Technische Universität München

Lehrstuhl für Logistik und Supply Chain Management

Rebalancing in Shared Mobility Systems

Competition, Feature-Based Mode Selection and Technology Choice

Layla Martin, M.Sc.

Vollständiger Abdruck der von der Fakultät für Wirtschaftswissenschaften der Technischen Universität München zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaften (Dr. rer. pol.) genehmigten Dissertation.

Vorsitzende:		Prof. Dr. Gudrun Kiesmüller
Prüfer der Dissertation:	1.	Prof. Dr. Stefan Minner
	2.	Prof. Dr. Richard Hartl Universität Wien, Österreich

Die Dissertation wurde am 17.09.2020 bei der Technischen Universität München eingereicht und durch die Fakultät für Wirtschaftswissenschaften am 15.11.2020 angenommen.

Acknowledgements

I would like to thank everyone who supported me throughout my doctoral studies.

First and foremost, I am sincerely grateful to my supervisor Professor Stefan Minner for continuously supporting me and giving me constructive feedback. His insights and knowledge are invaluable to the progress of this thesis. I would also like to thank him for encouraging and supporting me to spend five months at Stanford University. I also want to thank Professor Marco Pavone for inviting me to Stanford, and the fruitful collaboration which resulted in a chapter of this thesis. I would like to express my gratitude to Professor Richard Hartl for being a member of my examination committee, and Professor Gudrun Kiesmüller for being the chairwoman of the committee.

Further, I would like to thank Professor Andreas S. Schulz and Dr. Diogo Poças for our joint work on a chapter of this thesis. I am grateful to Professor Maximilian Schiffer for his collaboration on a chapter of this thesis, connecting me with Professor Marco Pavone leading to a great stay abroad, and his continuous mentoring on (paths into) academia. I am also grateful to my collaborators Professor M. Grazia Speranza, Michael Wittmann and Xinyu Li, even though these projects are not part of this thesis.

I want to thank all current and former colleagues at the chair of Logistics and Supply Chain Management: Yuka Akasaka, Dr. Szymon Albiński, Tobias Crönert, Jun.-Prof. Dr. Pirmin Fontaine, Emanuel Herrmann, Michael Keilhacker, Dr. Miray Közen, Eunji Lee, Sebastian Malicki, Dr. Christian Mandl, Melika Mohsenizadehkamou, Vinh Ngo, Santiago Nieto-Isaza, Thitinan "Kate" Pholsook, Dr. Dennis Prak, Moritz Rettinger, Dr. Patricia Rogetzer, Caroline Spieckermann, Josef Svoboda and Francesco Zangaro. I would also like to thank my colleagues from the RTG Advanced Optimization in a Networked Economy: Ramin Barzanji, Eleni Batziou, Carolin Bauerhenne, Giacomo Dall'Olio, Alexander Eckl, Alexandre Forel, Anja Kirschbaum, Stefan Kober, Marilena Leichter, Richard Littmann, Alexandros Tsigonias-Dimitriadis and Donghao Zhu. Thank you for all the research discussions and also for the after-work activities. I would also like to thank Dr. Yoshimi von Felbert, Evelyn Gemkow and Dr. Isabel Koch for their continous support.

I am grateful to Deutsche Forschungsgemeinschaft (DFG) for funding my research as part of the GRK 2201 ("Advanced Optimization in a Networked Economy"). I would also like to thank TUM School of Management for funding my stay at Stanford and various conference visits via the Research Excellence Program and the Diversity Promotion Fund.

Finally, I would like to thank my parents and my sisters for their constant encouragement, their love and their support. Special thanks go to Uli, for his support, motivation and patience.

Abstract

Shared mobility operators such as carsharing and ride-hailing services commonly face the problem of unbalanced demand: The number of vehicles rented from a location does not necessarily equal the number of vehicles returned to this location. To counteract demand imbalances, operators *rebalance* their fleet, i.e., move vehicles from locations with an excess in supply to locations with an excess in demand. We investigate three extensions of the rebalancing problem: competition, modal selection, and autonomous vehicles. This thesis provides guidance for operators of shared mobility systems on how to increase their profitability by optimal rebalancing.

With an increasing competitiveness of the carsharing market, operators must consider the position where other operators currently have vehicles, as well as how the competitors rebalance their fleets. Existing models have so far ignored the aspect of competition in the optimization of rebalancing routes. We present a novel model called "Competitive Pickup and Delivery Orienteering Problem" (C-PDOP) that models competition in rebalancing. We solve the C-PDOP for Nash equilibria using two algorithms, Iterated Best Response and Potential Function Optimizer. The study reveals that operators can gain as much as 40% of their profit in a case study settled in Munich, Germany, due to considering competition. However, operators lose up to 12% of their profit in a Munich case study due to the presence of competition (compared to a merger).

