

All residual customers have one incoming arc.

$$\sum_{i \in R} X_{i,j} = 1 \quad \forall j \in V \tag{3.23}$$

The $current\_node$ has no incoming arc and one outgoing arc.

$$\sum_{i \in V} X_{i,current\_node} = 0 \tag{3.24}$$

$$\sum_{j \in V} X_{current\_node,j} = 1 \tag{3.25}$$

end_hub has one incoming arc and no outgoing arc.

$$\sum_{i \in V} X_{i,end\_hub} = 1 \tag{3.26}$$

$$\sum_{j \in V} X_{end\_hub,j} = 0 \tag{3.27}$$

A node can not be at both the start and end of the same arc.

$$X_{i,i} = 0 \quad \forall i \in R \tag{3.28}$$

Ensure time continuity during the trip. Where ST_i is the service time at customer i and $TT_{i,j}$ is the travel time from node i to node j.

$$T_j \geq \sum_{i \in R} X_{i,j} * (T_i + ST_i + TT_{i,j}) \quad \forall j \in V \tag{3.29}$$

Ensure time continuity at the current node. $T_{current\_node}$ represents the imaginary arrival time at the $current\_node$ because the residual service time at the $current\_node$ at the moment of re-optimization is equal to zero seconds. The $current\_time$ is passed from the simulation environment to the re-optimization model.

$$T_{current\_node} = current\_time \tag{3.30}$$

Ensure time continuity at the end_hub.

$$T_{end\_hub} \geq \sum_{\forall j \in V} X_{i,end\_hub} * (T_i + ST_i + TT_{i,end\_hub}) \tag{3.31}$$

Soft time-window constraints
The below constraints are used to determine the gap between the planned arrival time and the optimal arrival window. TWS_j and TWE_j respectively represent the start and end times of the delivery time window of customer j. SF represents the safety margin with respect to the boundaries of the delivery time window. The safety margin is set at 10 minutes.

Calculate seconds early with respect to the optimal arrival window.

$$E_j \geq (TWS_j + SF) - T_j \quad \forall j \in V \tag{3.32}$$

$$E_j \geq 0 \quad \forall j \in V \tag{3.33}$$

Calculate seconds late with respect to the optimal arrival window.

$$L_j \geq T_j - (TWE_j - SF) \quad \forall j \in V \tag{3.34}$$

$$L_j \geq 0 \quad \forall j \in V \tag{3.35}$$

Calculate the total deviation from the optimal arrival window.

$$D_j \geq E_j + L_j \quad \forall j \in V \tag{3.36}$$

3.5.2. Moment of re-optimization

As discussed in section 2.2.3 the frequency and moments of re-optimization have a large impact on the computational demand of the routing model. For this reason, the re-optimization is not performed upon departure from each node in the trip. The re-optimization is never performed upon departure from the $start\_hub$ because the trip departure time depends on the driver's estimate of the departure stem time, see chapter 4.3. In addition, re-optimization is never performed upon departure from the second-to-last and last customers in a trip because re-optimization at those nodes will never result in a different sequence of customers. In order to prevent situations where the driver has to wait until the re-optimization has completed, a time limit of 10 seconds is set on each re-optimization. As a consequence, the previously calculated sequence of customers is followed in case the re-optimization model does not find an optimal solution within the set time limit.

The re-optimization model uses the same travel times, service times and stem times at each moment of re-optimization. The only input parameter which can differ amongst different moments of re-optimization is the $current\_time$. If the $current\_time$ is the same as the planned departure time from that node according to a previously calculated planning, all input parameters to the re-optimization model are the same. In this case, re-optimization will not change the sequence of residual customers in a trip. For this reason, re-optimization is only started when the difference between the realised departure time and planned departure time from that node equals five minutes or more. A similar selective approach to the initiation of re-optimization is taken by Errico et al. (2016). They start re-optimization only when the trip runs behind schedule.

3.5.3. Model validation

An example of a trip where re-optimization has led to an improved sequence of residual customers is presented in figure 3.15. The green dot indicates the first re-optimization that resulted in a sequence of residual customers which deviates from the a-priori planned trip, the new trip planning is represented by the green line. At customer “e52s236oda” re-optimization resulted in yet another sequence of residual customers represented by the purple line. The sequence of residual customers was changed twice because the re-optimization model penalizes planned arrival times close to the start of the customers’ delivery time windows.

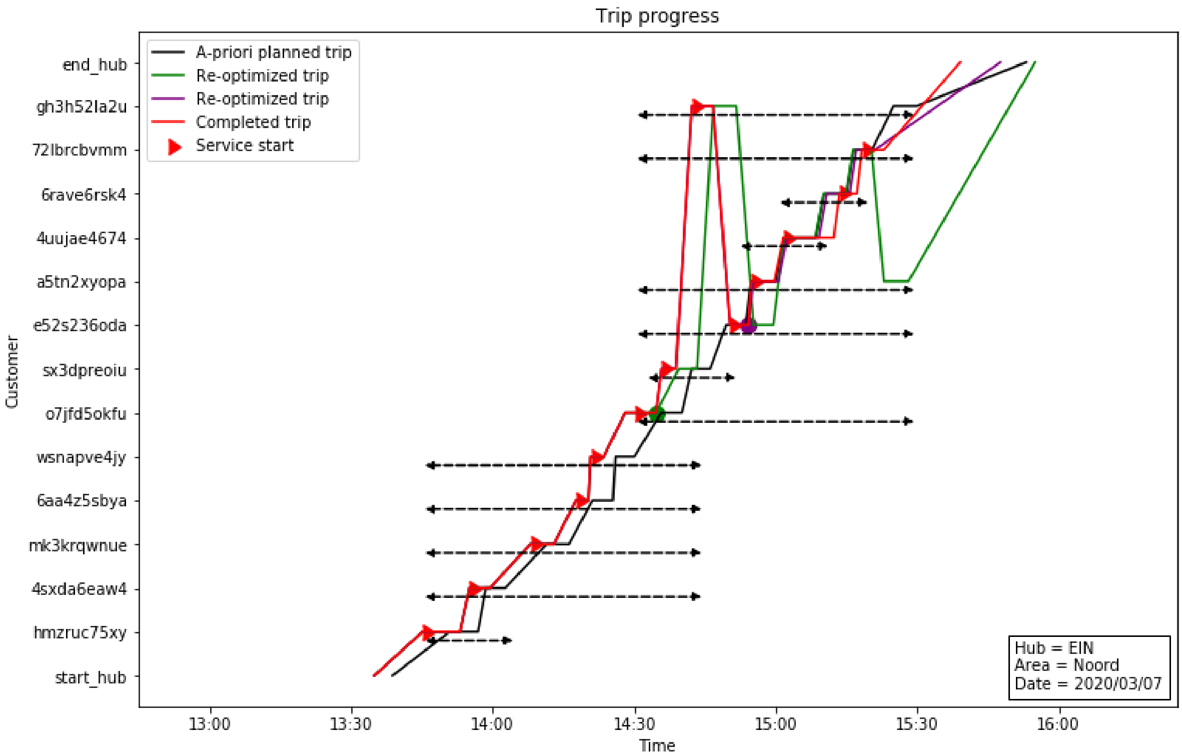

Figure 3.15: Example of a trip were re-optimization resulted in a changed sequence of customers. The double-sided arrows represent the delivery time windows for each customer.

When the re-optimization has not finished before the time limit, it is terminated and the trip planning remains unchanged. In order to gain an insight in the effects of this time-limit, the occurrence of failed re-optimizations is analysed. Figure 3.16 displays the correlation between the re-optimization failure frequency and the customer index. As the customer index in the trip increases, the number of re-optimization failures decreases. This can be explained by the fact that the size of the re-optimization problem decreases when the customer index increases. Moreover, re-optimization is initiated upon only 14% of the departures.

An experiment is performed to investigate the effects of the used approach to reduce the number of re-optimizations: Skipping re-optimization when the realised departure time lies within 5 minutes from the planned departure time. One day of operations from a single hub was simulated using two different re-optimization model variants. A variant where skipping is allowed and a variant where skipping is not allowed. The results are presented in figure 3.17. The results demonstrate that the on-time delivery rate and average time spent per delivery are not significantly affected by the selective skipping of re-optimizations. The fraction of trips were the a-priori sequence of customers was changed during re-optimization decreases due to skipping, i.e. reducing the number of re-optimizations that is performed. One can conclude that the optimal sequence of residual customers is sensitive for the deviation of the realised with respect to the planned departure time. However, those additional alterations of the customer sequence which were performed when skipping was not allowed, do not result in a significantly different on-time delivery rate or average time spent per delivery.

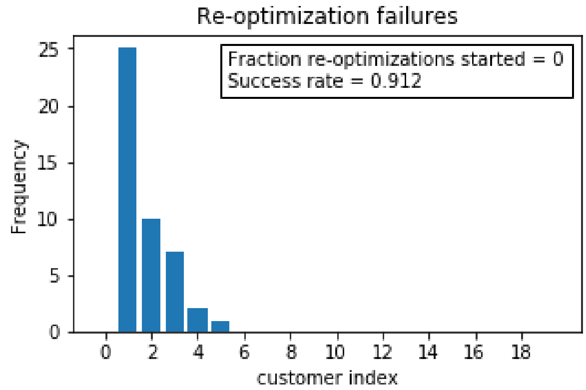

Figure 3.16: Re-optimization failure occurrences based on ten simulation iterations of all trips on one day of operation

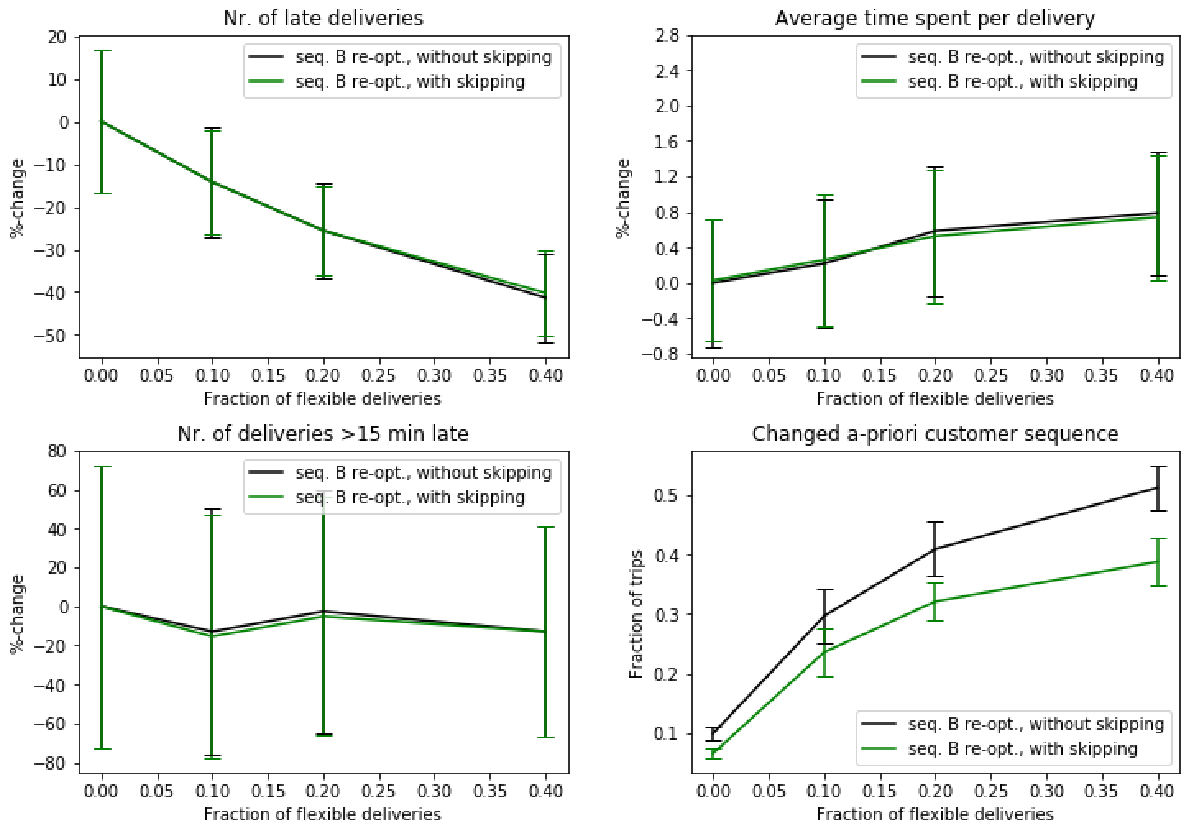

Figure 3.17: Comparison of experimental results for a model variant with skipping and a model variant without skipping. Ten iterations of the experiment are performed. Results are presented as relative to the variant without skipping and a fraction of flexible deliveries equal to 0.0.

4 Experimental method

In order to quantify the performance of the sub-models as presented in chapter 3, experiments are performed. This chapter describes the experimental method used to quantify the performance of different configurations.

An overview of the experimental method is provided in figure 4.1. Whether historic data is used in the a-priori sequencing model depends on the sequencing model that is tested. The simulation environment always uses historic data to sample realizations of trip departure times, stem times, travel times and service times. The completed simulated trips are analysed and their on-time delivery rate, rate of extreme lates and average time spent per delivery are used to quantify the performance of a configuration. All steps in the experimental method are executed on a personal computer with 16 Gb of RAM and a 1.90 GHz processor. Everything is programmed in Python 3.7. The a-priori sequencing model and the real-time re-optimization model call the Gurobi MILP solver to solve the mathematical problem and find the optimal sequence of customers.

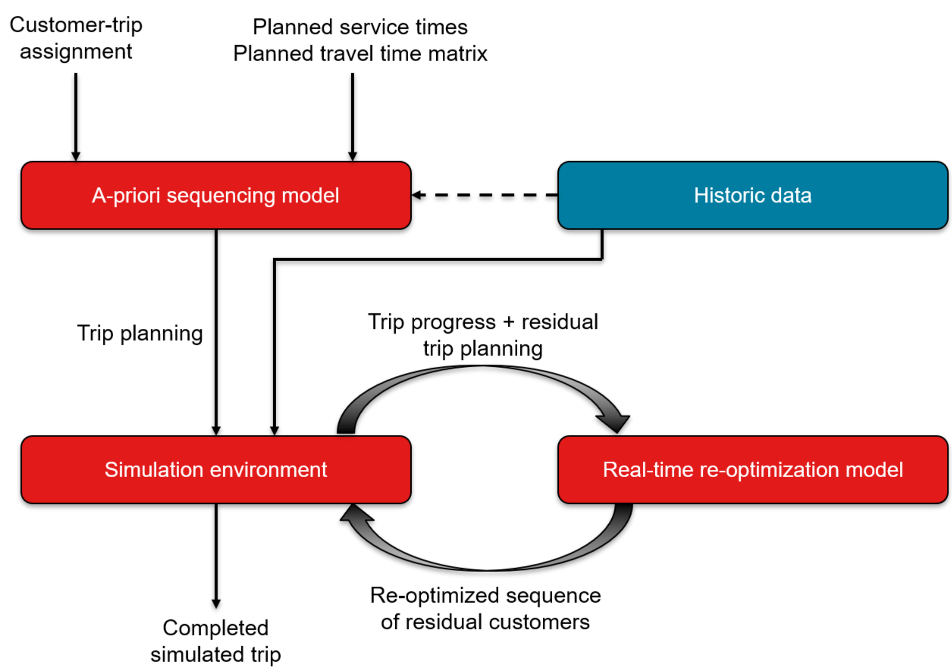

Figure 4.1: Overview of the experimental method

4.1 A-priori sequencing model

Inputs to the a-priori sequencing model are the customer-trip assignment, planned service times, planned travel time matrix and, depending on the model variant, historic data. These inputs are obtained from Picnic. The output is a trip planning consisting of the sequence of customers, planned departure and arrival times at the hub and at each customer and the communicated delivery time window for each customer. The investigated a-priori sequencing models are presented in sections 3.2, 3.3 and 3.4.

4.2 Historic data

The experimental setup makes use of historic data collected in the period between 17/Feb/2020 and 15/March/2020 (4 weeks). Both the planned and realised trip departure times, stem times, travel times and service times are collected from Picnics data warehouse.

4.2.1. Data filtering

First of all, the historic datasets are filtered in order to eliminate faulty measurements. Stem time data, travel time data and service time data are filtered based on the fraction of the realised over the planned time.

4.2.2. Data aggregation level

Most likely, each unique combination of area, weekday and shift will have its own distribution of trip departure time deviations, stem times, travel times and service times. However, for each of those unique combinations only a limited number of historic data points is available. Therefore, a statistical analysis is performed to determine the right data aggregation level for the distributions of historic data used in the simulation environment. The objective is to capture as many area-weekday-shift-specific peculiarities as possible because that makes the simulation environment more realistic. On the other hand, the used distributions should be statistically significant.

The statistical significance at different aggregation levels is studied by means of cutting a large dataset, containing data of all hubs over the period between 01/Dec/2019 and 15/March/2020, in smaller subsets each containing data of a period of 2 weeks. Histograms of stem times, travel times and service times are computed for each subset. Finally, the average bin size and corresponding 95% confidence interval for the mean is calculated. From this analysis it can be concluded the hub aggregation level is the most detailed level of aggregation which offers reliable data distributions. The hub level of aggregation is used to sample from in the simulation environment.

4.2.3. Data sampling

In the simulation environment, realisations of planned trip departure time deviations, stem times, travel times and service times are simulated by means of the corresponding filtered historic datasets. The sampling processes for the stem times, travel times and service times are identical and follow the following procedure:

1. The filtered realised times are divided by their planned counterparts to obtain factors.
2. These factors are stored in bins based on the value of their planned times. Factors corresponding to planned times that fall into the same bin are stored together. The number of bins depends on the type of data.
3. In the simulation environment planned times are used to compute simulated times; Using the planned time a multiplication factor is sampled from the bin within which the planned time falls.
4. This sampled factor is then multiplied by the planned time to obtain the simulated time.

An analysis of the trip departure time deviations shows that there is a historic correlation between the trip departure time and the difference between the planned and realised departure stem time. It can be concluded that, in general, when a delay occurs during the departure stem of a trip, the driver has departed earlier than planned. Therefore the trip departure time deviation is sampled based on the difference between the simulated departure stem time and the planned departure stem time.

4.3 Simulation environment

The simulation environment simulates the progress of a trip. For each tested customer-trip assignment, the a-priori sequencing model calculates the optimal sequence of customers once. The resulting trip planning is simulated ten times. Samples are drawn from historic datasets according to the procedure described in section 4.2.3. For the same iteration of a trip, the random samples that are drawn to calculate the simulated trip departure time deviation, stem times, travel times and service times are the same. In other words, if the sequence of customers in a trip is the same amongst different tested configurations, the simulated trip departure time deviation, stem times, travel times and service times are exactly the same. By means of fixing the random seed of each iteration, reproducibility of the experiment is guaranteed. The sequence of sampling the different data types in the simulation environment is graphically presented in figure 4.2.

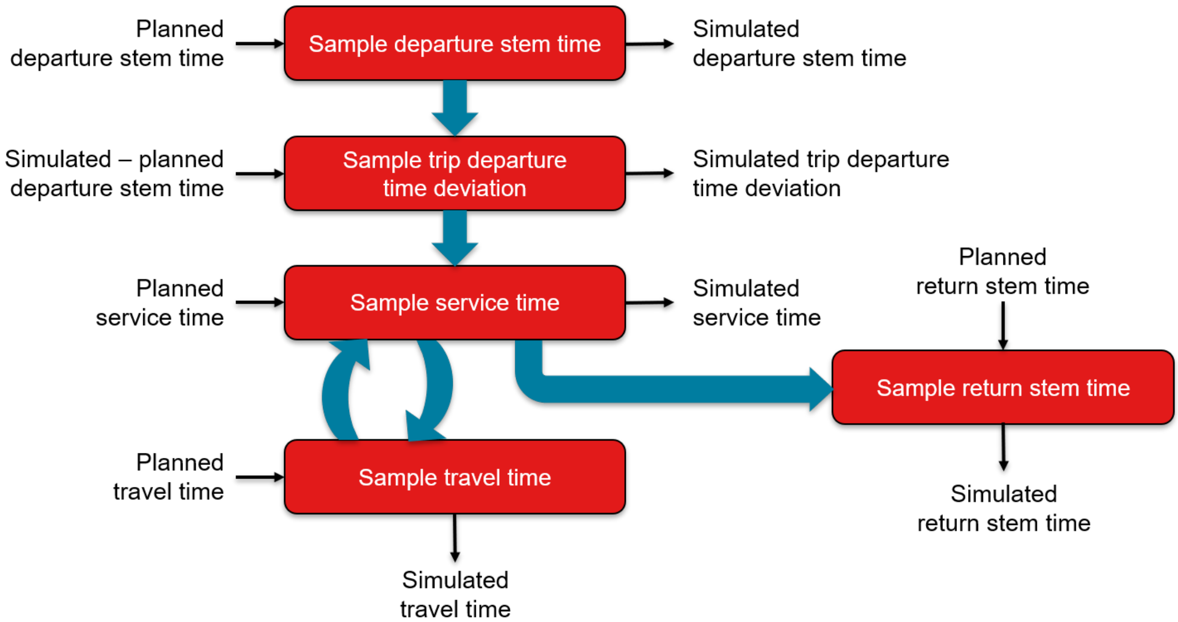

Figure 4.2: Sampling flow in the simulation environment; The number of customers in the simulated trip (n) determines the number of cycles (n-1), consisting of the sampling of a travel time and a service time, completed.