Vehicles can be rebalanced by loading them onto a truck, or by driving them. In the latter case, staff must be rebalanced as well, i.e., workers have to give each other lifts, bike or use public transit to reach the next vehicle. We study which features drive the choice for either of the modes. Therefore, we build classifiers based on multiple linear regression, multinomial logistic regression, and decision trees. The accuracy of linear and logistic regression is very high (above 90%), and in the misclassified instances, operators incur only little additional cost (less than 10% over all misclassified instances). This novel approach reveals that the modal choice is driven by wages for workers, and vehicle costs (car and truck).

The advent of driverless vehicles will directly impact the shared mobility market,

and operators consider whether to procure driverless vehicles (to completely or partially replace human-driven vehicles). We study the technology choice and mix problem operators face, balancing investment costs with operational costs and contribution margins. The operational rebalancing decision is modeled as a semi-Markov decision problem and a closed queueing network. This thesis provides profound insights into the optimal fleet composition, and gains due to progressing automation: In ride-hailing systems (the customer is chauffeured), driverless vehicles will quickly replace the entire fleet, and allow operators to offer their service in new business regions. In carsharing systems (the customer drives herself), operators often benefit of mixed fleets, and will not always replace their entire fleet.

Contents

List of Tables

List of Figures

List of Abbreviations

AMoD	Autonomous Mobility-on-Demand
C-PDOP	Competitive Pickup and Delivery Orienteering Problem
Coop-PDOP	Coopetitive Pickup and Delivery Orienteering Problem
CRP	Carsharing Relocation Problem
CRP-B	Carsharing Relocation Problem with Bike
CRP-C	Carsharing Relocation Problem with Car
CRP-CD	Carsharing Relocation Problem with Car and Dedicated Helpers
CRP-CS	Carsharing Relocation Problem with Car and Shared Helpers
CRP-P	Carsharing Relocation Problem with Public Transit
CRP-T	Carsharing Relocation Problem with Truck
DLP	Dual Linear Program
FE	Full Enumeration
FLP	Fluid-Based Approximation Linear Program
FN	false negatives
FP	false positives
IBR	Iterated Best Response
M-CRP	Multi-Mode Carsharing Relocation Problem
M-PDOP	Monopoly Pickup and Delivery Orienteering Problem
MaaS	Mobility-as-a-Service
MVA	Mean Value Analysis
PDOP	Pickup and Delivery Orienteering Problem
PDP	Pickup and Delivery Problem
PFO	Potential Function Optimizer
QMO	Optimistic Quasi-Monopolistic Pickup and Delivery Orienteering Problem
QMP	Pessimistic Quasi-Monopolistic Pickup and Delivery Orienteering Problem
R^2	coefficient of determination

RRMSE	relative root-mean-squared error
SMDP	Semi-Markov Decision Process
VRP	Vehicle Routing Problem
W-PDOP	Welfare-Maximizing Pickup and Delivery Orienteering Problem

Chapter 1

Introduction

1.1 Motivation

Shared mobility or vehicle sharing schemes comprise different mobility concepts, including, but not limited to, carsharing and ridesharing, mainly with a focus on short distances in an urban context. These concepts become more important globally, as they allow to address some of the key challenges our society is currently facing due to increased mobility: congestion and emissions. As a consequence of implementing vehicle sharing, the total number of vehicles decreases, and oftentimes, carbon dioxide emissions can be reduced due to using newer and alternative fuel vehicles (Bellos et al., 2017; Firnkorn and Müller, 2011; Shaheen and Cohen, 2013). In the near future, driverless shared vehicles will further contribute to lowering emissions (Greenblatt and Saxena, 2015; PriceWaterhouseCoopers, 2017). Less parking spaces are required if shared vehicles replace privately owned ones (Boston Consulting Group, 2020). For many users, sharing vehicles is more cost-efficient than owning them (Baptista et al., 2014), in particular if vehicle procurement prices decrease over time (Ostrovsky and Schwarz, 2019). Due to the societal benefits, and the possibility to open a new market segment, new carsharing and ride-hailing services emerge frequently (Perboli et al., 2018). Thus, the shared mobility market is increasingly competitive, and everchanging (Kortum et al., 2016).

As Basciftci et al. (2020), Nair and Miller-Hooks (2014), and Shaheen et al. (2006) state, demand for carsharing services depends on the availability of the service: Customers only rely on the service if the probability of finding a vehicle nearby is high (similar trends have been reported for other modes of shared mobility, including bikesharing (Kabra et al., 2020)). If the availability is low, customers cannot reduce their expected travel time which is one of the most important drivers for adopting carsharing

(Schaefers, 2013). Operators can increase the fleet size to increase the availability of the service (Boyacı et al., 2015), but larger fleets incur higher fixed costs and can result in very low utilization levels which partially alleviate the benefits attributed to shared mobility. Thus, operators of one-way and free-floating shared mobility services have to rebalance vehicles. In all one-way and free-floating systems, randomness can cause imbalances in the vehicle distribution. Demand imbalance can be caused by random differences in supply and demand (Boyacı et al., 2015), by systematic spatial demand imbalances, e.g. due to integration with other modes (Wagner et al., 2016), and by spatio-temporal imbalances such as during the rush hour (Ampudia-Renuncio et al., 2020; Huang et al., 2018; Schmöller et al., 2015). Without interfering, vehicles agglomerate in "cold spots" while customers cannot be served in "hot spots" which incurs lost sales. Thus, carsharing operators either send employees to move vehicles, or incentivize users to adapt their travel patterns. The former is commonly referred to as "operator-based" rebalancing, while the latter is called "user-based" rebalancing. Vasconcelos et al. (2017) report that rebalancing increases the profitability of carsharing systems, and in many cases is necessary to even reach profitability. Even if user-based rebalancing is employed, operator-based rebalancing remains necessary: (i) User-based rebalancing introduces additional uncertainty, (ii) for some remote locations, user-based rebalancing is more expensive than operator-based rebalancing, and (iii) operator-based rebalancing can easily be integrated with cleaning and maintenance operations. Analogously, ridesharing operators either send empty vehicles directly to the destination (if they own the fleet and employ the drivers), or use dynamic pricing and remuneration schemes (often referred to as "surge pricing") to steer supply and demand. However, surge pricing alone is no sufficient measure to achieve balanced demand. For example, the expected waiting times for an Uber can be as high as 43 minutes in San Francisco and Manhattan, despite the availability of surge pricing (Chen et al., 2015). Consequentially, surge pricing alone cannot fully address spatial supply-demand imbalances. If the demand imbalance is not too high (Guda and Subramanian, 2019) and during off-peak hours (Ming et al., 2020), paying independent workers bonuses to move to a different location is beneficial over surge pricing with respect to operator profit and/or social welfare. Further, paying bonuses rather than employing surge pricing is risk averse behavior (Jiang et al., 2020).

In carsharing systems, operators usually pool multiple relocation operations, often during the course of the night when demand is low (Huang et al., 2018; Weikl and Bogenberger, 2015). Vehicles that shall be rebalanced are either driven by employees,

or loaded onto a truck. If vehicles are driven by operators, workers have to continue to the next vehicle, and use bikes or public transportation, walk, or hitch rides with colleagues in carsharing vehicles. In ride-hailing systems and Autonomous Mobility-on-Demand (AMoD) systems, the fleet is rebalanced continuously, since no clear pooling benefit emerges.

1.2 Contribution and Research Questions

Despite increasing competition in shared mobility markets, this aspect has been mostly ignored when optimizing a carsharing fleet (with exceptions of Albiński and Minner (2020) and Balac et al. (2019)). We investigate how a carsharing provider can adapt her routing and servicing decision to maximize profits if one or more competitors are present. We answer the following research questions:

RQ 1.1 How much can operators gain from considering the presence of competition in their rebalancing operations with regards to gross profits? Put differently, what is the price of ignoring the presence of competition?

RQ 1.2 How much is lost by competing in comparison to jointly optimizing fleet rebalancing with regard to gross profits, and how do alternative business models under competition compare to each other?

RQ 1.3 Which features drive the gains from considering competition, and the losses due to the presence of competition?

Further, we present a novel model formulation for rebalancing carsharing systems that permits marginally decreasing payoffs, and contribute two algorithms (Iterated Best Response and Potential Function Optimizer) for finding pure strategy Nash equilibria in rebalancing problems.

Existing literature on carsharing rebalancing either does not specify a rebalancing mode, or assumes that only one mode is used (usually bike, hitching rides, or truck). However, the optimal mode vastly depends on the structure of the city, e.g., in terms of population density, quality of the road and public transport network, as well as the fleet, e.g., in terms of size and imbalance, and it might be beneficial to use multiple modes in the same city, in particular if the structure of city and fleet is not homogeneous. We answer the following research questions:

RQ 2.1 Can a good mode be selected a-priori based upon features of the fleet and city?

RQ 2.2 Which features drive the choice of the optimal rebalancing mode?