After the different trip times are simulated separately, they are used to compute the simulated trip. A completed simulated trip consists of simulated departure times and arrival times at each customer and the hub. A trip is simulated using algorithm 1. When the configuration to be tested includes a real-time re-optimization model, the trip-progress along with the residual trip-planning is passed to the real-time re-optimization model after each simulated delivery. The re-optimization model returns the re-optimized residual trip planning to the simulation environment.

Algorithm 1 Simulation algorithm

```
function execute re-optimization(current_time, residual trip planning)
    Calculate the optimal residual trip planning
    return residual trip planning
end function

                    If not stated otherwise, trip times are simulated trip times.

hub departure time = planned trip departure time + trip departure time deviation
first customer arrival time = hub departure time + departure stem time
for cust ∈ customers do
    if departure time previous node + simulated travel time < delivery time window start then
        arrival time = delivery time window start
    else
        arrival time = departure time previous node + travel time
    end if
    departure time = arrival time + service time
    if tested configuration includes re-optimization model then
        if current node != hub and number of residual customers > 1 then
            execute re-optimization(departure time, residual trip planning)
            residual trip planning = re-optimized residual trip planning
        end if
    end if
end for
hub arrival time = departure time last customer + return stem time
```

4.4 Validation of the simulation environment

The simulation environment is used to quantify the performance of different routing model configurations. In order to ensure that the results obtained by means of the presented experimental method are representative for the results that would be obtained in real operations, the simulation environment has to be validated. For this reason, the simulation environment is compared to the environment that Picnic encounters in reality.

4.4.1. Trips used for validation

In order to validate the simulation environment, the simulation outputs have to be compared to the trip performances as realised by Picnic. For this purpose the exact same trips as were planned by Picnic are simulated. Four weeks of operations from a single hub are used for validation (17/Feb/2020 - 15/March/2020). An example of a simulated trip is presented in figure 4.3.

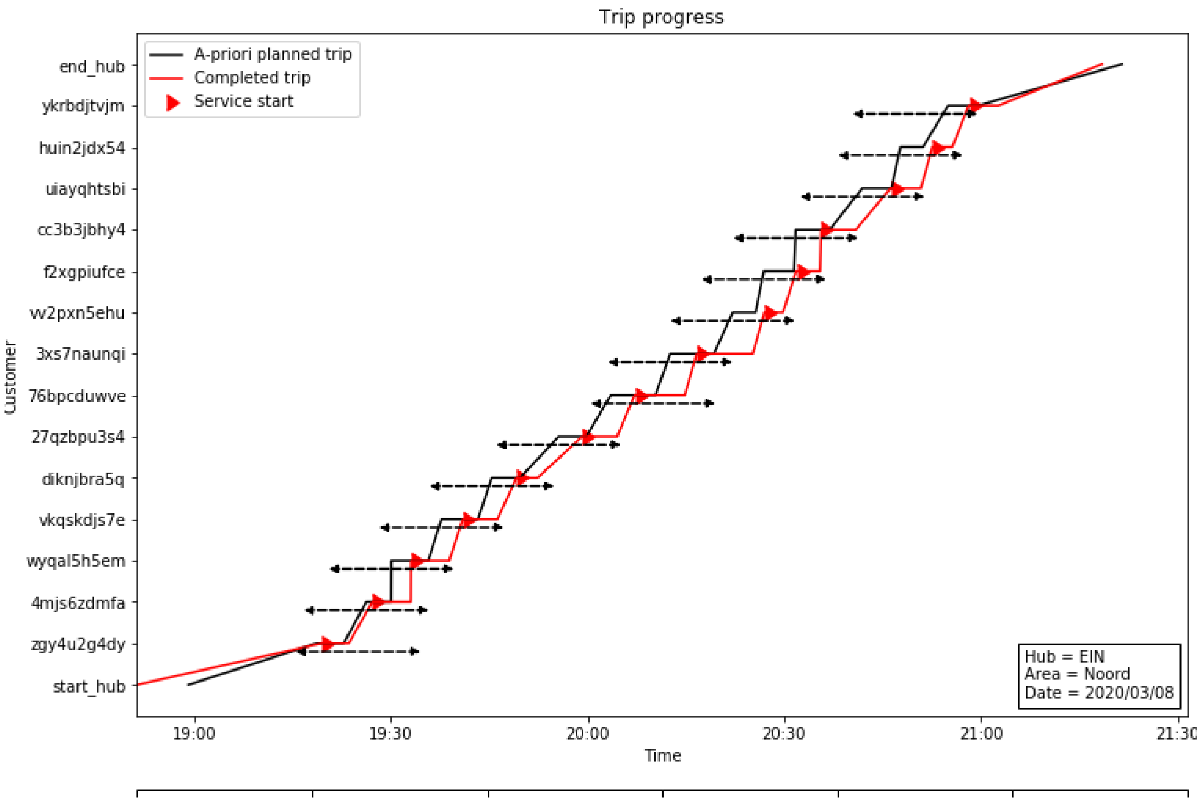

Figure 4.3: Example of a simulated trip used for validation; The double sided arrows represent the communicated delivery time windows.

4.4.2. Comparison on KPIs

The ultimate goal of this experimental method is to realistically predict the performance of different routing model variants on the KPIs as defined in chapter 1. Therefore, the on-time delivery rate and average time spent per delivery of simulated trips are compared to historic trips.

A meaningful comparison of the KPIs requires the selection of a subset of trips to be compared. Picnic's database keeps track of the on-time delivery of groceries. It would be inconsistent to filter out the historic service times measured for a portion of all deliveries but include those data points in the calculation of the historic on-time delivery rate. For this reason it is decided to only include deliveries with a reliable service time measurement in the calculation of the historic on-time delivery rate.

The average time spent per delivery can not be collected from Picnic's historic trip database directly, and therefore has to be calculated based on historic trip times. In order to calculate the historic average time spent per delivery, a subset of all historic trips was selected because trips which contain too many measurement errors do not provide a representative average time per delivery. Only trips which satisfy the below conditions are used for validation of the average time spent per delivery.

- No departure stem time measurement error
- No return stem time measurement error
- Less than two travel time measurement errors
- Less than two service time measurement errors

This filtering yields a subset of trips. The comparison of the historic and simulated average time spent per delivery is presented in figure 4.5. Ten iterations of the simulations are performed, yielding the presented 95% confidence intervals. Since historic trips were only completed once, no confidence intervals for historic KPIs could be calculated.

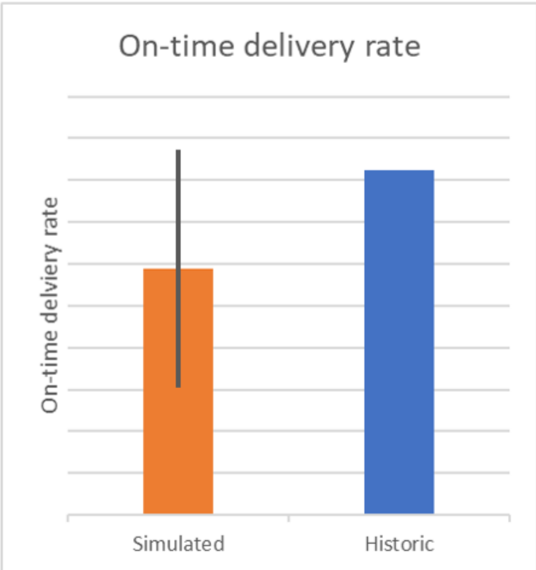

Figure 4.4: Comparison of the on-time delivery rate. The brackets represent the 95% confidence interval of the mean.

Figure 4.5: Comparison of the average time spent per delivery. The brackets represent the 95% confidence interval of the mean.

4.4.3. Analysis of validation results

The validation results demonstrate that the distributions of simulated times are similar to those of historic data. This was expected because samples are drawn from historic datasets. The difference between the historic and simulated trip times can be explained by the sampling process. In the simulation environment for each planned time a multiplication factor is drawn from a bin containing many historic multiplication factors. Since sampling is random, not all historic multiplication factors are used in the simulation environment the same number of times. Some are not used at all, others are used multiple times for different planned times (corresponding to the same bin). This also explains that the difference between historic and simulated averages is most evident for bins containing few data points.

The historic on-time delivery rate lies just within the 95% confidence interval of the simulated trips. This means that the simulation environment is likely to generate an outcome with an on-time delivery rate lower than the historic on-time delivery rate. There are two main explanations for the gap between the simulated and historic on-time delivery rate. First of all, the simulation environment does not consider a correlation between the working speed of a driver and the driver's progress with respect to the trip planning. In real operations, a driver's working speed might improve when running behind schedule.

Figure 4.6: Comparison of the size of the two vehicly types experimented with. A Toyota Prius is displayed for the sake of comparison.

Secondly, amongst other performance indicators, drivers are rated based on the on-time delivery rate they accomplish. At the same time, drivers confirm the start of a delivery through the app on their PDA, this might be cause for incorrect data measurements.

4.5 Test instances

This experiment simulates trips which where completed between 02/March/2020 and 09/March/2020. In order to limit the size of the experiment, one hub is selected. This choice is based on the hub-specific stem time distributions, travel time distribution, service time distribution and on-time delivery rate distribution. One hub with a relatively poor on-time delivery rate was selected.

Experiments are performed using two different types of customer-trip assignments: “Small vehicle” customer-trip assignments and “large vehicle” customer-trip assignments. The two vehicle types are compared in figure 4.6. The small vehicle is specialized at urban areas. The large vehicle is more suitable for suburban areas. For both vehicle types the set of customers in a specific shift-area combination is identical. The planned travel times and stem times for the same trip legs are 10% longer for the small vehicle. For each customer, the planned service time is the same for both vehicles types.

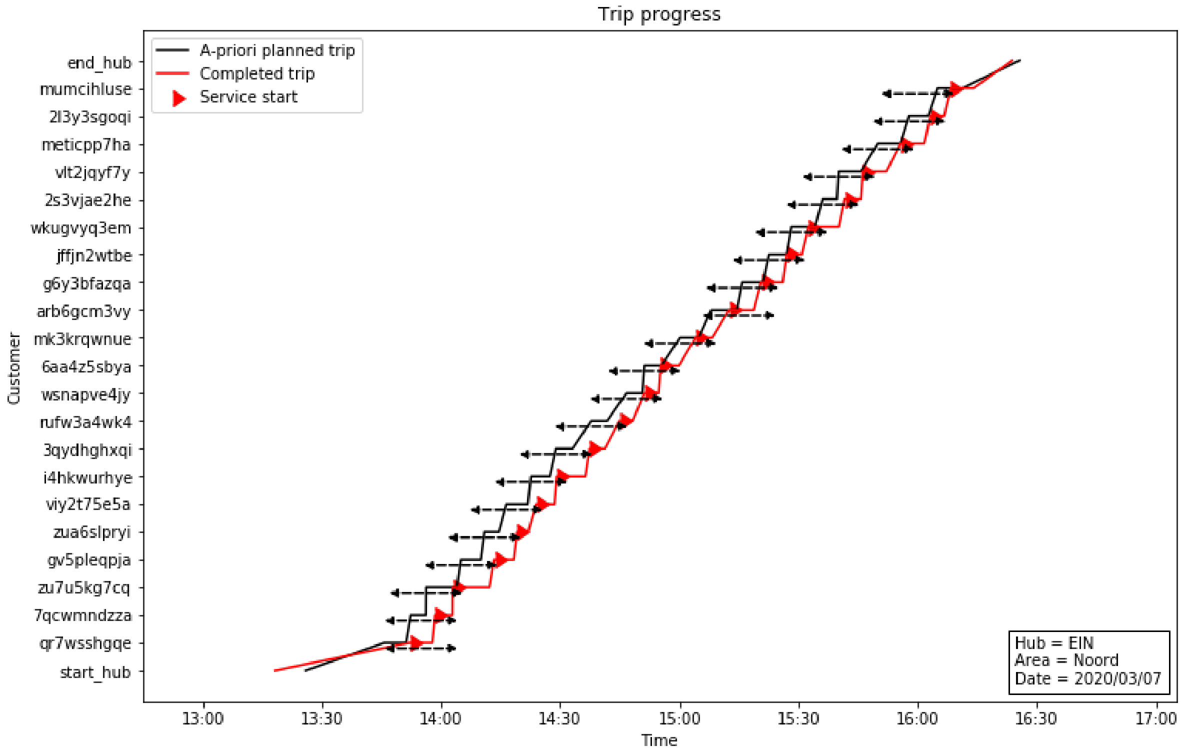

Figure 4.7: Example of a simulated trip for a large vehicle

Moreover, the performance of the concept of flexible deliveries depends on the size of the flexible delivery windows. Therefore, two different types of flexible delivery windows are experimented with.

- *Default*: The flexible delivery window is the same as the order window.
- *Extended*: The flexible delivery window is an extension of the order window. The start-time is the same, yet the flexible delivery window is 15 minutes longer.

Lastly, an important parameter of the test instances is the fraction of customers that opts for a flexible delivery. The correlation between this fraction of flexible deliveries and the performance of a configuration should be quantified. If a configuration only performs well on the KPIs when the fraction of flexible deliveries is very large, it might not be suitable for use in real operations. Offering customers a larger delivery time window comes at the price of either a compensation for customers who opt for a flexible delivery (which translates into operational costs for the e-grocer) or a worse customer experience. The following fractions of flexible deliveries are investigated: 0.0, 0.05, 0.1, 0.2 and 0.4.

4.6 Tested configurations

In order to answer sub-question 4, “How does the proposed routing model perform?”, different ways to apply the concept of flexible deliveries have to be studied. This section presents the specific configurations that are evaluated in the experiment and explains their relevance for answering sub-question 4. An overview of the tested configurations is provided in table 4.1.

Benchmark sequencing model without re-optimization

This configuration is the static and deterministic benchmark configuration to which other configurations can be compared. It consists of the benchmark sequencing model and does not include the real-time re-optimization model. This comparison yields an insight in the performance of different variants of stochastic and dynamic routing models with respect to a static deterministic variant. This configuration also resembles the routing model currently used at Picnic.

Simulation-based sequencing model without re-optimization

The simulation-based sequencing model is a stochastic sequencing model and predicts the optimal solution out of a set of good solutions by means of an a-priori simulation. Because the a-priori simulation takes into account the presence of flexible deliveries, this sequencing model indirectly considers the flexible deliveries when deciding amongst the set of good solutions. This configuration gives an insight in the performance of a stochastic and static routing model variant.

Simulation-based sequencing model with re-optimization

In order to quantify the performance of a stochastic and dynamic routing model variant, a configuration consisting of the simulation-based sequencing model and the re-optimization model is tested.

Heuristics-based sequencing model with re-optimization

In order to compute the optimal a-priori sequence of customers the heuristics-based sequencing model uses empirical correlations between the position of a flexible delivery in a trip and the effectiveness of re-optimization. This configuration also includes the real-time re-optimization model. The results of the static version of this configuration (without re-optimization) are not presented because the heuristics-based sequencing model is solely aimed at improving the effectiveness of re-optimization.

Table 4.1: Overview of the tested configurations

Configurations		**Flexible delivery windows**		**Vehicle size**	
Sequencing model	Re-optimization model	Default	Extended	Small	Large
Benchmark	No	✓	✓	✓	✓
Simulation-based	No	✓	✓	✓	✓
Simulation-based	Yes	✓	✓	✓	✓
Heuristics-based	Yes	✓	✓	✓	✓

5

Experimental results

This chapter presents and discusses the results of the experiments performed according to the experimental method as explained in section 4. First, the combinations of vehicle type and flexible delivery window are presented one by one in order to create an insight in how the performances of the tested routing model configurations compare. In section 5.5 the effects of the size of the flexible delivery windows are discussed. Next, section 5.6 explains how the number of deliveries in a trip affects the effectiveness of the studied configurations. In section 5.7 all of the presented results are integrated to formulate an answer to the sub-question: “How does the proposed routing model perform?”. For each configuration the results are presented as percentage deviations relative to the results of the benchmark configuration with a fraction of flexible deliveries of 0.0. In the figures presented in this chapter the brackets represent the 95% confidence intervals of the mean. In the figures the sequencing model variants are abbreviated as following: benchmark (B), simulation-based (S), heuristics-based (H).

5.1 Small vehicle and default flexible delivery windows

Figure 5.1 illustrates that the simulation-based sequencing model does not improve the on-time delivery rate of the last-mile distribution system significantly with respect to the benchmark configuration. In combination with the re-optimization model, it improves the on-time delivery rate marginally for fractions of flexible deliveries larger than 0.10. The performance of the heuristics-based sequencing model improves once the fraction of flexible deliveries is 0.10 or more. Regarding the fraction of extreme lates, no significant difference in performance is shown amongst the different configurations. The average time spent per delivery increases due to re-optimization of trips. This can be explained by the fact that if the objective of the a-priori sequencing model is the minimization of trip duration, re-optimization always results in an increase in the planned trip duration. Therefore, an increase in the simulated average trip duration is observed when the number of simulated trips is large. When looking at the configurations including the re-optimization model, it becomes clear that the heuristics-based sequencing model results in a larger increase in the average time spent per delivery than the simulation-based sequencing model. This can be explained by the fact that the dominant objective of the mathematical model of the heuristics-based sequencing model is to optimize the position of flexible deliveries in the trip, as opposed to the minimization of total trip duration for the simulation-based sequencing model.

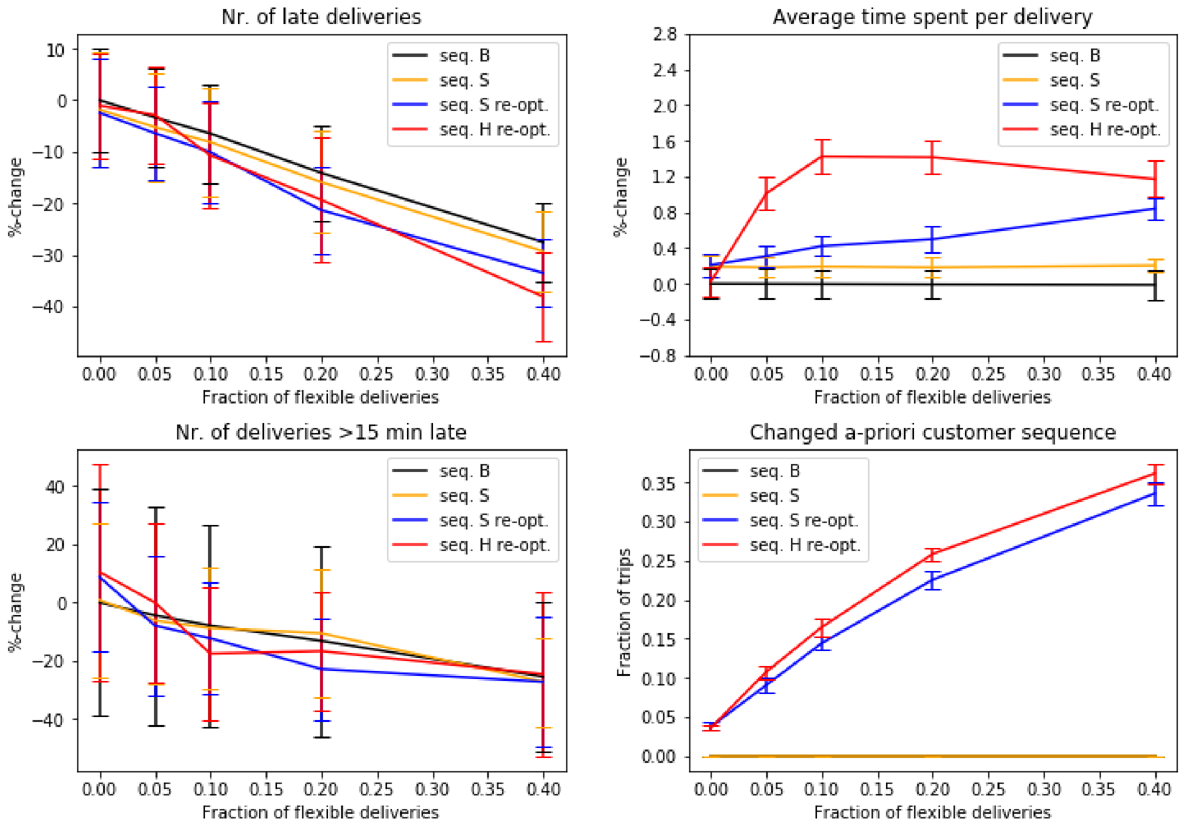

Figure 5.1: Results for the small vehicle type with default flexible delivery windows

5.2 Small vehicle and extended flexible delivery windows

When the flexible delivery windows are extended by 15 minutes, the re-optimization model has a larger impact on the performance of the last-mile distribution system. Figure 5.2 shows a significant reduction in the number of late deliveries for the configurations with the re-optimization model. In addition, the number of extreme lates is reduced. This improved on-time delivery performance comes at the price of a larger average time spent per delivery. The difference in terms of average time spent per delivery between the two configurations with re-optimization model is not significantly affected by extension of the flexible delivery windows. The heuristics-based sequencing model results in an at most 1% larger increase with respect to the benchmark configuration than the simulation-based sequencing model.

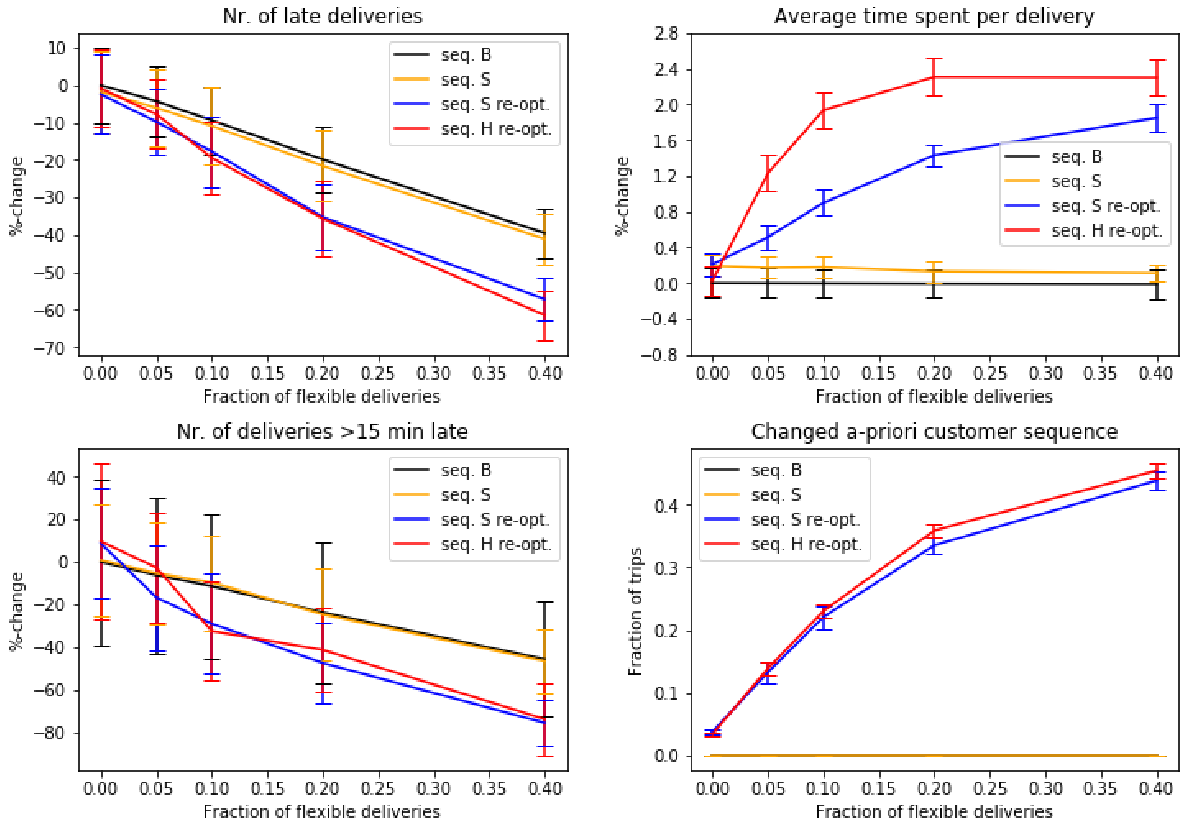

Figure 5.2: Results for the small vehicle type with extended flexible delivery windows

5.3 Large vehicle and default flexible delivery windows

For the large vehicle type, the differences between the performances of the investigated configurations are more profound. For all fractions of flexible deliveries the on-time delivery rate is significantly improved by the simulation-based sequencing model. In case a re-optimization model is used, the simulation-based sequencing model reduces the number of late deliveries by 13% with respect to the benchmark configuration. When looking at extreme lates, the simulation-based sequencing model needs the re-optimization model to achieve a performance similar to the benchmark configuration. The reason being that the simulation-based sequencing model selects the best solution out of a set of good solutions based on the highest on-time delivery rate. This approach is reflected in the experimental results regarding the on-time delivery rate. However, the experimental results show that this approach results in customer sequences which are at a higher risk of extreme lates. This can be explained by the fact that the simulation-based sequencing model does not distinguish between gradations of lateness.

The number of successful re-optimizations shows a remarkable result; Despite the fact that the heuristics-based sequencing model is specifically designed to improve the effectiveness of re-optimization, the configuration with the simulation-based sequencing model results in a larger number of successful re-optimizations than the configuration with the heuristics-based sequencing model. A possible explanation is that the heuristics are formulated based on an exploratory research performed using the small vehicle type only. It appears to be the case that the heuristics for maximizing the effect of re-optimization are not the same for different vehicle sizes.

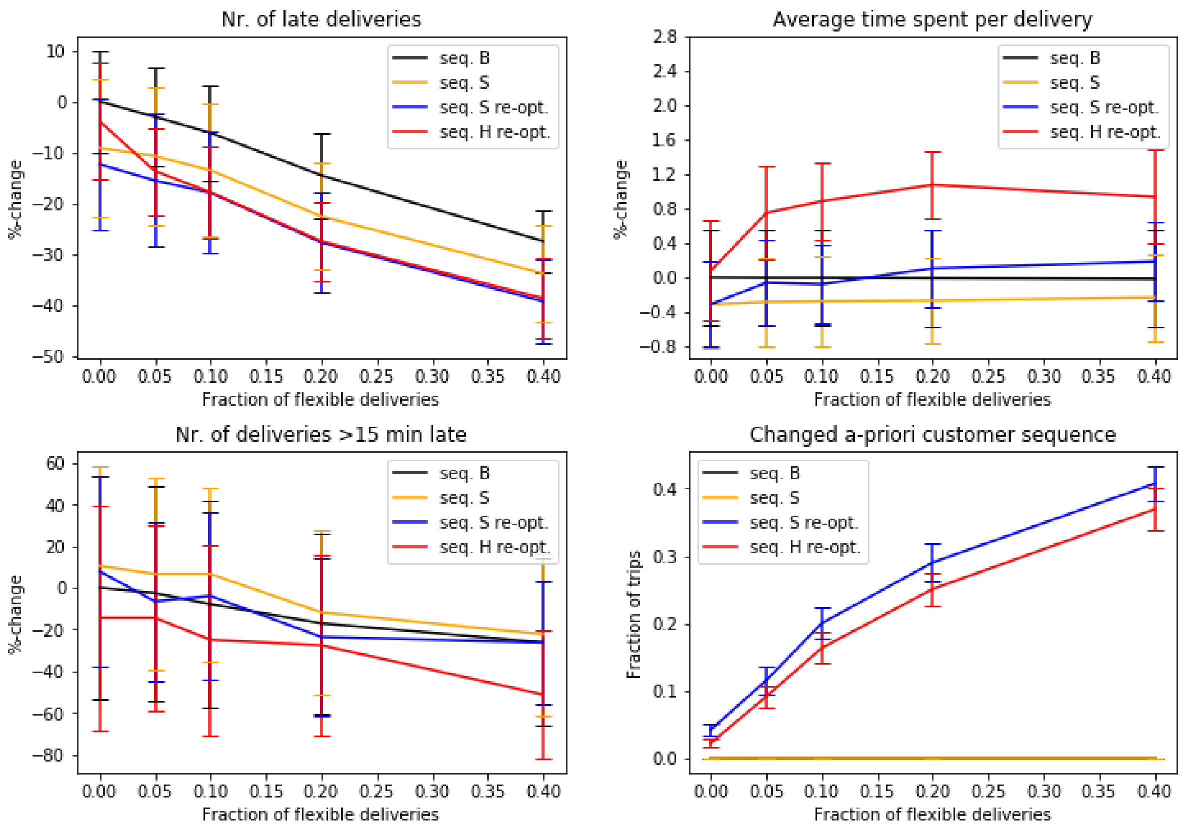

Figure 5.3: Results for the large vehicle type with default flexible delivery windows

5.4 Large vehicle and extended flexible delivery windows

Figure 5.4 shows that for extended flexible delivery windows the average time spent per delivery increases relative to the benchmark configuration for the configurations including the re-optimization model. When comparing the two configurations which include the re-optimization model it can be found that the heuristics-based sequencing model outperforms the simulation-based sequencing model in terms of the number of extreme lates. Interestingly, the simulation-based sequencing model without re-optimization model results in a decrease in the average time spent per delivery compared to the benchmark configuration. This suggests that when an a-priori solution is picked based on a combination of the expected on-time delivery rate and the planned trip duration, the realised trip duration is shorter than when an a-priori customer sequence is selected based on the planned total trip duration only. A possible explanation for these findings could be that the on-time rate suffers from the use of unreliable travel links in the trip. The use of unreliable travel links in the trip also increases the average time spent per delivery. Therefore, a maximization of the expected on-time delivery rate also reduces the average time spent per delivery.

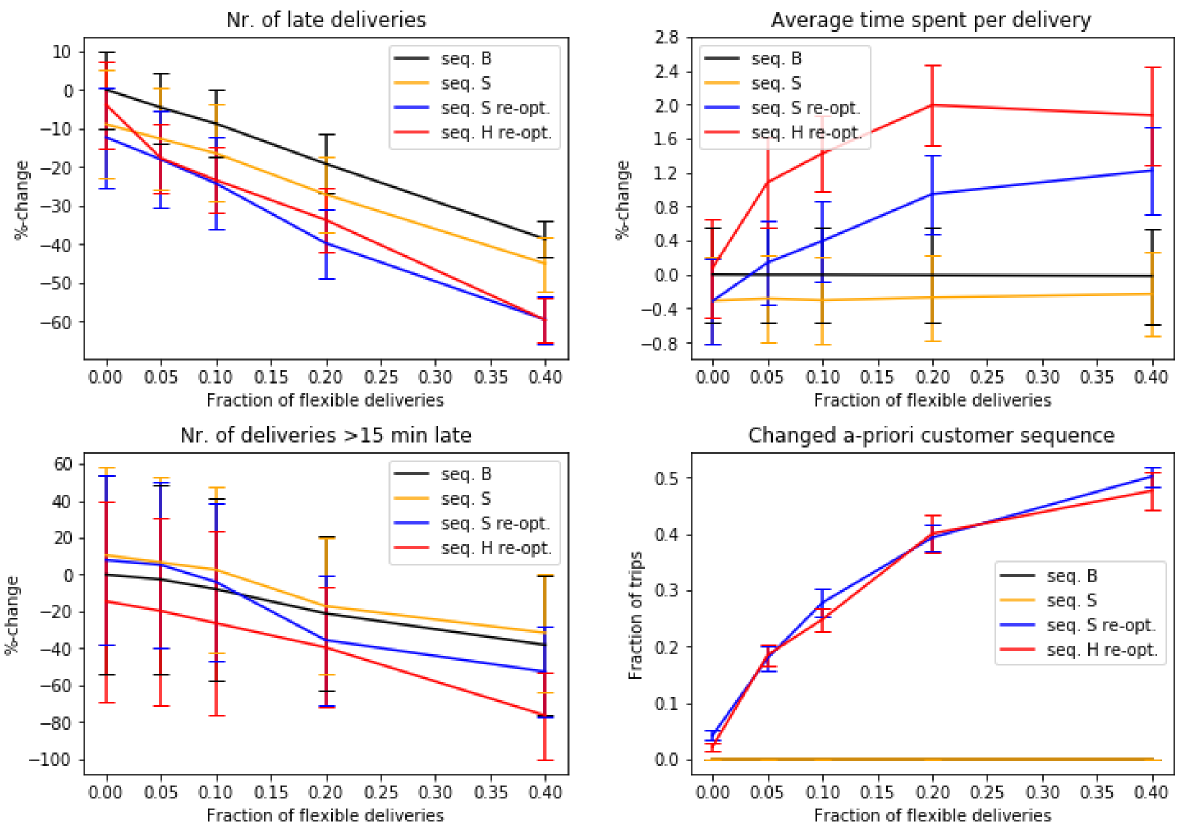

Figure 5.4: Results for the large vehicle type with extended flexible delivery windows

In order to illustrate the effects of the size of the flexible delivery window and the vehicle size on the performance of the routing model configurations, the simulation-based sequencing model is used in sections 5.5 and 5.6 as an example.

5.5 Comparison of flexible delivery window sizes

A comparison of the experimental results for different flexible delivery window sizes is presented in figure 5.5 for the small vehicle. The results show that the size of the flexible delivery window significantly affects the performance of the routing model. Larger flexible delivery windows increase the number of successful re-optimizations and thereby reduce the number of late deliveries and increase the average time spent per delivery. When comparing the effects of the different flexible delivery windows for the large vehicle, see figure 5.6, a similar pattern can be recognized. However, the difference in performance between the two types of flexible delivery windows is not as big as for the small vehicle.

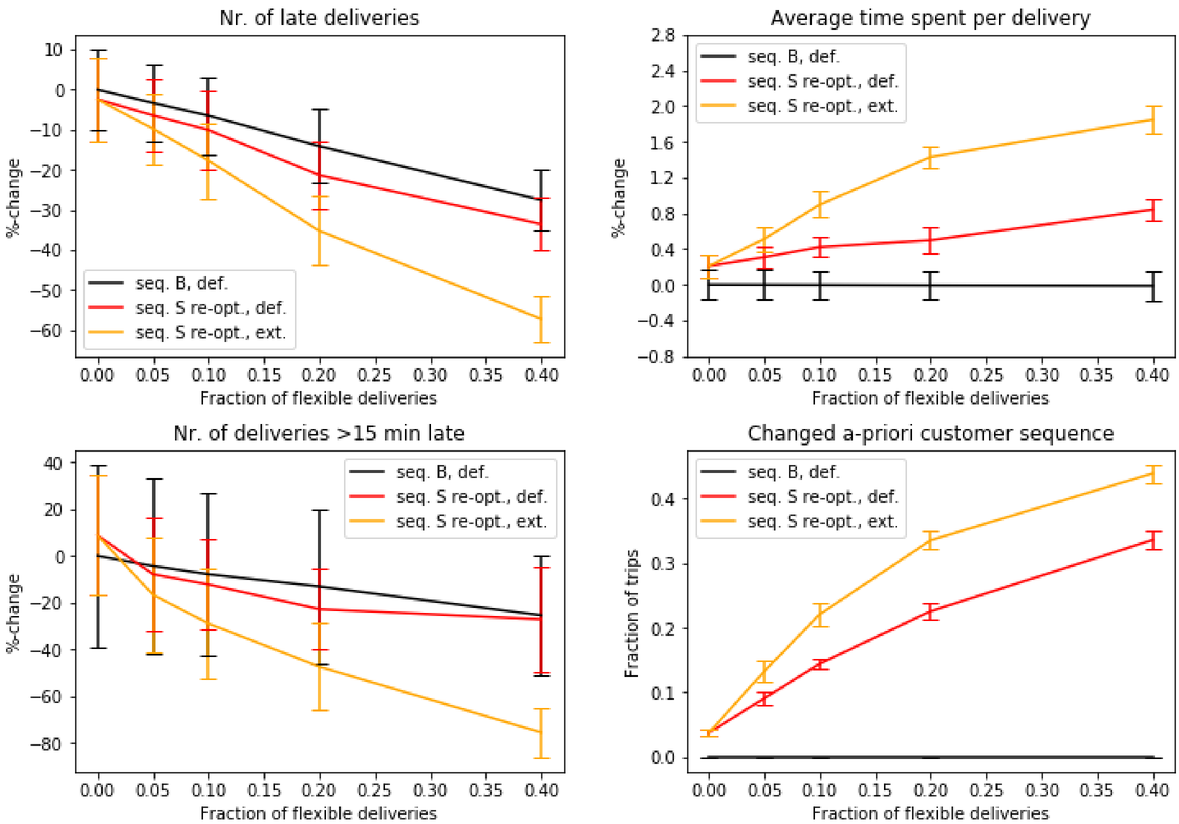

Figure 5.5: Results for the small vehicle type; comparison of default and extended flexible delivery windows

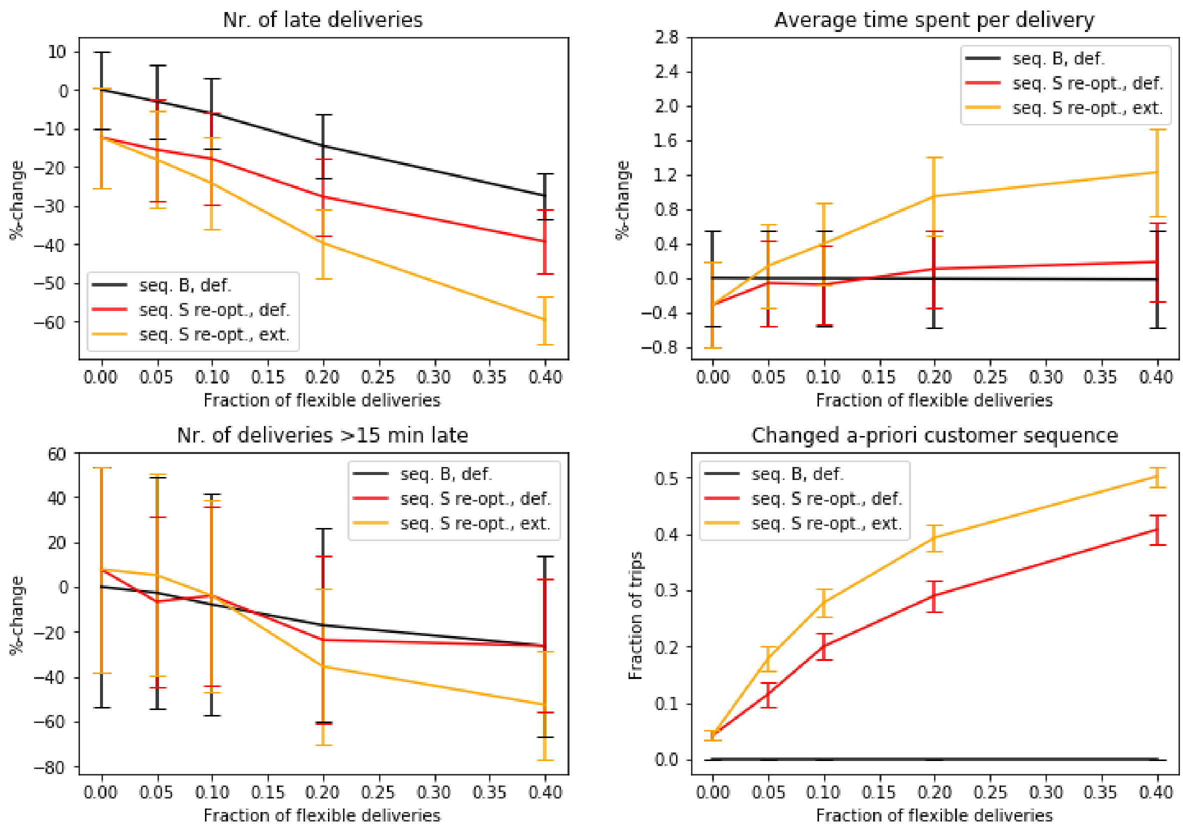

Figure 5.6: Results for the large vehicle type; comparison of default and extended flexible delivery windows

5.6 Comparison of vehicle types

For both types of vehicles, the configuration consisting of the simulation-based sequencing model and the re-optimization model is compared to the benchmark configuration for that same vehicle type. The results for both vehicle types are plotted in figure 5.7. The set of test instances used for experiments with the large vehicle type is a subset of those used for the small vehicle type (see section 4.5). In order to make a valid comparison, that same subset of test instances is used to generate the results as presented in figure 5.7 for the small vehicle.

The stochastic and dynamic routing model seems to improve the on-time delivery rate more for the large vehicle, however, the confidence intervals for the results obtained for the small vehicle are large. The investigated configuration is more effective at reducing the number of extreme lates for the small vehicle and increases the average time spent per delivery more for the small vehicle. The a-priori sequence of customers is changed after re-optimization in a larger fraction of trips for the large vehicle. This can be explained by the fact that more customers are served in one trip by the large vehicle, therefore on average more re-optimizations will be performed per trip.

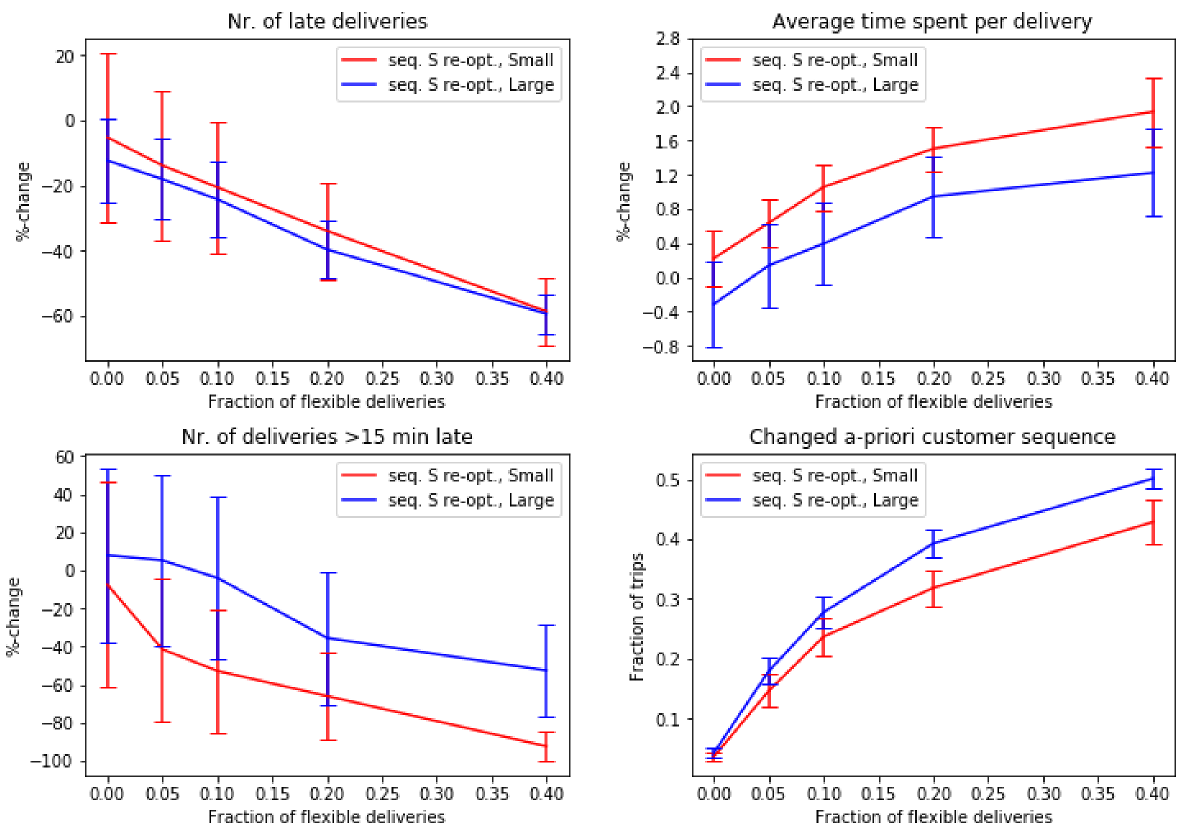

Figure 5.7: Comparison of the performance of the routing model for the small vehicle and the large vehicle with default flexible delivery windows

5.7 Conclusion

The presented analysis of the experimental results provides insights in the performance of the concept of flexible deliveries. The results demonstrate that the concept of flexible deliveries can significantly affect the on-time delivery performance of an e-grocer's routing model when the fraction of flexible deliveries equals 0.10 or more. The difference in terms of on-time performance between the benchmark configuration and the configurations which include the re-optimization model increases with the fraction of flexible deliveries. Regarding the average time spent per delivery, for these configurations the %-change increases rapidly for small fractions of flexible deliveries (< 0.10) and slowly for large fractions of flexible deliveries (> 0.10). The results regarding the number of late deliveries, number of extreme lates and average time spent per delivery are presented in tables 5.1, 5.2 and 5.3 for a fraction of flexible deliveries of 0.10.

It can be concluded that the configurations including a re-optimization model outperform the static and deterministic benchmark configuration in terms of on-time delivery performance. However, re-optimization comes at the price of an increased average time spent per delivery. The simulation-based sequencing model proves to be effective without re-optimization for the large vehicle type. For the small vehicle type the simulation-based sequencing model needs the re-optimization model to significantly outperform the benchmark configuration. When the re-optimization model is included in the configuration, the simulation-based and heuristics-based configurations show similar performance in terms of late deliveries. The heuristics-based approach appears to outperform the simulation-based approach in terms of the rate of extreme lates, especially for the large vehicle. However, the configuration including the simulation-based sequencing model results in a 1% lower average time spent per delivery, no matter the vehicle size or the size of the flexible delivery windows.

In order to decide which routing model configuration performs best regarding the specifications defined in chapter 1, a trade-off has to be made between the fraction of extreme lates and the average time spent per delivery. An example of such a trade-off is provided for a fraction of flexible deliveries equal to 0.10: When opting for the configuration with the heuristics-based sequencing model instead of the simulation-based sequencing model, for the large vehicle the number of extreme lates can be reduced by 24% in exchange for a 1% increase in the average time spent per delivery. For the small vehicle, the fraction of extreme lates can be reduced by 5% in exchange for a 1% increase in the average time spent per delivery. It can be concluded that a larger reduction of the fraction of extreme lates can be achieved at the same increase of operational costs for the large vehicle. In conclusion, it depends on the value proposition of the e-grocer which configuration performs best. In other words: How are the operational costs allowed to increase in order to decrease the number of extreme? Most e-grocers do not have such a figure available. Instead both solutions are assessed qualitatively and that assessment will result in the decision for either the simulation-based or heuristics-based sequencing model, both in combination with the re-optimization model.

Extension of the flexible delivery windows results in a significant improvement of the on-time delivery performance of the configurations. However, the increased number of successful re-optimizations burdens the average time spent per delivery. The added value of extended flexible delivery windows is largest for the small vehicle type. When comparing the two types of vehicles studied, it becomes evident that the choice for a specific routing model configuration has a larger effect on the KPIs for a large vehicle than for a small vehicle.

It has to be mentioned that the fraction of trips which are affected by the heuristics behind the heuristics-based sequencing model amounts to just 36% (see section 3.4.3). In order to gain insights in how the sequencing model performs on that subset of trips specifically, the performances of the different routing model configurations on that subset of trips is compared. This analysis is presented in appendix A. It can be concluded that, in general, the configuration with the heuristics-oriented sequencing model does not perform differently relative to the other configurations when only that subset of trips is considered. However, when looking at the trip instances regarding a large vehicle and extended flexible delivery windows, the on-time delivery performance of the heuristics-oriented sequencing model is relatively better when only the subset of trips is considered. This means that for these test instances the on-time delivery performance of the heuristics-oriented sequencing model can be improved when the fraction of trips which is affected by the heuristics behind this sequencing model is increased.

Table 5.1: Number of late deliveries for a fraction of deliveries equal to 0.10. For both vehicle types results are presented relative to the results obtained using default flexible delivery windows and the benchmark configuration.

Late deliveries					
Configuration		Small vehicle		Large vehicle	
Sequencing model	Re-optimization model	Default FDW[1]	Extended FDW[1]	Default FDW[1]	Extended FDW[1]
Benchmark	No	0.0 ± 10 %	-3.1 ± 9.6 %	0.0 ± 9.9 %	-2.8 ± 9.3 %
Simulation-based	No	-1.8 ± 11 %	-4.7 ± 11 %	-7.8 ± 14 %	-11 ± 13 %
Simulation-based	Yes	-3.9 ± 10 %	-12 ± 10 %	-13 ± 13 %	-19 ± 13 %
Heuristics-based	Yes	-4.5 ± 11 %	-14 ± 11 %	-12 ± 10 %	-18 ± 9 %

Table 5.2: Average time spent per delivery for a fraction of deliveries equal to 0.10. For both vehicle types results are presented relative to the results obtained using default flexible delivery windows and the benchmark configuration.

Average time spent per delivery					
Configuration		Small vehicle		Large vehicle	
Sequencing model	Re-optimization model	Default FDW[1]	Extended FDW[1]	Default FDW[1]	Extended FDW[1]
Benchmark	No	0.0 ± 0.16 %	0.0 ± 0.16 %	0.0 ± 0.56 %	0.0 ± 0.56 %
Simulation-based	No	0.2 ± 0.12 %	0.2 ± 0.12 %	-0.3 ± 0.52 %	-0.3 ± 0.52 %
Simulation-based	Yes	0.4 ± 0.11 %	0.9 ± 0.15 %	-0.1 ± 0.46 %	0.4 ± 0.48 %
Heuristics-based	Yes	1.4 ± 0.20 %	1.9 ± 0.2 %	0.9 ± 0.44 %	1.4 ± 0.45 %

Table 5.3: Number of extreme lates for a fraction of deliveries equal to 0.10. For both vehicle types results are presented relative to the results obtained using default flexible delivery windows and the benchmark configuration.

Extreme lates					
Configuration		Small vehicle		Large vehicle	
Sequencing model	Re-optimization model	Default FDW[1]	Extended FDW[1]	Default FDW[1]	Extended FDW[1]
Benchmark	No	0 ± 38 %	-4 ± 37 %	0 ± 54 %	0 ± 54 %
Simulation-based	No	-1 ± 23 %	-2 ± 24 %	16 ± 45 %	11 ± 48 %
Simulation-based	Yes	-5 ± 20 %	-23 ± 25 %	4 ± 43 %	4 ± 47 %
Heuristics-based	Yes	-11 ± 25 %	-27 ± 25 %	-19 ± 50 %	-20 ± 54 %

[1] Flexible delivery window

6
Conclusion

By means of answering the sub-questions as defined in section 1.4 an answer to the main research question can be formulated. The main research question of this thesis is: **"How can stochastic and dynamic routing models improve the delivery service of an e-grocer?"** Below the answers to the sub-questions are formulated one-by-one before an answer to the main research question is given.

"What are the approaches for stochastic and dynamic routing models used in literature?"

Stochastic and dynamic routing models have gained increased attention from the research community over the past decade. Although the literature on stochastic and dynamic routing models suited to the needs of e-grocers is very sparse, researchers have developed stochastic and dynamic routing models for many other applications. Stochastic routing models can be categorized based on the nature of their stochastic element(s). Most researchers assume independent stochastic elements and log-concave probability distributions in order to limit the computation time of the model. Dynamic routing models can be categorized based on the nature of their dynamic element: Based on which development is re-optimization of the a-priori route initiated? The majority of the dynamic routing models uses deterministic parameters during re-optimization because the available computation time is limited compared to a-priori optimization. The combination of stochastic and dynamic modelling approaches has not received a lot of attention yet. However, the added value of this combination of modelling approaches is widely acknowledged. A well-thought use of stochastic and dynamic elements in a routing model can relieve the burden of model tractability to a large extent. Careful consideration should be made concerning the performance improvement due to each added complexity.

"What is a promising approach to apply stochastic and dynamic routing models for e-grocers?"

When the findings of other researchers in the field of stochastic and dynamic modelling approaches are combined with the requirements which an e-grocers poses for its routing model, the back-bone of a promising routing model can be designed. The routing model should consist of three separate sub-models which are run in a sequential manner. First the customer-trip assignment model determines which customers are combined in the same trip. It makes use of deterministic input parameters because of the short computation time available for this model. Next, the a-priori sequencing model uses stochastic service times and stochastic and inter-dependent travel times to compute the optimal sequence of customers in a given customer-trip assignment. Lastly, during the trip, the real-time re-optimization model uses stochastic service times and real-time deterministic travel times to reconsider the sequence of residual customers upon departure from the customers.

Inspired by stochastic and dynamic modelling approaches used by other researchers, the concept of flexible deliveries is introduced. This solution approach relies on customers who opt for flexible deliveries. Flexible deliveries have a larger communicated delivery time window within which the groceries have to be delivered. When a fraction of all deliveries in a trip is flexible, the sequence of residual customers in a trip can be changed during the trip in order to mitigate the effects of running late or running early. The concept of flexible deliveries is further investigated in this thesis.

“How does the proposed routing model perform?”

In order to quantify the performance of a routing model making use of flexible deliveries, different routing model configurations are investigated by means of a computational experiment. Three different types of sequencing models are investigated: A deterministic sequencing model, a stochastic sequencing model which uses a-priori simulation to choose the best solution out of a set of good solutions, and a deterministic sequencing model which optimizes the positions of flexible deliveries in the customer sequence in such a way that the effectiveness of re-optimization is maximized. The experiment uses test instances which were encountered in reality by e-grocer Picnic and the simulation environment uses historic data from Picnic to simulate realisations of planned trip times. Two different types of customer-trip assignments are investigated: for a small vehicle and for a large vehicle. Moreover, two different sizes of the flexible delivery window are studied: 60 minutes and 75 minutes.

From the experimental results it can be concluded that the configurations including a re-optimization model outperform the static and deterministic benchmark configuration in terms of on-time delivery performance. However, re-optimization comes at the price of an increased average time spent per delivery. The use of a re-optimization model significantly improves the on-time delivery performance when at least 10% of all deliveries is flexible; With this percentage of flexible deliveries, for the small vehicle type the number of late deliveries can be reduced by 14% and the number of extreme lates by 27%. Moreover, for the large vehicle type the number of lates can be reduced by 18% and the number of extreme lates by 20%. Below the relative performances of the different routing model configurations are discussed for the case in which 10% of all deliveries are flexible.

The simulation-based sequencing model proves to be effective at improving the on-time delivery performance for the large vehicle type. Even without the re-optimization model both the number of late deliveries and the number of extreme lates are reduced by 11%. For the small vehicle type the simulation-based sequencing model needs the re-optimization model to significantly outperform the benchmark configuration, reducing the number of late deliveries by 12% and the number of extreme lates by 23%.

When the re-optimization model is included in the configuration, the simulation-based and heuristics-based sequencing models show similar performance in terms of on-time delivery rate. The heuristics-based sequencing model outperforms the simulation-based sequencing model in terms of the number of extreme lates by 5% for the small vehicle and by 16% for the large vehicle. However, the configuration with the simulation-based sequencing model results in a 1% lower average time spent per delivery. In other words, the heuristics-based sequencing model significantly reduces the number of extreme lates in exchange for a 1% increase in the average time spent per delivery.

Extending the flexible delivery windows results in a significant improvement of the on-time delivery performance of the configurations, especially for the small vehicle type: up to 10% reduction of the number of late deliveries, up to 16% reduction of the number of extreme lates. However, the increased number of successful re-optimizations increases the average time spent per delivery by 1% compared to the case with default flexible delivery windows. The added value of extended flexible delivery windows is largest for the small vehicle type. When comparing the two types of vehicles studied, it becomes evident that the choice for a specific routing model configuration has a larger effect on the KPIs for a large vehicle than for a small vehicle.

When an e-grocer is very costs-sensitive a 1% increase in the operational costs would tilt the balance in the favour of the configuration consisting of the simulation-based sequencing model and the re-optimization model. For e-grocer’s which have more financial possibilities for an increase in the operational costs in exchange for a significant improvement of the customer experience, the configuration consisting of the heuristics-based sequencing model and the re-optimization model is favourable.

“How can stochastic and dynamic routing models improve the delivery service of an e-grocer?”

In conclusion, stochastic and dynamic routing models can significantly improve the on-time delivery performance of an e-grocer’s last-mile distribution system. The concept of flexible deliveries was proven to successfully add a degree of freedom to the routing system which is required to make effective use of real-time re-optimization. The experimental results demonstrate that the use of a stochastic a-priori sequencing model improves the performance of the routing model on the KPIs as defined by an e-grocer. This sequencing model uses a-priori simulation to choose a solution out of a set of good solutions. The effectiveness of the concept of flexible deliveries depends on the type of vehicles used, the fraction of flexible deliveries and the size of the flexible delivery windows. When compared to the static and deterministic benchmark configuration, the potential gain through the use of a stochastic and dynamic configuration is most profound for trips completed with the large vehicle. When the fraction of flexible deliveries increases, the effects of re-optimization become more evident. When 10% of all deliveries is flexible, real-time re-optimization can significantly improve the performance of the last-mile distribution system. Extending the flexible delivery window results in further improvement of the on-time delivery performance of the routing model, however, it also increases the average time spent per delivery and thereby increases the operational costs for delivery.

7 Discussion

This thesis investigates how the concept of flexible deliveries affects the on-time delivery performance and the operational costs of an e-grocer's last mile delivery service. However, when considering the implementation of a new type of routing model other aspects have to be taken into account as well. Therefore, the organisational impact of the concept of flexible deliveries is discussed in section 7.1. Lastly, section 7.2 critically reviews the conducted research and formulates suggestions for further research.

7.1 Organisational impact

In order for the concept of flexible deliveries to significantly improve the performance of the last-mile distribution system, 10% of all customers should opt for a flexible delivery. Customers require some type of incentive to opt for a flexible delivery instead of a regular delivery. In order to provide a matching incentive for the selection of a flexible delivery, e-grocers have to know how much the customers value the size of their delivery time window. The value of a regular delivery time window as opposed to a flexible delivery window might also vary between different types of customers. For example, for business-customers increasing the size of delivery time window might not harm the customer-experience because a store or office will be open the whole day anyhow. However, for working parents with young kids who have tight schedules outside working hours as well, increasing the size of the delivery time window might result in a large depreciation of the customer experience. For this reason, an e-grocer should analyse its customer-base in order to propose an adequate incentive for opting for a flexible delivery. Examples of incentives are a discount or a membership which allows for extra ordering options.

In order to improve the customer experience for a flexible delivery, a notification could be send to the customer shortly before the delivery will start (e.g. upon departure from the previous customer). Such a notification adds more value when the delivery time window is larger; When the initial uncertainty of the delivery time is large (e.g. 60 minutes) a customer is likely to appreciate a warning shortly before the delivery starts because the customer will be undertaking other activities in that delivery time window as well. A shorter delivery time window will probably cause a customer to schedule other activities outside of the delivery time window. Therefore, a notification will not increase the customer experience as much. The effect of such a notification on the customer experience has to be investigated in more detail.

Taking into account the type of incentive that is used to attract a sufficient number of flexible deliveries, the financial consequences of the concept should be evaluated. The increased average time spent per delivery and potentially the rewards offered to customers who opt for flexibility result in higher costs per delivery. However, the improved customer experience due to the improved on-time delivery performance might (indirectly) increase the number of customers and the average order size resulting in a decrease of the delivery cost per item. The financial consequences of implementing the concept of flexible deliveries have to be researched before the concept can be taken to the implementation stage.

When an e-grocer replaces the benchmark sequencing model by the simulation-based sequencing model the complexity of the routing model will increase. First of all, multiple solutions have to be calculated for each trip. Secondly, each solution has to be simulated a number of times in order to select the optimal solution out of the set of good solutions. The complexity of the simulation-based sequencing model will require more effort for implementation, and the computational power required to run the model will increase. On the contrary, replacing the benchmark sequencing model by the heuristics-based sequencing model only requires an adaptation of the benchmark sequencing model. The computational power needed will remain similar and implementation will not be very time consuming. For this reason, the heuristics-based sequencing model is easier to implement.

Implementation of the real-time re-optimization model requires more elaborate changes to the current routing system. In the current operation, the driver's hand-held device already sends information to the central server. In order to make the real-time re-optimization model work, reverse communication should be enabled as well; a re-optimized trip planning has to be sent from the central server to the hand-held device. The re-optimizations have to be performed on a central server. Although the size of the re-optimization problem is small, a large number of re-optimization problems has to be simultaneously. Depending on the capacity of the central server, this might require an investment in server capacity. In conclusion, implementation of the re-optimization model requires a significant upfront investment in the form of server capacity and IT-development hours.

In the current operation, drivers know the full trip planning before departure. When the re-optimization model is used, it would cause confusion when the driver knows the a-priori planned trip. Therefore it is suggested to only communicate the next customer in the trip with the driver. It is interesting to investigate how drivers would experience this uncertainty. Positive effects could be that drivers are more focussed on the next customer, instead of already worrying about the trip planning further ahead. This could relieve the working pressure experienced drivers. Moreover, the re-optimization model improves the on-time delivery performance, therefore drivers have to worry less about arriving late at customers. A relaxed driver is likely to provide a better customer experience at the front door. Furthermore, a relaxed driver drives safer than a driver in a hurry. On the other hand, only communicating the next customer in the trip takes away a sense of responsibility from the driver. In general, a driver feels responsible for arriving on time at the customers. When information regarding all deliveries except the next delivery is shielded from the driver, the driver's responsibility for arriving on-time is reduced because the driver's freedom of own interpretation of the trip planning is taken away. Therefore other measures should be used to increase the sense of responsibility of the drivers when a customer-by-customer re-optimization approach is used.

7.2 Critical reflection and directions for further research

In this study, different approaches to the use of the concept of flexible deliveries are investigated. However, it is not researched how the concept of flexible deliveries compares to alternative solution approaches. For this reason, the performance of the concept of flexible deliveries should be compared to the performance of stochastic and dynamic routing models proposed by other researchers. In order to perform a meaningful comparison, the solution approaches from literature should be adaptable to the use-case of e-grocers. The performance of different types of solution approaches can be compared in two different ways: either the solution approach of flexible deliveries is tested on other (benchmark) test instances, or the solution approaches from literature are tested on Picnic's case-study test instances.

The literature study as presented in chapter 2 yields the promising solution approach of flexible deliveries. In more detail, the outcome of the literature study suggests a specific modelling approach including the modelling of inter-dependent travel-times for determining the a-priori sequence of customers and the use of stochastic service times and deterministic real-time travel times for re-optimization. The different configurations investigated in this thesis do no include these three model features. Inter-dependent travel times are left out of the scope of this thesis because of the added complexity in the simulation environment that is required to evaluate the performance of modelling travel times as inter-dependent. In this research stochastic service times were not taken into account during re-optimization because it results in unworkably long run times for the experiments. Other researchers could build upon this thesis by investigating the use of these stochastic service times during re-optimization by means of an experiment comprising of fewer test instances or by means of a faster computer to run the experiments. Lastly, the use of real-time travel times during re-optimization is not included in this thesis because it is not a viable solution from a business perspective. E-grocers pay travel time data providers for the travel times they request. If for each re-optimization real-time travel times are requested from a travel time data provider the costs would become very high.

The simulation-based sequencing model uses a-priori simulation of a set of good solutions to find the optimal solution. The a-priori simulation does not include re-optimization. It is interesting to investigate the added value of including the possibility for re-optimization in the a-priori simulation. However, attention should be paid to the computational demand of such a sequencing model because the re-optimization steps during a-priori simulation are expected to significantly increase the computation time. Moreover, sequencing model 1 only compares three solutions. However, as demonstrated by figure 3.7(b), the performance of the sequencing model might be improved when more solutions are investigated. However, increasing the number of solutions to investigate also increases the computational demand. Therefore, careful attention should be paid to the trade-off between improvement of the model's performance and computational demand.

The simulation-based sequencing model picks the best solution out of a set of good solutions based solely on the simulated average on-time delivery rate. This is confirmed by the results of the configuration including the simulation-based sequencing model; when comparing the number of late deliveries and the number of extreme late deliveries it is demonstrated that the gap between the simulation-based sequencing model and the heuristics-based sequencing model is larger regarding the fraction of extreme lates. For the simulation-based sequencing model it has to be investigated whether the performance regarding extreme lates can be improved if the extreme lates are considered when the "best" solution is picked out of a the set of good solutions.

The heuristics behind the heuristics-based sequencing model are formulated based on an exploratory research using the small vehicle type only. In order to maximize the effectiveness of the heuristics-based sequencing model for the large vehicle type, another exploratory research should be performed. That exploratory research would yield heuristics specifically effective for the large vehicle type.

The simulation environment used in the computational experiment does not simulate a correlation between the working speed of a driver and the progress of the trip with respect to the trip planning. In order to close the gap between the simulation and real operations, it should be investigated whether such a correlation exists. However, it can be argued that a working speed independent of the progress of a trip might be preferred; A driver in a hurry does not improve the customer experience at the front door and is likely to drive less safe (Pandi et al., 2019, Rendón-Vélez et al., 2016).

The customer-trip assignment model is left out of the scope of this research. Customer-trip assignments which are optimized for use with a static deterministic routing model are used as test instances for the experimental method. It would be interesting to further investigate the potential of customer-trip assignment models. For example, an additional objective of the customer-trip assignment model could be to spread the flexible deliveries evenly over all trips. Moreover, the customer-trip assignment model could be adjusted in such a way that the flexible deliveries are spread more evenly amongst the different planning windows within each trip. Investigation of such changes in the customer-trip assignment model might lead to further improvement of the routing model.

The computational experiments presented in this thesis use test instances from a single hub. It should be investigated whether the relative performance of the routing model configurations is the same for different hubs. Furthermore, the experiment should be performed using test instances from a variety of e-grocers in order to conclude whether the results obtained in this thesis are representative for the e-grocery business as a whole.

A
Results subset of trips

This appendix presents a comparison of the experimental results obtained for the subset of trips for which the heuristics-based sequencing model can yield a different sequence of customers compared to the benchmark sequencing model. In other words, this is the subset of trips for which the heuristics behind the heuristics-based sequencing model are designed. This subset consists of trips for which at least one of the planning windows contains exactly one flexible delivery. For the small vehicle type this amounts to 36% of all tested trip instances. For the large vehicle type this amounts to 45%.

The results as presented in figures A.1, A.2, A.3 and A.4 show that the correlation between the fraction of flexible deliveries and the performance on the KPIs is less evident when only a subset of all trips is investigated. This can be expected because the total set of trips is filtered based on the number of flexible deliveries in each planning window. For large fractions of flexible deliveries this means that the number of flexible deliveries in a trip will be lower on average. For small fractions of flexible deliveries the average number of flexible deliveries in a trip will be higher.

A.1 Small vehicle and default flexible delivery windows

For default flexible delivery windows and the small vehicle type (see figure A.1), the relative on-time delivery performance of the configurations with the re-optimization model is not significantly different compared to the results when all tested trips are considered. However, the average time spent per delivery shows different relative performances. When only the subset of trips is considered the increase in terms of average time spent per delivery is much larger for the configurations with a re-optimization model.

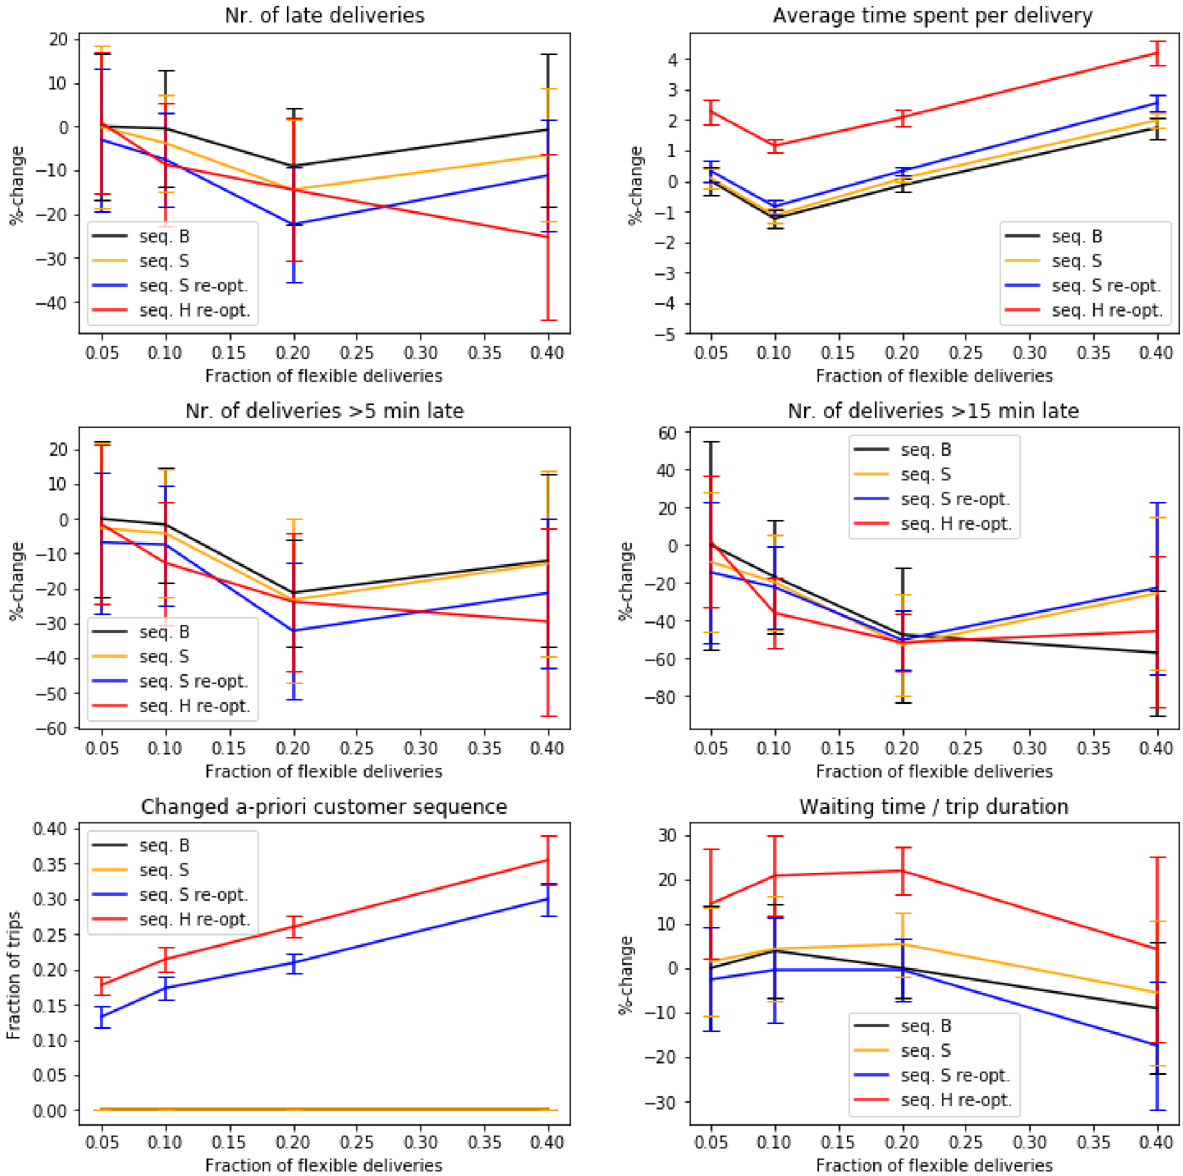

Figure A.1: Results for the small vehicle type with default flexible delivery windows

A.2 Small vehicle and extended flexible delivery windows

Extension of the flexible delivery windows results in a small relative improvement of the configuration consisting of the heuristics-based sequencing model and the re-optimization model. The average time spent per delivery is still larger for this sequencing model but the gap with the simulation-based sequencing model is smaller.

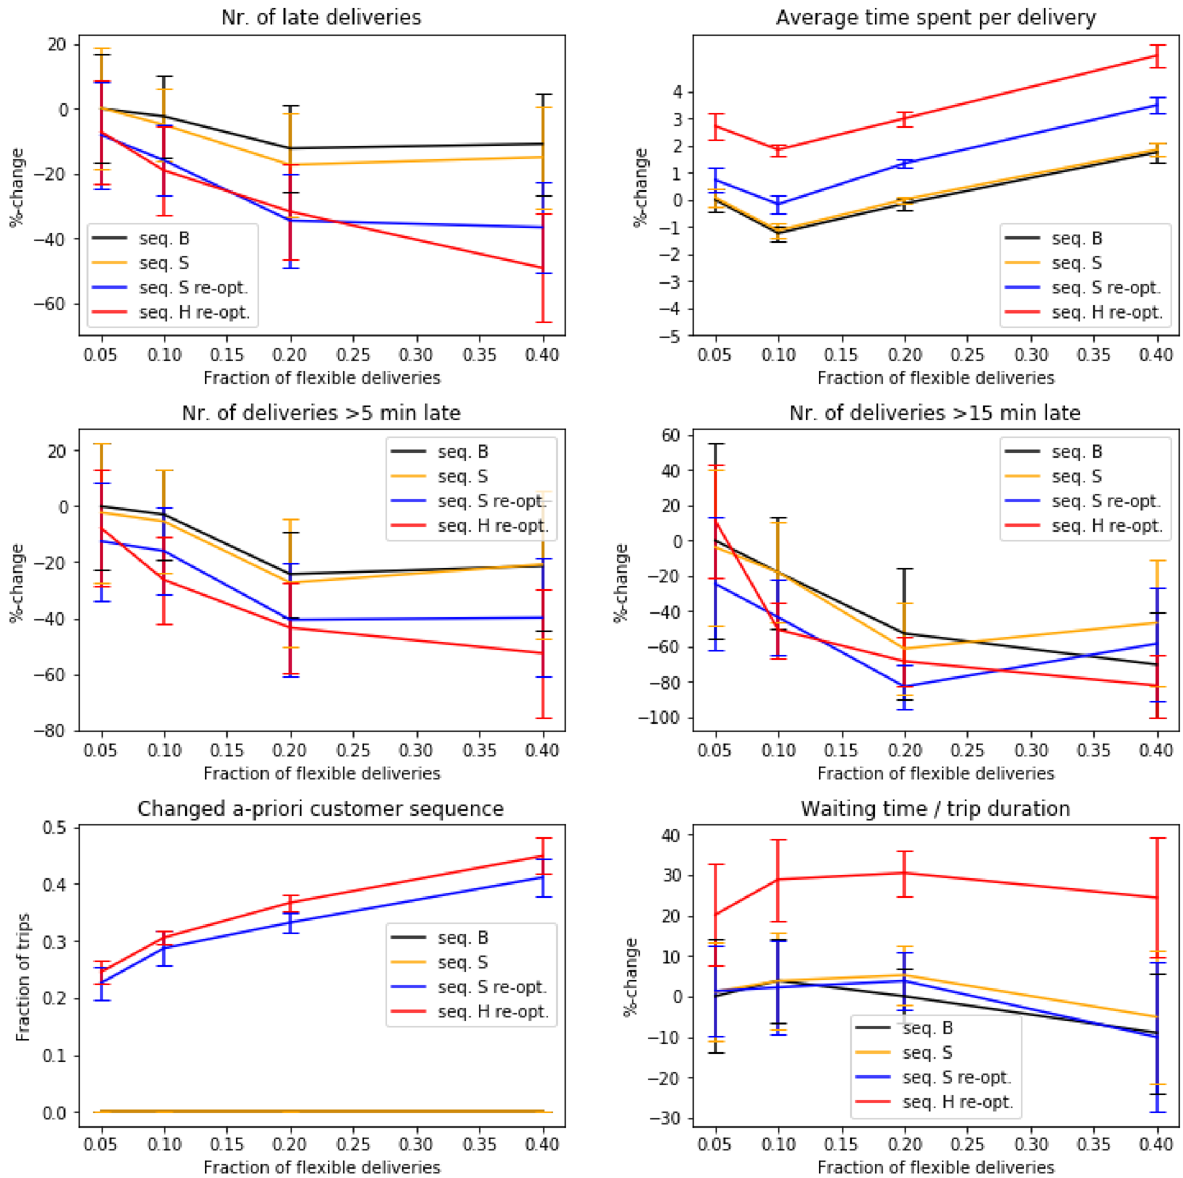

Figure A.2: Results for the small vehicle type with extended flexible delivery windows

A.3 Large vehicle and default flexible delivery windows

For the large vehicle type (see figure A.3), the heuristics-based sequencing model shows an improvement in terms of extreme lates for this subset of trips compared to the other sequencing models. Apparently, the heuristics-based sequencing model benefits more from the extension of the flexible delivery windows. However, the difference in terms of average time spent per delivery increases as well, resulting in a much higher average time spent per delivery for this sequencing model.

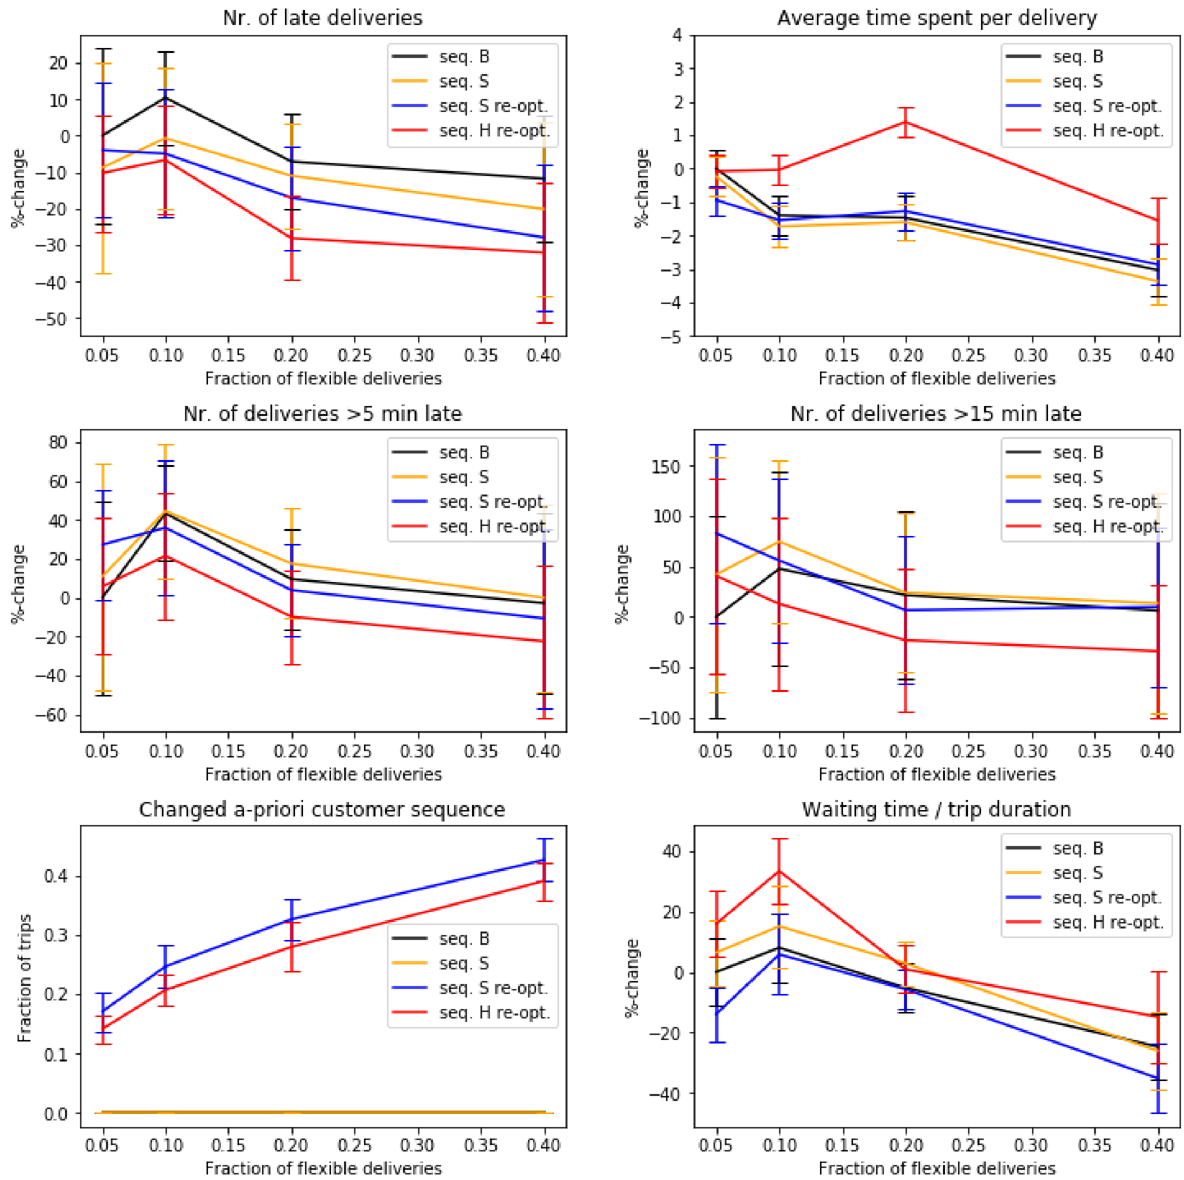

Figure A.3: Results for the large vehicle type with default flexible delivery windows

A.4 Large vehicle and extended flexible delivery windows

For the large vehicle type the relative performance of the configurations is not changed significantly when the flexible delivery windows are extended, see figure A.4.

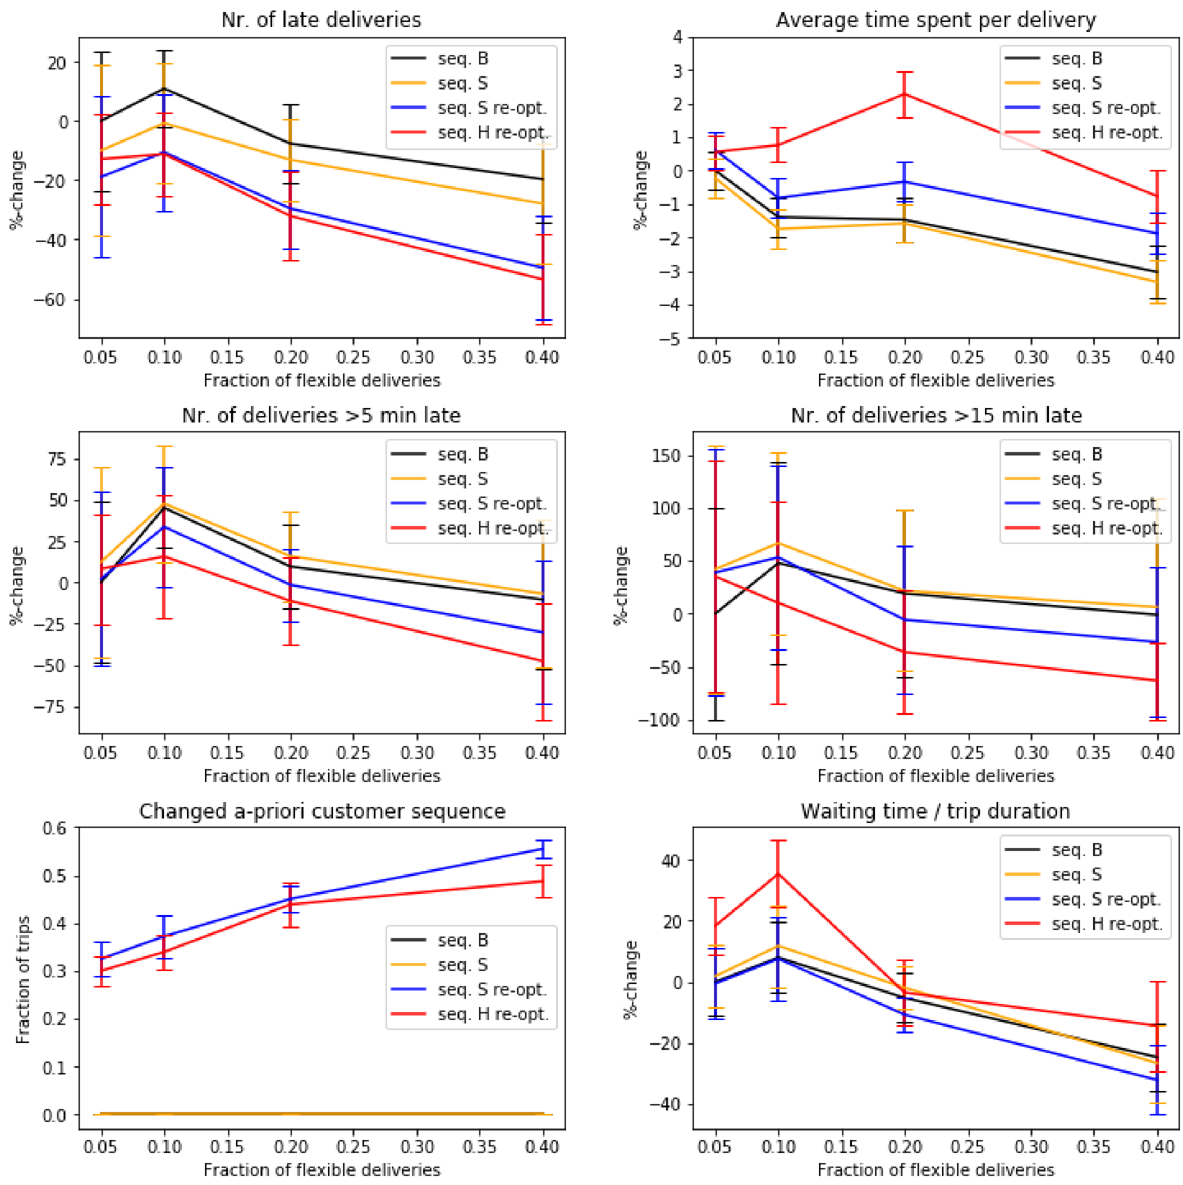

Figure A.4: Results for the large vehicle type with extended flexible delivery windows

B

Paper

Stochastic and dynamic routing models for improving the on-time performance of an e-grocer's delivery service

P. S. Bouwstra[a], Dr. B. Atasoy[a], Dr. Ir. G. Correia[b], Ir. P. M. Bijl[c], Prof. Dr. R. R. Negenborn[a]

[a]*Department of Marine and Transport Technology, Delft university of Technology, Delft, The Netherlands*
[b]*Department of Transport and Planning, Delft university of Technology, Delft, The Netherlands*
[c]*Picnic Technologies B.V., Amsterdam, The Netherlands*

***Abstract*—E-grocers offer their assortment online and deliver a customer's groceries at home. The quality of their delivery service is a crucial asset for an e-grocer to create and maintain a loyal customer-base. Because the online share of the grocery market has grown rapidly over the last decade, there is an urgent need for e-grocer specific routing systems. Although recently stochastic and dynamic routing models are studied for a wide range of applications, it is yet unresearched how stochastic and dynamic routing models can be used to improve the last-mile delivery service for e-grocers specifically. This research investigates the concept of flexible deliveries which introduces differentiated time window sizes. This creates possibilities for real-time re-optimization of the sequence of customers in a trip. The potential of flexible deliveries is investigated by means of computational experiments in which historic trip instances from Dutch e-grocer Picnic are used. It can be concluded that the configurations including a re-optimization model outperform the static and deterministic benchmark configuration in terms of on-time delivery performance. However, re-optimization comes at the price of an increased average time spent per delivery. The effects of re-optimization become significant for percentages of flexible deliveries larger than or equal to 10%. It was found that when 10% of the deliveries is flexible, the number of late deliveries can be reduced by upto 18% and the number of extreme lates ($\geq$ 15 minutes late) by upto 27%. This improved on-time delivery performance comes at the price of a maximum 2% increase in the average time spent per delivery.**

I. Introduction

E-grocers offer their customers a comfortable grocery shopping experience by allowing them to order from virtually any location and within moments. The groceries are delivered at the customer's front door or even into the kitchen. When compared to other last-mile distribution systems, e-grocers experience large service times relative to their travel times. Moreover, an important requirement for a good customer experience is the interaction between the customer and the driver upon delivery. The e-grocers' efforts to perform well on this aspect lead to a large uncertainty in these service times.

The first e-grocer businesses emerged in the late 1990s [1] and since the 2010s [2] e-grocers have started to seriously compete with traditional grocery stores. Because of the success of e-grocer start-ups, many traditional food retail market players have started to invest in an online grocery delivery service as well [3].

Because e-grocers have only recently started to gain a significant market share in the total grocery market [4], not much research has been dedicated to the development of routing models for e-grocers specifically. Using an of-the-shelve routing model does not accurately take into account the peculiarities concerning the last-mile distribution of groceries, such as the perishability of the commodity and the importance of the customer's trust in the quality of an e-grocer's delivery service [5], [6], [7].

In order to design routing models tailored to the needs of e-grocers, first the objectives and requirements for an e-grocer's routing model are determined by means of a case-study at Dutch e-grocer Picnic. The strict requirements for an e-grocer's routing model are: (1) Offer customers a free one-hour delivery time window when they place an order. (2) Offer free communication of a 20-minute delivery time window at the morning of the delivery. (3) Maximum calculation time for combinations of customers in a trip is one hour. (4) Maximum calculation time for the trip planning is six hours. When these requirements are met, the objectives for the routing model are to maximize the on-time delivery performance and to minimize the operational costs. In the context of this research, the on-time delivery performance is quantified by means of the on-time delivery rate and the rate of extreme lates ($\geq$ 15 minutes late). The operational costs are quantified by means of the average time spent per delivery.

The research community has made rapid developments in the fields of stochastic and dynamic routing models during the past decade [8], [9]. The results of these studies suggest that the use of stochastic and dynamic routing models improves the performance of routing models for a large variety of applications, e.g. [10], [11], [12]. However, it is yet unresearched how these stochastic and dynamic routing models can be used to improve the last-mile delivery service for e-grocers specifically.

The objective of this research is the formulation of a promising modelling approach to use stochastic and dynamic routing model elements to improve the performance of an e-grocer's routing model. The effectiveness of the proposed modelling approach is assessed by means of a set of computational experiments that use e-grocer specific test instances.

The different approaches encountered in literature are tested against the specifications for an e-grocer's routing model. This yields the promising solution approach of flexible deliveries, which is further investigated in this paper. "Flexible" deliveries have a larger delivery time window than "regular" deliveries.

These larger delivery time windows allow for dynamic re-optimization of the residual sequence of customers in a trip and thereby can mitigate the effects of running early or running late with respect to the trip planning. This strategy has the potential to improve the on-time delivery rate and reduce the waiting time caused by early arrivals at customers. Moreover, it does not incur significant additional operational costs because (1) the number of vehicles remains the same and (2) changing the sequence of customers in an e-grocer trip does not result in a large increase in total travel time because the travel times between any pair of customers served in the same trip are short.

The structure of this paper is as follows. First, related work on the topic of dynamic and stochastic routing models is presented in Section II. Different types of routing models that make use of the concept of flexible deliveries are designed and explained in Section III. Section IV explains how the performance of these routing model variants (“configurations”) is evaluated and presents the experimental results. The paper is rounded of with conclusions and recommendations for future research in Section V.

II. Literature review

This section discusses approaches to the use of stochastic and dynamic routing model elements used by other researchers. Studying their approaches has led to the formulation of the concept of flexible deliveries. Their work is compared to the concept of flexible deliveries and the similarities and differences are discussed.

[13] researched an approach to solve the capacitated vehicle routing problem with stochastic demands and introduce the concept of premium customers. For premium customers, the probability that their demand is met is larger than for other customers. This idea of making a distinction between customers is also investigated in this paper. However, whereas [13] make a distinction with regards of the probability of a customer’s demand being met, the concept of flexible deliveries makes a distinction with regards of the delivery time window size.

Rerouting strategies which concern the exchange of certain customers in the remaining sequence of a trip were investigated by [14]. They use real-time traffic data to make these adjustments to the trip planning. This paper also investigates the added value of re-optimization of the sequence of residual customers in a trip. However, the problem studied in this paper involves time window constraints during the re-optimization stage whereas the problem addressed by [14] does not.

[15] investigate the vehicle routing problem with hard time windows and stochastic service times. The main difference between the problem they studied and the problem studied in this paper is the complexity of the recourse action. [15] investigate two simple recourse actions: skip the service at the current customer or skip the service at the next customer. In the context of e-grocers such recourse actions are unviable. In this paper a more sophisticated re-optimization of the sequence of residual customers is investigated.

In the context of an urban freight transport problem [16] look into the potential of using real-time traffic times for re-optimization of the allocation of customers to trucks and the sequence of customers in a trip. They found that incorporating real-time traffic information results in an improved reliability of the arrival times at customers. As opposed to the problem studied in this paper their problem includes multiple degrees of freedom which can be re-optimized: the number of trucks, trip departure times and sequence of customers in each trip. In the routing problem investigated in this paper only the sequence of customers in each trip can be re-optimized. This makes the problem more constrained and therefore it is more challenging to make a significant impact by means of re-optimization.

From Table I it can be concluded that the combination of using a re-optimization model to reconsider the optimal sequence of residual customers in a trip and the presence of time window constraints is precedented by [16]. However, [16] only considers stochastic travel times while this paper also considers stochastic service times. Moreover, this paper focusses on trip instances specifically encountered by e-grocers. The other studies listed in Table I address a general application of their problem type. Lastly, most researchers in the field of stochastic and dynamic routing models address a variant of the vehicle routing problem, whereas this paper focusses on the travelling salesman problem. As a consequence the degree of freedom of the problem posed in this paper is smaller compared to most studies found in the scientific literature.

III. Routing models

The VRP solved by e-grocers is large in size. The resulting customer-trip assignments have to be computed within one hour. However, calculation of the planned arrival times at each customer is allowed to take up to six hours. Therefore, it is interesting to split the traditional VRP into two separate problems: The customer-trip assignment problem and the sequencing problem. Because of the short computation time available to solve the assignment problem, a deterministic approach is suggested. In the sequencing problem it is feasible to use stochastic travel and service times in order to determine the optimal sequence of customers because the size of this problem is limited and the computation budget is large. The a-priori sequence of customers should be computed in such a way that online re-optimization of the customer sequence is effective at maximizing the performance of the delivery service. When the trip has begun, re-optimization takes place shortly before the driver departs from a customer. The structure of the routing model is summarized in Figure 1. The scope of this study is limited to the a-priori sequencing model and the real-time re-optimization model, a combination of which is referred to as a “configuration”.

In order to investigate the performance of different approaches to the concept of flexible deliveries, three different sequencing models and one re-optimization model are designed which are presented in respectively Subsections III-A, III-B, III-C and III-D.

TABLE I: Related literature

Author(s)	**Problem type**	**Stochastic element**	**Dynamically re-optimized element(s)**	**Time constraints**
[13]	VRPSD[a]	Quantity of customer demand	None	None
[14]	DVRP[b]	Travel times	Sequence of residual customers in a trip	None
[15]	DVRPTW[c]	Service times	Skipping service at a customer	Yes
[16]	DVRPTW[c]	Travel times	1. Number of trips 2. Trip departure times 3. Sequence of customers in each trip	Yes
This paper	DTSPTW[d]	1. Service times 2. Travel times	Sequence of residual customers in a trip	Yes

[a]Vehicle Routing Problem with Stochastic Demands
[b]Dynamic Vehicle Routing Problem
[c]Dynamic Vehicle Routing Problem with Time Windows
[d]Dynamic Travelling Salesman Problem with Time Windows

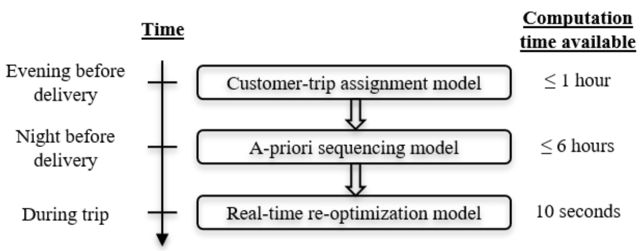

Fig. 1: Structure of the routing model as proposed based on the literature review

A. Benchmark sequencing model

This model is referred to as the benchmark sequencing model because of its simplicity and deterministic nature. The notation of symbols in this mathematical model is adopted from [17]. All trips depart from and end at the "hub".

The sets are defined as:

- N: Set of all nodes. The hub is represented as two separate nodes: the $start\_hub$ node and the end_hub node.
- V: Subset of N containing only customer nodes.

The decision variables are defined as:

- $X_{i,j} \quad \forall i \in N, \forall j \in N$: Binary variable; 1 if the arc from node i to node j is included in the trip, 0 otherwise.
- $T_j \quad \forall j \in N$: Arrival time at node j.

The optimal sequence of customers is determined based on the minimization of the total trip duration. However, this objective does not define the optimal departure time from the hub. A trip can depart on several moments in time resulting in the same total trip duration. In this case, the departure time is determined based on the optimization of time-window overlap; The model prefers solutions in which the planned moments of delivery lie in the middle of the delivery time windows. The delivery time window is constrained by the order window. In order to accomplish this objective, a planned arrival time within $0.5RDW$ minutes from the start or end of the order window is penalized by means of the below decision variables, where RDW is the size of a regular delivery window.

- $SEA_j \quad \forall j \in V$: Early arrival with respect to the start of the customer order time window + $0.5RDW$.
- $SLA_j \quad \forall j \in V$: Late arrival with respect to the end of the customer order time window - $0.5RDW$.

The dominant objective of this sequencing model is to minimize the duration of a trip. Subordinately it minimizes the number of deliveries with a planned arrival time within $0.5RDW$ from the bounds of the order window. The latter results in a higher probability of arriving on time at the customers.

$$Minimize \quad M * (T_{end\_hub} - T_{start\_hub}) + \sum_{j \in V} (SEA_j^2 + SLA_j^2) \tag{1}$$

Subject to:

$$\sum_{j \in N} X_{i,j} = 1 \quad \forall i \in V \tag{2}$$

$$\sum_{i \in N} X_{i,j} = 1 \quad \forall j \in V \tag{3}$$

$$\sum_{i \in N} X_{i,start\_hub} = 0 \tag{4}$$

$$\sum_{j \in N} X_{start\_hub,j} = 1 \tag{5}$$

$$\sum_{i \in N} X_{i,end\_hub} = 1 \tag{6}$$

$$\sum_{j \in N} X_{end\_hub,j} = 0 \tag{7}$$

$$X_{i,i} = 0 \quad \forall i \in N \tag{8}$$

$$T_j \geq \sum_{i \in N} X_{i,j} * (T_i + ST_i + TT_{i,j}) \quad \forall j \in V \tag{9}$$

$$T_{start\_hub} \leq \sum_{\forall j \in V} X_{start\_hub,j} * (T_j - TT_{start\_hub,j}) \tag{10}$$

$$T_{end\_hub} \geq \sum_{\forall j \in V} X_{i,end\_hub} * (T_i + ST_i + TT_{i,end\_hub}) \tag{11}$$

$$PTWS_j \leq T_j \leq PTWE_j \quad \forall j \in V \tag{12}$$

$$T_j - OTWS_j + SEA_j \geq 0.5 * RDW \quad \forall j \in V \tag{13}$$

$$OTWE_j - T_j + SLA_j \geq 0.5 * RDW \quad \forall j \in V \tag{14}$$

Constraints 2 till 7 ensure spatial continuity. Time continuity is guaranteed by Constraints 9, 10 and 11. Here ST_i is the planned service time at customer i and $TT_{i,j}$ is the planned travel time from node i to node j. Constraint 12 ensures a planned arrival time after the start of a customer's planning window ($PTWS_j$) and before the end of a customer's planning window ($PTWE_j$). The planning window is a shortened version of the order window which prevents a planned arrival time in the closing minutes of a customer's order window. Constraints 13 and 14 calculate the seconds early (SEA_j) and seconds late (SLA_j) for the arrival time at each customer with respect to the optimal arrival window. Here $OTWS_j$ and $OTWE_j$ are, respectively, the start and end times of customer j's order time window.

B. Simulation-based sequencing model

The simulation-based sequencing model is a stochastic sequencing model. This model uses the benchmark sequencing model as its basis. It concerns a stochastic sequencing model because the optimal sequence of customers is selected based on a comparison of performance predictions of multiple solutions. A similar approach in the context of the Vehicle Routing Problem with Stochastic Demands (VRPSD) is taken by [18]. They search for the optimal number of trips to deliver uncertain amounts of goods to a set of customers. For different numbers of trips, they solve numerous iterations with customer demands generated according to Monte-Carlo simulation. The results for the different numbers of trips are compared, and this comparison yields an optimal number of trips.

This sequencing model completes *n* runs of the benchmark sequencing model. In each next run, the solution(s) found in the previous run(s) are eliminated. This results in a different solution for each run. The mathematical model is nearly identical to the mathematical model defined for the benchmark sequencing model, see Section III-A. In order to eliminate previously found solutions from the solution space, the following sets are added to the model.

- S: Set of all previously computed solutions
- $AS_s \quad \forall s \in S$: Set of arcs used in solution s. The arc from node i to node j is denoted as (i, j).

In order to arrive at different solutions Constraint 15 is added to the model. The *n* best solutions are looked for. In case less than *n* feasible solutions are available a smaller set of solutions is compared.

$$\sum_{(i,j) \in AS_s} X_{(i,j)} \leq len(N) - 1 \quad \forall s \in S \tag{15}$$

The best solution out of the set of good solutions is picked based on the a-priori simulated average on-time delivery rate over 25 iterations of the solution. This simulation is performed using the experimental method as explained in Section IV-A.

C. Heuristics-based sequencing model

In order to maximize the potential of real-time re-optimization it is interesting to investigate how the a-priori sequencing model can contribute to creating possibilities for effective re-optimization. In order to gain an insight in the correlation between the index of flexible deliveries in a trip and the effectiveness of re-optimization, an exploratory research is conducted. In this exploratory research, flexible deliveries are assigned to specific indices in a trip. Historic Picnic trips from one day of operations are used. Next, simulations of those trips are performed, including the re-optimization model (see Section III-D). By means of this approach an insight is gained in effective positions of flexible deliveries within a trip. The exploratory research makes use of trip instances which contain two one-hour order windows. This means that amongst the customers who are served in the same trip at most two different yet successive one-hour order windows were selected. The customers select a one-hour order window when placing an order. In case of a regular delivery that one-hour window is reduced to a smaller delivery time window in the morning of the delivery. If the customer opts for a flexible delivery, the delivery time window is the same as the order window.

The results from the exploratory research suggest the following:

- In case there is only one flexible delivery in an order window, the position of the flexible delivery within that window correlates with the performance on the KPIs.
- For the first order window in a trip, the possibility for re-optimization is most effective when the flexible delivery is positioned at the middle of that window.
- Regarding the second order window in a trip, the possibility for re-optimization is most effective when the flexible delivery is positioned at the start of that window.
- In case there are multiple flexible deliveries in the same order window, their positions within that window do not significantly affect the performance on the KPIs.

The results from the exploratory research point out that the positions of flexible deliveries within the trip sequence significantly affect the performance of the sequencing model. Therefore, the findings from the exploratory research are translated into the mathematical model of the heuristics-based sequencing model in order to maximize the re-optimization possibilities offered by flexible deliveries.

The majority of this mathematical model is identical to that formulated for the benchmark sequencing model (see Section III-A). Therefore only the adapted objective function, added set, decision variables and constraints are presented in this subsection.

The added set is defined as:

- OD: Set of all nodes of which the index within the trip has to be optimized. This amounts to a set consisting of a maximum of two nodes; one customer node in the first order window and one customer node in the second order window of a trip.

The added decision variables are defined as:

- $I_j \quad \forall j \in N$: Index of node j in the trip
- $OID_j \quad \forall j \in OD$: Deviation from the optimal index for node j

In the first place, the objective function minimizes the flexible deliveries' deviation from the optimal delivery index. Subordinately it minimizes the total trip duration. Lastly it minimizes the number of deliveries with a planned arrival time within $0.5RDW$ from the bounds of the order window, resulting in planned delivery times closer to the center of the corresponding communicated delivery time windows.

$$\begin{aligned} Minimize \quad & M * \sum_{j \in V} OID_j + M/1000 \\ & * (T_{end\_hub} - T_{start\_hub}) + \sum_{j \in V} (DWS_j^2 + DWE_j^2) \end{aligned} \tag{16}$$

Constraint 17 calculates the index of each customer in the trip. Constraints 19 and 20 determine the absolute value of the difference (OID_j) between the optimal index of node j (OI_j) and the index node j is assigned to in the solution (I_j). The optimal index of each node is determined based on the heuristics that were obtained by means of the exploratory research.

$$\sum_{j \in N} X_{i,j} * (I_j - I_i) = 1 \quad \forall i \in N \tag{17}$$

$$I_{start\_hub} = 0 \tag{18}$$

$$OID_j \leq OI_j - I_j \quad \forall j \in OD \tag{19}$$

$$OID_j \leq -OI_j + I_j \quad \forall j \in OD \tag{20}$$

D. Re-optimization model

This re-optimization model is ran upon departure from each customer in a trip, except the last two customers. Inputs are the currently followed trip planning (planned stem times, travel times and service times), delivery time windows and the trip progress. The output is the re-optimized customer sequence. Similar to [9], a route-oriented approach is used instead of a customer-by-customer approach. This means that during re-optimization the complete residual trip is re-optimized instead of merely the next customer in the trip.

The sets are defined as:

- R: Set containing the current node, all residual customer nodes and the end_hub node
- V: Set of residual customer nodes

The decision variables are defined as:

- $X_{i,j} \quad \forall i \in N, \forall j \in N$: Binary variable; 1 if the arc from node i to node j is included in the trip, 0 otherwise.
- $T_j \quad \forall i \in N$: Arrival time at node j

As opposed to the sequencing models which take into account the one hour order windows, the re-optimization model has to take into account smaller delivery time windows. When the progress of a trip falls behind the trip planning calculated a-priori, the use of hard time window constraints might result in an infeasible problem. For this reason, the re-optimization model makes use of soft time window constraints. Opposed to other researchers who penalize lateness or earliness in a linear fashion [19], [20], [21], this model penalizes earliness and lateness in a quadratic fashion. This approach ensures that a solution with two planned arrival times which are five minutes late, is selected instead of a solution with one planned arrival time which is ten minutes late. In order to optimize the chance that a driver arrives on time, the planned arrival time should be close to the middle of the delivery time window. The optimal arrival window is defined as following: [delivery time window start + SM, delivery time window end - SM]. SM represents the safety margin with respect to the boundaries of the delivery time window. In order to include the deviation of the planned arrival time from the optimal arrival window in the objective function the below decision variables are introduced.

- $L_j \quad \forall j \in V$: Lateness w.r.t. the optimal arrival window at customer j, ≥ 0
- $E_j \quad \forall j \in V$: Earliness w.r.t. the optimal arrival window at customer j, ≥ 0
- $D_j \quad \forall j \in V$: Absolute deviation of the planned arrival time from the optimal arrival window at customer j

The objective function minimizes the cumulative penalty due to planned arrival times outside of the optimal arrival windows at customers. Secondly, the duration of the residual trip is minimized. dev_opt describes the weight of the penalty of arriving outside of the optimal arrival window. The value of the parameter dev_opt depends on the application-specific trade-off between the two elements in the objective function.

$$Minimize \quad dev\_opt * \sum_{j \in V} D_j * D_j + T_{end\_hub} - T_{current\_node} \tag{21}$$

subject to:

$$\sum_{i \in R} X_{i,j} = 1 \quad \forall j \in V \tag{22}$$

$$\sum_{j \in R} X_{i,j} = 1 \quad \forall i \in V \tag{23}$$

$$\sum_{i \in V} X_{i,current\_node} = 0 \tag{24}$$

$$\sum_{j \in V} X_{current\_node,j} = 1 \tag{25}$$

$$\sum_{i \in V} X_{i,end\_hub} = 1 \tag{26}$$

$$\sum_{j \in V} X_{end\_hub,j} = 0 \tag{27}$$

$$X_{i,i} = 0 \quad \forall i \in R \tag{28}$$

$$T_j \geq \sum_{i \in R} X_{i,j} * (T_i + ST_i + TT_{i,j}) \quad \forall j \in V \tag{29}$$

$$T_{current\_node} = current\_time \tag{30}$$

$$T_{end\_hub} \geq \sum_{\forall j \in V} X_{i,end\_hub} * (T_i + ST_i + TT_{i,end\_hub}) \tag{31}$$

$$E_j \geq (TWS_j + SM) - T_j \quad \forall j \in V \tag{32}$$

$$L_j \geq T_j - (TWE_j - SM) \quad \forall j \in V \tag{33}$$

$$D_j \geq E_j + L_j \quad \forall j \in V \tag{34}$$

Constraints 22 till 27 ensure spatial continuity of a trip. Time continuity is ensured by constraints 29, 30 and 31. Where ST_i is the service time at customer i and $TT_{i,j}$ is the travel time from node i to node j. Constraints 32, 33 and 34 are used to determine the gap between the planned arrival time and the optimal arrival window. TWS_j and TWE_j respectively represent the start and end times of the delivery time window of customer j.

IV. Numerical experiments

The performances of the designed routing model configurations are evaluated by means of case-study test instances. The design of the computational experiment is presented in Subsection IV-A after which the results are discussed in Subsection IV-B.

A. Experimental method

In order to quantify the performance of the models as presented in Section III different routing model configurations are tested. An overview of the experimental method used to quantify the performance of those configurations is provided in Figure 2. In the experiments the size of the regular delivery windows (RDW) is set at 20 minutes, the weight of sub-optimal arrival times (dev_opt) is set at 100 and the safety margin (SM) is set at 10 minutes.

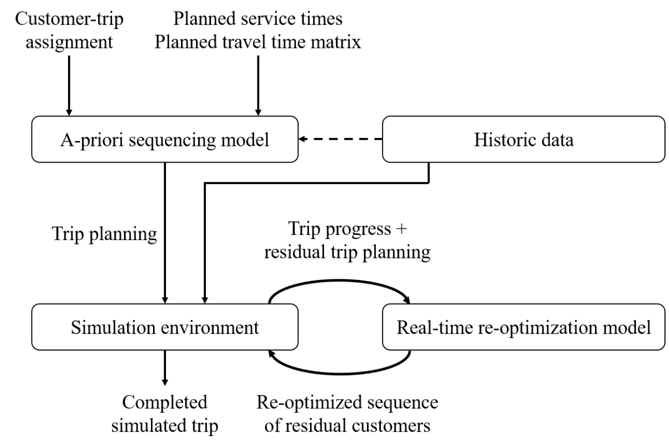

Fig. 2: Overview of the experimental method

Whether historic data is used in the a-priori sequencing model depends on the sequencing model that is tested. The simulation environment always uses historic data from e-grocer Picnic to sample realizations of travel times and service times. The completed simulated trips are analysed and their on-time delivery rate, rate of extreme lates and average time spent per delivery are used to quantify the performance of a configuration. All steps in the experimental method are executed on a personal computer with 16 Gb of RAM and a 1.90 GHz processor. Everything is programmed in Python 3.7. The a-priori sequencing model and the real-time re-optimization model call the Gurobi MILP solver to solve the mathematical problem and find the optimal sequence of customers.

The performance of four different routing model configurations is tested using customer-trip assignments from e-grocer Picnic. Two types of customer-trip assignments are investigated; for a small vehicle and a large vehicle. The large vehicle can carry 33% more groceries compared to the small vehicle. Figure 3 illustrates the vehicles sizes and compares them to a standard Toyota Prius. The experiment is performed for two different sizes of the flexible delivery time windows; 60 minutes and 75 minutes. An overview of the performed experiments is presented in Table II.

B. Experimental results

The results obtained using small vehicle customer-trip assignments and default flexible delivery windows are presented in Figure 4. Here the sequencing models are abbreviated as following: Benchmark (B), simulation-based (S), heuristics-based (H). Figure 4 shows that the simulation-based sequenc-

TABLE II: Overview of the tested configurations

Configurations		**Flexible delivery windows**		**Vehicle size**	
Sequencing model	Re-optimization model	60 min	75 min	Small	Large
Benchmark (B)	No	✓	✓	✓	✓
Simulation-based (S)	No	✓	✓	✓	✓
Simulation-based (S)	Yes	✓	✓	✓	✓
Heuristics-based (H)	Yes	✓	✓	✓	✓

Fig. 3: Scale of the vehicle sizes: small vehicle type (top), large vehicle type (bottom) and a reference car (middle)

ing model does not improve the on-time delivery rate of the last-mile distribution system significantly with respect to the benchmark configuration. In combination with the re-optimization model, it improves the on-time delivery rate marginally for fractions of flexible deliveries larger than 0.20. Regarding the fraction of extreme lates, no significant difference in performance is shown amongst the different configurations. The average time spent per delivery increases due to real-time re-optimization of trips. When looking at the configurations including the re-optimization model, it becomes clear that the heuristics-based sequencing model results in a larger increase in the average time spent per delivery than the simulation-based sequencing model.

The results for the different combinations of vehicle sizes and flexible delivery window sizes are presented in Tables III, IV and V for a fraction of flexible deliveries of 0.10. It can be concluded that the configurations including a re-optimization model outperform the static and deterministic benchmark configuration in terms of on-time delivery performance. However, re-optimization comes at the price of an increased average time spent per delivery. The simulation-based sequencing model proves to be effective without re-optimization for the large vehicle type. For the small vehicle type the simulation-based sequencing model needs the re-optimization model to significantly outperform the benchmark configuration. When the re-optimization model is included in the configuration, the simulation-based and heuristics-based sequencing models show similar performance in terms of late deliveries. The heuristics-based approach appears to outperform the simulation-based approach in terms of the rate of extreme lates, especially for the large vehicle. However, the configuration including the simulation-based sequencing model results in a 1% lower average time spent per delivery, no matter the vehicle size or the size of the flexible delivery windows.

Extension of the flexible delivery windows results in a significant improvement of the on-time performance of the configurations. However, the increased number of successful re-optimizations burdens the average time spent per delivery. The added value of extended flexible delivery windows is largest for the small vehicle type. When comparing the two types of vehicles studied, it becomes evident that the choice for the routing model configuration has a larger effect on the KPIs for a large vehicle than for a small vehicle.

V. Conclusions and future research directions

In conclusion, stochastic and dynamic routing models can significantly improve the on-time delivery performance of an e-grocer's last-mile distribution system. The concept of flexible deliveries was proven to successfully add a degree of freedom to the routing system which is required to make effective use of real-time re-optimization. The experimental results demonstrate that the use of a stochastic a-priori sequencing model improves the performance of the routing system on the KPIs as defined by an e-grocer. This sequencing model uses a-priori simulation to choose the best solution out of a set of good solutions. The effectiveness of the concept of flexible deliveries depends on the type of vehicles used, the fraction of flexible deliveries and the size of the flexible delivery windows. When compared to the static and deterministic benchmark configuration, the potential gain through the use of a stochastic and dynamic configuration is most profound for the large vehicle. When the fraction of flexible deliveries increases, the effects of re-optimization become more evident. When 10% of all deliveries is flexible, real-time re-optimization can significantly improve the performance of the last-mile distribution system. Extending the flexible delivery window results in further improvement of the on-time delivery performance of the routing system, however, it also increases the average time spent per delivery and thereby increases the operational costs for delivery.

In order to compare the performance of the concept of flexible deliveries to approaches studied by other researchers, the routing model configurations presented in this paper should be tested on benchmark test instances such as an adapted version of Solomon's test instances [22]. Alternatively, benchmark routing models could be tested on the e-grocer specific test instances used in this paper.

In this study the research scope is limited to the sequencing model and the re-optimization model. For the computational

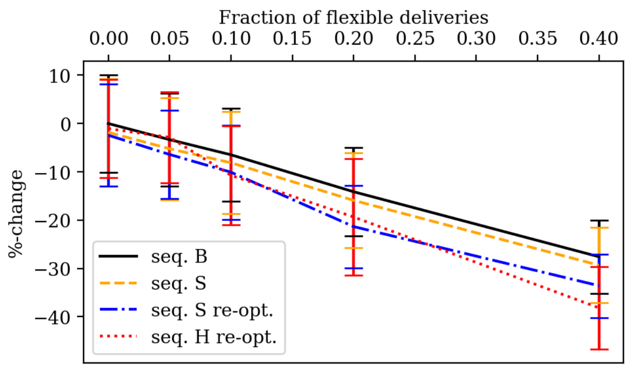

(a) Nr. of late deliveries

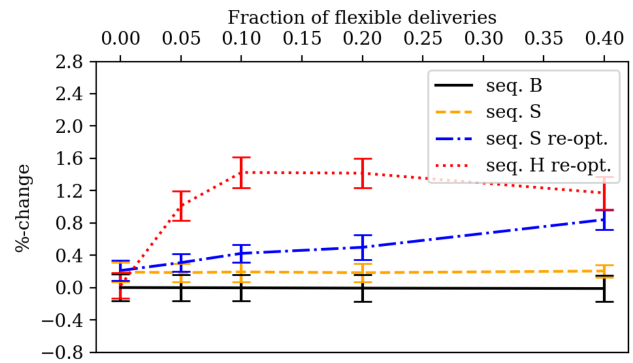

(b) Average time spent per delivery

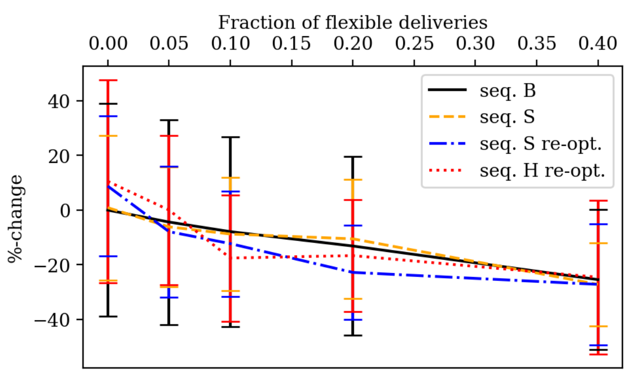

(c) Nr. of deliveries $\geq$ 15 min late

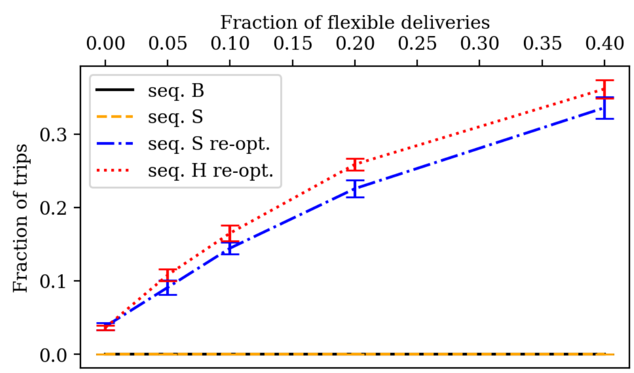

(d) Changed a-priori customer sequence

Fig. 4: Results for the small vehicle with default flexible delivery windows

TABLE III: Results for late deliveries

Late deliveries (fraction of flexible deliveries = 0.10)					
Configuration		Small vehicle		Large vehicle	
Sequencing model	Re-optimization model	Default FTW*	Extended FTW*	Default FTW*	Extended FTW*
Benchmark	No	0.0 ± 10 %	-3.1 ± 9.6 %	0.0 ± 9.9 %	-2.8 ± 9.3 %
Simulation-based	No	-1.8 ± 11 %	-4.7 ± 11 %	-7.8 ± 14 %	-11 ± 13 %
Simulation-based	Yes	-3.9 ± 10 %	-12 ± 10 %	-13 ± 13 %	-19 ± 13 %
Heuristics-based	Yes	-4.5 ± 11 %	-14 ± 11 %	-12 ± 10 %	-18 ± 9 %

TABLE IV: Results for extreme lates

Fraction of extreme lates (fraction of flexible deliveries = 0.10)					
Configuration		Small vehicle		Large vehicle	
Sequencing model	Re-optimization model	Default FTW*	Extended FTW*	Default FTW*	Extended FTW*
Benchmark	No	0 ± 38 %	-4 ± 37 %	0 ± 54 %	0 ± 54 %
Simulation-based	No	-1 ± 23 %	-2 ± 24 %	16 ± 45 %	11 ± 48 %
Simulation-based	Yes	-5 ± 20 %	-23 ± 25 %	4 ± 43 %	4 ± 47 %
Heuristics-based	Yes	-11 ± 25 %	-27 ± 25 %	-19 ± 50 %	-20 ± 54 %

TABLE V: Results for average time spent per delivery

Average time spent per delivery (fraction of flexible deliveries = 0.10)					
Configuration		Small vehicle		Large vehicle	
Sequencing model	Re-optimization model	Default FTW*	Extended FTW*	Default FTW*	Extended FTW*
Benchmark	No	0.0 ± 0.16 %	0.0 ± 0.16 %	0.0 ± 0.56 %	0.0 ± 0.56 %
Simulation-based	No	0.2 ± 0.12 %	0.2 ± 0.12 %	-0.3 ± 0.52 %	-0.3 ± 0.52 %
Simulation-based	Yes	0.4 ± 0.11 %	0.9 ± 0.15 %	-0.1 ± 0.46 %	0.4 ± 0.48 %
Heuristics-based	Yes	1.4 ± 0.20 %	1.9 ± 0.2 %	0.9 ± 0.44 %	1.4 ± 0.45 %

experiments a simplistic customer-trip assignment model was used which is not optimized for use with the concept of flexible deliveries. Therefore it would be interesting to investigate the effects of a customer-trip assignment model tailored to perform well in combination with this concept. A possible research topic would be the design of a customer-trip assignment model which spreads the flexible deliveries over different trips more evenly.

The heuristics-based sequencing model is designed around the findings of an exploratory research which was performed using the small vehicle type only. The optimal positions of flexible deliveries in a trip could be different for the two vehicle types. Therefore, another exploratory research should be performed that uses large vehicle customer-trip assignments to optimize the performance of the heuristics-based sequencing model for that vehicle type.

References

[1] N. Saunders and GlobalData, "Online grocery: Lessons from history," May 2019. [Online]. Available: https://www.onespace.com/blog/2018/10/online-grocery-lessons-history/

[2] N. Saunders, "Online grocery; food shopping statistics," May 2019. [Online]. Available: https://www.onespace.com/blog/2018/08/online-grocery-food-shopping-statistics/

[3] M. Mkansi and A. L. Nsakanda, "Leveraging the physical network of stores in e-grocery order fulfilment for sustainable competitive advantage," *Research in Transportation Economics*, p. 100786, 2019.

[4] J. van der Weerd, T. Leusden, and H. Tellingen, *Online supermarkt omzet op waarde geschat*. Supermarkt en Ruimte, 2019.

[5] M. Joerss, J. Schröder, F. Neuhaus, C. Klink, and F. Mann, "Parcel delivery: The future of last mile," *McKinsey & Company*, 2016. [Online]. Available: https://www.mckinsey.com/ /media/mckinsey/industries/travel transport and logistics/our insights/how customer demands are reshaping last mile delivery/parcel_delivery_the_future_of_last_mile.ashx

[6] N. Galante, E. G. López, and S. Monroe, "The Future of Online Grocery in Europe," *McKinsey & Company*, pp. 22–31, 2013.

[7] M. Wilson-Jeanselme and J. Reynolds, "Understanding shoppers' expectations of online grocery retailing," *International Journal of Retail and Distribution Management*, vol. 34, no. 7, pp. 529–540, 2006.

[8] J. Oyola, H. Arntzen, and D. L. Woodruff, "The stochastic vehicle routing problem, a literature review, part I: models," *EURO Journal on Transportation and Logistics*, vol. 7, no. 3, pp. 193–221, 2018.

[9] M. W. Ulmer, J. C. Goodson, D. C. Mattfeld, and B. W. Thomas, "Dynamic Vehicle Routing: Literature Review and Modeling Framework," no. June 2018, 2017. [Online]. Available: http://www.slu.edu/ goodson/papers/RouteBasedMDPs.pdf

[10] D. Li, C. Antoniou, H. Jiang, Q. Xie, W. Shen, and W. Han, "The Value of Prepositioning in Smartphone-Based Vanpool Services under Stochastic Requests and Time-Dependent Travel Times," *Transportation Research Record*, vol. 2673, no. 2, pp. 26–37, 2019.

[11] A. Toriello, W. B. Haskell, and M. Poremba, "A dynamic traveling salesman problem with stochastic arc costs," *Operations Research*, vol. 62, no. 5, pp. 1107–1125, 2014.

[12] C. Archetti, D. Feillet, A. Mor, and M. G. Speranza, "Dynamic traveling salesman problem with stochastic release dates," *European Journal of Operational Research*, vol. 280, no. 3, pp. 832–844, 2020.

[13] L. Chen, W. C. Chiang, R. Russell, J. Chen, and D. Sun, "The Probabilistic Vehicle Routing Problem with Service Guarantees," *Transportation Research Part E: Logistics and Transportation Review*, vol. 111, no. August 2017, pp. 149–164, 2018.

[14] K. K. Ng, C. K. Lee, S. Z. Zhang, K. Wu, and W. Ho, "A multiple colonies artificial bee colony algorithm for a capacitated vehicle routing problem and re-routing strategies under time-dependent traffic congestion," *Computers and Industrial Engineering*, vol. 109, pp. 151–168, 2017.

[15] F. Errico, G. Desaulniers, M. Gendreau, W. Rei, and L.-M. Rousseau, "A priori optimization with recourse for the vehicle routing problem with hard time windows and stochastic service times," *European Journal of Operational Research*, vol. 249, no. 1, pp. 55–66, 2016.

[16] E. Taniguchi and H. Shimamoto, "Intelligent transportation system based dynamic vehicle routing and scheduling with variable travel times," *Transportation Research Part C: Emerging Technologies*, vol. 12, no. 3-4 SPEC.ISS., pp. 235–250, 2004.

[17] Y. J. D. Dumans, M. M. Solomon, and F. Soumis, "Handbooks in operations research and management science, vol. 9. Finance," *European Journal of Operational Research*, vol. 8, no. 3, pp. 35–139, 1995.

[18] A. A. Juan, J. Faulin, S. E. Grasman, D. Riera, J. Marull, and C. A. Mendez, "Using safety stocks and simulation to solve the vehicle routing problem with stochastic demands," *Transportation Research Part C: Emerging Technologies*, vol. 19, no. 5, pp. 751–765, 2011.

[19] N. Balakrishnan, "Simple heuristics for the vehicle routeing problem with soft time windows," *Journal of the Operational Research Society*, vol. 44, no. 3, pp. 279–287, 1993.

[20] D. Taş, N. Dellaert, T. Van Woensel, and T. De Kok, "Vehicle routing problem with stochastic travel times including soft time windows and service costs," *Computers and Operations Research*, vol. 40, no. 1, pp. 214–224, 2013.

[21] H. I. Calvete, C. Galé, M. J. Oliveros, and B. Sánchez-Valverde, "A goal programming approach to vehicle routing problems with soft time windows," *European Journal of Operational Research*, vol. 177, no. 3, pp. 1720–1733, 2007.

[22] M. M. Solomon, "Algorithms for the Vehicle Routing and Scheduling Problems With Time Window Constraints." *Operations Research*, vol. 35, no. 2, pp. 254–265, 1987.

Bibliography

Yossiri Adulyasak and Patrick Jaillet. Models and algorithms for stochastic and robust vehicle routing with deadlines. *Transportation Science*, 50(2):608–626, 2016. ISSN 15265447. doi: 10.1287/trsc.2014.0581.

Albert Heijn. Wat betaal ik voor het laten thuisbezorgen of ophalen van boodschappen bij een pickup point?, 2020. URL https://www.ah.nl/klantenservice/vraag/kA00Y000000Mtx8SAC.

Claudia Archetti, Dominique Feillet, Andrea Mor, and Maria G. Speranza. Dynamic traveling salesman problem with stochastic release dates. *European Journal of Operational Research*, 280(3):832–844, 2020. ISSN 03772217. doi: 10.1016/j.ejor.2019.07.062.

Philippe Augerat. *Approche polyèdrale du problème de tournées de véhicules. (Polyhedral approach of the vehicle routing problem)*. PhD Thesis, 1995.

Nagraj Balakrishnan. Simple heuristics for the vehicle routeing problem with soft time windows. *Journal of the Operational Research Society*, 44(3):279–287, 1993. ISSN 14769360. doi: 10.1057/jors.1993.53.

Vahid Baradaran, Amir Shafaei, and Amir Hossein Hosseinian. Stochastic vehicle routing problem with heterogeneous vehicles and multiple prioritized time windows: Mathematical modeling and solution approach. *Computers and Industrial Engineering*, 131(September 2018):187–199, 2019. ISSN 03608352. doi: 10.1016/j.cie.2019.03.047.

Dimitris Bertsimas, Patrick Jaillet, and Sébastien Martin. Online vehicle routing: The edge of optimization in large-scale applications. *Operations Research*, 67(1):143–162, 2019. ISSN 15265463. doi: 10.1287/opre.2018.1763.

Yossi Borenstein, Nazaraf Shah, Edward Tsang, Raphael Dorne, Abdullah Alsheddy, and Christos Voudouris. On the partitioning of dynamic workforce scheduling problems. *Journal of Scheduling*, 13(4):411–425, 2010. ISSN 10946136. doi: 10.1007/s10951-009-0152-6.

Kris Braekers, Katrien Ramaekers, and Inneke Van Nieuwenhuyse. The vehicle routing problem: State of the art classification and review. *Computers and Industrial Engineering*, 99:300–313, 2016. ISSN 03608352. doi: 10.1016/j.cie.2015.12.007.

Herminia I. Calvete, Carmen Galé, María José Oliveros, and Belén Sánchez-Valverde. A goal programming approach to vehicle routing problems with soft time windows. *European Journal of Operational Research*, 177(3):1720–1733, 2007. ISSN 03772217. doi: 10.1016/j.ejor.2005.10.010.

Wujun Cao and Wenshui Yang. A survey of vehicle routing problem. *MATEC Web of Conferences*, 100:1–6, 2017. ISSN 2261236X. doi: 10.1051/matecconf/201710001006.

Abraham Charnes and William W. Cooper. Chance-Constrained Programming. *Management Science*, 6(1):73–79, 1959. ISSN 0025-1909. doi: 10.1287/mnsc.6.1.73.

Lijian Chen, Wen Chyuan Chiang, Robert Russell, Jun Chen, and Dengfeng Sun. The Probabilistic Vehicle Routing Problem with Service Guarantees. *Transportation Research Part E: Logistics and Transportation Review*, 111(August 2017):149–164, 2018. ISSN 13665545. doi: 10.1016/j.tre.2018.01.012.

Xin Chen. An improved efficient algorithm for time dependent vehicle routing. *Operations and Supply Chain Management*, 11(2):55–65, 2018. ISSN 25799363. doi: 10.31387/oscm0310201.

Zhi Long Chen and Hang Xu. Dynamic column generation for dynamic vehicle routing with time windows. *Transportation Science*, 40(1):74–88, 2006. ISSN 15265447. doi: 10.1287/trsc.1050.0133.

Harold Creusen, Arno Meijer, Gijsbert Zwart, and Henry Wiel. *Static efficiency in Dutch supermarket chain*. CPB Netherlands Bureau for Economic Policy Analysis, 2008. doi: 978-90-5833-362-9.

Yvan J. D. Dumans, Marius M. Solomon, and François Soumis. Handbooks in operations research and management science, vol. 9. Finance. *European Journal of Operational Research*, 8(3):35–139, 1995. ISSN 03772217. doi: 10.1016/s0377-2217(97)83331-0.

Jan Fabian Ehmke, Ann Melissa Campbell, and Timothy L. Urban. Ensuring service levels in routing problems with time windows and stochastic travel times. *European Journal of Operational Research*, 240(2):539–550, 2015. ISSN 03772217. doi: 10.1016/j.ejor.2014.06.045.

Fausto Errico, Guy Desaulniers, Michel Gendreau, Walter Rei, and Louis-Martin Rousseau. A priori optimization with recourse for the vehicle routing problem with hard time windows and stochastic service times. *European Journal of Operational Research*, 249(1):55–66, 2016. ISSN 03772217. doi: 10.1016/j.ejor.2015.07.027.

Fausto Errico, Guy Desaulniers, Michel Gendreau, Walter Rei, and Louis-Martin Rousseau. The vehicle routing problem with hard time windows and stochastic service times. *EURO Journal on Transportation and Logistics*, 7(3):223–251, 2018. ISSN 21924384. doi: 10.1007/s13676-016-0101-4.

Nicolò Galante, Enrique García López, and Sarah Monroe. The Future of Online Grocery in Europe. *McKinsey & Company*, pages 22–31, 2013. ISSN 0034-3617. doi: 10.1016/S0034-3617(11)70034-0.

Michel Gendreau, Ola Jabali, and Walter Rei. 50th Anniversary invited article future research directions in stochastic vehicle routing. *Transportation Science*, 50(4):1163–1173, 2016. ISSN 15265447. doi: 10.1287/trsc.2016.0709.

Gianpaolo Ghiani, Emanuele Manni, Antonella Quaranta, and Chefi Triki. Anticipatory algorithms for same-day courier dispatching. *Transportation Research Part E: Logistics and Transportation Review*, 45(1):96–106, 2009. ISSN 13665545. doi: 10.1016/j.tre.2008.08.003.

Pasquale Grippa, Doris A. Behrens, Friederike Wall, and Christian Bettstetter. Drone delivery systems: job assignment and dimensioning. *Autonomous Robots*, 43(2):261–274, 2019. ISSN 15737527. doi: 10.1007/s10514-018-9768-8.

Patrick Oliver Groß, Michael Geisinger, Jan Fabian Ehmke, and Dirk Christian Mattfeld. Interval Travel Times for More Reliable Routing in City Logistics. *Transportation Research Procedia*, 12(June 2015): 239–251, 2016. ISSN 23521465. doi: 10.1016/j.trpro.2016.02.062.

Ali R. Güner, Alper Murat, and Ratna Babu Chinnam. Dynamic routing for milk-run tours with time windows in stochastic time-dependent networks. *Transportation Research Part E: Logistics and Transportation Review*, 97:251–267, 2017. ISSN 13665545. doi: 10.1016/j.tre.2016.10.014.

Shuihua Han, Ling Zhao, Kui Chen, Zong wei Luo, and Deepa Mishra. Appointment scheduling and routing optimization of attended home delivery system with random customer behavior. *European Journal of Operational Research*, 262(3):966–980, 2017. ISSN 03772217. doi: 10.1016/j.ejor.2017.03.060.

Sin C. Ho, W. Y. Szeto, Yong Hong Kuo, Janny M.Y. Leung, Matthew Petering, and Terence W.H. Tou. A survey of dial-a-ride problems: Literature review and recent developments. *Transportation Research Part B: Methodological*, 111:395–421, 2018. ISSN 01912615. doi: 10.1016/j.trb.2018.02.001.

Patrick Jaillet, Jin Qi, and Melvyn Sim. Routing optimization under uncertainty. *Operations Research*, 64(1):186–200, 2016. ISSN 15265463. doi: 10.1287/opre.2015.1462.

Chenlu Ji, Anne Liret, Gilbert Owusu, Jiyin Liu, Raphael Dorne, and Rupal Rana. Service scheduling to minimise the risk of missing appointments. *Proceedings of Computing Conference 2017*, 2018-Janua (July):971–980, 2018. doi: 10.1109/SAI.2017.8252211.

Martin Joerss, Jürgen Schröder, Florian Neuhaus, Christoph Klink, and Florian Mann. Parcel delivery: The future of last mile. *McKinsey & Company*, 2016. URL https://www.mckinsey.com/{~}/media/mckinsey/industries/traveltransportandlogistics/ourinsights/howcustomerdemandsarereshapinglastmiledelivery/parcel{_}delivery{_}the{_}future{_}of{_}last{_}mile.ashx.

Angel A. Juan, Javier Faulin, Scott E. Grasman, Daniel Riera, Joan Marull, and Carlos A. Mendez. Using safety stocks and simulation to solve the vehicle routing problem with stochastic demands. *Transportation Research Part C: Emerging Technologies*, 19(5):751–765, 2011. ISSN 0968090X. doi: 10.1016/j.trc.2010.09.007.

Jumbo. Thuisbezorgen, 2020. URL https://www.jumbo.com/service/online-bestellen/thuisbezorgen.

Amela Karahasanović, Kristoffer T. Holm, and Aria Nejad. Design for trust - online grocery shopping. pages 239–243. IADIS press, 2017. ISBN 978-989-8533-64-7.

Phillip Kilby, Patrick Prosser, and Paul Shaw. Dynamic {VRP}s: A Study of Scenarios. (APES-06-1998), 1998.

Chi Guhn Lee, Marina A. Epelman, Chelsea C. White, and Yavuz A. Bozer. A shortest path approach to the multiple-vehicle routing problem with split pick-ups. *Transportation Research Part B: Methodological*, 40(4):265–284, 2006. ISSN 01912615. doi: 10.1016/j.trb.2004.11.004.

Hongtao Lei, Gilbert Laporte, and Bo Guo. The vehicle routing problem with stochastic demands and split deliveries. *Infor*, 50(2):59–71, 2012. ISSN 03155986. doi: 10.3138/infor.50.2.059.

Donghui Li, Constantinos Antoniou, Hai Jiang, Qianyan Xie, Wei Shen, and Weijian Han. The Value of Prepositioning in Smartphone-Based Vanpool Services under Stochastic Requests and Time-Dependent Travel Times. *Transportation Research Record*, 2673(2):26–37, 2019. ISSN 21694052. doi: 10.1177/0361198118822815.

Xiao Liang, Gonçalo Homem de Almeida Correia, Kun An, and Bart van Arem. Automated taxis' dial-a-ride problem with ride-sharing considering congestion-based dynamic travel times. *Transportation Research Part C: Emerging Technologies*, 112(March 2019):260–281, 2020. ISSN 0968090X. doi: 10.1016/j.trc.2020.01.024.

logistiek.nl. Albert heijn breidt hubnetwerk voor thuisbezorging uit, Sep 2016. URL https://www.logistiek.nl/supply-chain/nieuws/2016/09/albert-heijn-breidt-hubnetwerk-voor-thuisbezorging-uit-101147383.

Marcia Mkansi and Aaron Luntala Nsakanda. Leveraging the physical network of stores in e-grocery order fulfilment for sustainable competitive advantage. *Research in Transportation Economics*, page 100786, 2019. ISSN 07398859. doi: 10.1016/j.retrec.2019.100786.

Marcia Mkansi, Chukuakadibia Eresia-Eke, and Oyetola Emmanuel-Ebikake. E-grocery challenges and remedies: Global market leaders perspective. *Cogent Business and Management*, 5(1), 2018. ISSN 23311975. doi: 10.1080/23311975.2018.1459338.

Masahiro Nakagawa, Hiroshi Hasegawa, and Ken Ichi Sato. Dynamic path provisioning and disruption-free reoptimization algorithms for bandwidth on-demand services considering fairness. *IEICE Transactions on Communications*, E100B(4):536–547, 2017. ISSN 17451345. doi: 10.1587/transcom.2016EBP3202.

Kam K.H. Ng, Carman K.M. Lee, Shuzhu Z. Zhang, Kan Wu, and William Ho. A multiple colonies artificial bee colony algorithm for a capacitated vehicle routing problem and re-routing strategies under time-dependent traffic congestion. *Computers and Industrial Engineering*, 109:151–168, 2017. ISSN 03608352. doi: 10.1016/j.cie.2017.05.004.

Viet A. Nguyen, Jun Jiang, Kien M. Ng, and Kwong M. Teo. Satisficing measure approach for vehicle routing problem with time windows under uncertainty. *European Journal of Operational Research*, 248(2):404–414, 2016. ISSN 03772217. doi: 10.1016/j.ejor.2015.07.041.

Ocado. How much does delivery cost?, 2020. URL https://help.ocado.com/hc/en-us/articles/360017551912-How-much-does-delivery-cost-.

Jorge Oyola, Halvard Arntzen, and David L. Woodruff. The stochastic vehicle routing problem, a literature review, Part II: solution methods. *EURO Journal on Transportation and Logistics*, 6(4):349–388, 2017. ISSN 21924384. doi: 10.1007/s13676-016-0099-7.

Jorge Oyola, Halvard Arntzen, and David L. Woodruff. The stochastic vehicle routing problem, a literature review, part I: models. *EURO Journal on Transportation and Logistics*, 7(3):193–221, 2018. ISSN 21924384. doi: 10.1007/s13676-016-0100-5.

Paazl. The 5 key trends in e-commerce delivery for 2018, May 2018. URL https://www.paazl.com/blog/the-5-key-trends-in-e-commerce-delivery-for-2018/.

Dimitrios G. Pandelis, Constantinos C. Karamatsoukis, and Epaminondas G. Kyriakidis. Finite and infinite-horizon single vehicle routing problems with a predefined customer sequence and pickup and delivery. *European Journal of Operational Research*, 231(3):577–586, 2013. ISSN 03772217. doi: 10.1016/j.ejor.2013.05.050.

Ramesh Ramasamy Pandi, Song Guang Ho, Sarat Chandra Nagavarapu, and Justin Dauwels. Adaptive algorithm for dial-a-ride problem with vehicle breakdown. *2019 18th European Control Conference, ECC 2019*, pages 2682–2688, 2019. doi: 10.23919/ECC.2019.8795619.

Picnic. Veelgestelde vragen: Picnic online supermarkt, 2020. URL https://picnic.app/nl/veelgestelde-vragen.

Victor Pillac, Michel Gendreau, Christelle Guéret, and Andrés L. Medaglia. A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225(1):1–11, 2013. ISSN 03772217. doi: 10.1016/j.ejor.2012.08.015.

Harilaos N. Psaraftis, Min Wen, and Christos A. Kontovas. Dynamic Vehicle Routing Problems: Three Decades and Counting. *Wiley Online Library*, 2015. ISSN 1097-0037. doi: 10.1002/net.

Rabobank. Supermarkten, Dec 2019. URL https://www.rabobank.nl/bedrijven/cijfers-en-trends/food/supermarkten/.

Elizabeth Rendón-Vélez, Imre Horváth, and Wilhelm Frederik Van Der Vegte. Motivating subjects to drive in haste using time pressure in a simulated environment. *International Journal of Computer Aided Engineering and Technology*, 8(1-2):99–124, 2016. ISSN 17572665. doi: 10.1504/IJCAET.2016.073264.

Briseida Sarasola, Karl F. Doerner, Verena Schmid, and Enrique Alba. Variable neighborhood search for the stochastic and dynamic vehicle routing problem. *Annals of Operations Research*, 236(2):425–461, 2016. ISSN 15729338. doi: 10.1007/s10479-015-1949-7.

Neil Saunders and GlobalData. Online grocery: Lessons from history, May 2019a. URL https://www.onespace.com/blog/2018/10/online-grocery-lessons-history/.

Neil Saunders and GlobalData. Online grocery; food shopping statistics, May 2019b. URL https://www.onespace.com/blog/2018/08/online-grocery-food-shopping-statistics/.

Michael Schilde, Karl F. Doerner, and Richard F. Hartl. Integrating stochastic time-dependent travel speed in solution methods for the dynamic dial-A-ride problem. *European Journal of Operational Research*, 238(1):18–30, 2014. ISSN 03772217. doi: 10.1016/j.ejor.2014.03.005.

Marius M. Solomon. Algorithms for the Vehicle Routing and Scheduling Problems With Time Window Constraints. *Operations Research*, 35(2):254–265, 1987. ISSN 0030364X. doi: 10.1287/opre.35.2.254.

Duygu Taş, Nico Dellaert, Tom Van Woensel, and Ton De Kok. Vehicle routing problem with stochastic travel times including soft time windows and service costs. *Computers and Operations Research*, 40(1):214–224, 2013. ISSN 03050548. doi: 10.1016/j.cor.2012.06.008.

Eiichi Taniguchi and Hiroshi Shimamoto. Intelligent transportation system based dynamic vehicle routing and scheduling with variable travel times. *Transportation Research Part C: Emerging Technologies*, 12(3-4 SPEC.ISS.):235–250, 2004. ISSN 0968090X. doi: 10.1016/j.trc.2004.07.007.

Alejandro Toriello, William B. Haskell, and Michael Poremba. A dynamic traveling salesman problem with stochastic arc costs. *Operations Research*, 62(5):1107–1125, 2014. ISSN 15265463. doi: 10.1287/opre.2014.1301.

Marlin W. Ulmer. Horizontal combinations of online and offline approximate dynamic programming for stochastic dynamic vehicle routing. *Central European Journal of Operations Research*, 28(1): 279–308, 2020. ISSN 16139178. doi: 10.1007/s10100-018-0588-x.

Marlin W Ulmer, Justin C Goodson, Dirk C Mattfeld, and Barrett W Thomas. Dynamic Vehicle Routing: Literature Review and Modeling Framework. (June 2018), 2017. URL http://www.slu.edu/{~}goodson/papers/RouteBasedMDPs.pdf.

Jeroen van der Weerd, Tjerk Leusden, and Hans Tellingen. *Online supermarkt omzet op waarde geschat*. Supermarkt en Ruimte, 2019.

Bapi Raju Vangipurapu, Rambabu Govada, and Narayana Rao Kandukuri. Development Of A New Genetic Algorithm For Solving Capacitated Vehicle Routing Problem In Short Time. *International Journal of Innovative Technology and Exploring Engineering*, 8(11):903–907, 2019. doi: 10.35940/ijitee.k1545.0981119.

Pradeep Varakantham, Akshat Kumar, Hoong Chuin Lau, and William Yeoh. Risk-sensitive stochastic orienteering problems for trip optimization in urban environments. *ACM Transactions on Intelligent Systems and Technology*, 9(3), 2018. ISSN 21576912. doi: 10.1145/3080575.

Neža Vodopivec and Elise Miller-Hooks. An optimal stopping approach to managing travel-time uncertainty for time-sensitive customer pickup. *Transportation Research Part B: Methodological*, 102: 22–37, 2017. ISSN 01912615. doi: 10.1016/j.trb.2017.04.017.

Walmart. Walmart+, 2020. URL https://grocery.walmart.com/delivery-unlimited/landing.

Mirko Warschun, Werner Delfmann, and Sascha Albers. A fresh look at online grocery, 2012. URL https://www.de.kearney.com/consumer-goods/article?%2Fa%2Fa-fresh-look-at-online-grocery.

Muriel Wilson-Jeanselme and Jonathan Reynolds. Understanding shoppers' expectations of online grocery retailing. *International Journal of Retail and Distribution Management*, 34(7):529–540, 2006. ISSN 09590552. doi: 10.1108/09590550610673608.

Yu Wu, Bo Zeng, and Siming Huang. A Dynamic Strategy for Home Pick-Up Service with Uncertain Customer Requests and Its Implementation. *Sustainability*, 11(7):2060, 2019. doi: 10.3390/su11072060.

Baris Yildiz and Martin Savelsbergh. Pricing for delivery time flexibility. *Transportation Research Part B: Methodological*, 133:230–256, 2020. ISSN 01912615. doi: 10.1016/j.trb.2020.01.004.

Guodong Yu and Yu Yang. Dynamic routing with real-time traffic information. *Operational Research*, 19(4):1033–1058, 2019. ISSN 18661505. doi: 10.1007/s12351-017-0314-9.

Fuli Zhou, Yandong He, and Lin Zhou. Last Mile Delivery with Stochastic Travel Times Considering Dual Services. *IEEE Access*, 7:159013–159021, 2019. ISSN 21693536. doi: 10.1109/ACCESS.2019.2950442.