To answer these research questions we integrate data analytics methods with operational (and strategic) decision making. Due to this combination, features driving the modal choice are readily available from the models.

Carsharing and ride-hailing sharing operators will soon face the decision of whether or not to offer driverless vehicles, and how many vehicles of each type they shall procure. Driverless vehicles are beneficial from an operational standpoint, making rebalancing easier, but they are also more expensive to procure. We study a technology choice problem called "fleet sizing and composition" that balances operational profits and procurement costs, and find which aspects of a fleet drive the usage of driverless vehicles and which factors deter operators from using them.

RQ 3.1 Should shared mobility operators use fleets with only one vehicle type, or mix among driverless and human-driven vehicles?

RQ 3.2 Under which circumstances can shared mobility operators benefit from introducing driverless vehicles in their fleet?

RQ 3.3 How much can operators gain with respect to total profits from using mixed fleets comprised of driverless vehicles?

This thesis contributes a bound-and-enumerate algorithm for solving the technology choice problem. The optimal rebalancing policy is derived using Semi-Markov Decision Processes (SMDPs) or a fluid approximation.

1.3 Outline

Chapter 2 reviews related literature, focusing on rebalancing in shared mobility systems (both in the static/deterministic and dynamic/stochastic case, with and without competition) and fleet sizing and composition in shared mobility systems, as well as on the broader literature on technology choice.

The three main chapters of this thesis are based on three working papers.

Chapter 3 investigates the impact of competition on the routing and servicing decision. We find that competition severely impacts the profitability of carsharing operators, and ignoring competition in the rebalancing and servicing decision can frequently be worse than no rebalancing at all. Thus, operators who face competition should take the current position of the competitors' fleets as well as their rebalancing activities into consideration. We further observe that market segmentation (two or more operators who consider competition) is substantially less profitable than a single operator, since rebalancing operations cannot be pooled as easily. Chapter 3 is based on Martin et al. (2020b).

In Chapter 4, we study in-depth which modes are necessary in a solution. We build classifiers that – given properties of the fleet and the city the fleet is deployed in – estimate individual costs for each mode using a linear regression classifier and the probability that a mode is cost-minimizing using logistic regression, and also predict the best mode using decision trees. We compare different setups of these classifiers, and find that they perform well with respect to cost of misclassification and accuracy. The classifiers help to establish the most important features in the modal choice. Chapter 4 is based on Martin and Minner (2020).

Chapter 5 focuses on the optimal fleet size and composition of a shared mobility operator (e.g., carsharing or ridesharing). When driverless vehicles are first introduced in the fleet mix, they provide substantial operational benefits (lower rebalancing and operational costs), but usually impose a significantly higher investment cost upfront. We model this problem as a two-stage problem with first fixing fleet size and composition, and subsequently the routing and rebalancing decision. In this chapter, the rebalancing differs from the previous chapters, since rather than considering rebalancing during the night, we focus on dynamic rebalancing during the day. Chapter 5 is based on Martin et al. (2020a).

Chapter 6 concludes this thesis.

Chapter 2

Related Literature

This literature review focuses on optimal operations of shared mobility (see Banerjee and Johari (2019), He et al. (2019), and Laporte et al. (2015, 2018) for other recent reviews). We restrict ourselves to car-based shared mobility, i.e., carsharing and ride-hailing, and omit bikesharing (see Freund et al. (2019) for a recent review). More specifically, we restrict ourselves to literature on rebalancing operations in the deterministic static case (related to Chapters 3 and 4) and stochastic dynamic case (related to Chapter 5), as well as competitive extensions (related to Chapter 3), the underlying routing models (related to Chapter 4), and the strategic decision on the optimal fleet size and composition (related to Chapter 5).

2.1 Rebalancing in Shared Mobility

This thesis is closely related to literature on rebalancing in shared mobility. In this field, Illgen and Höck (2019) provide a recent overview. Rebalancing literature can be classified by different features: (i) static and deterministic vs. dynamic and stochastic (usually pooled rebalancing operations during the night or individual rebalancing operations during the day), (ii) the planning horizon (tactical planning of the number of vehicles in a district or at a station vs. operational planning of the tours necessary to rebalance the fleet) (iii) user-based vs. operator-based, (iv) the objective (maximization of operational profits, minimization of costs, maximization of a service level), and (v) by type of system (e.g., carsharing, bikesharing, ridesharing; e.g., station-based one-way, station-based round-trip, one-way free-floating). Most existing literature (e.g., He et al. (2020), Kek et al. (2009), and Nair and Miller-Hooks (2011)) studies the tactical problem of setting proper inventory levels, as opposed to this thesis which investigates the routing and servicing aspects of the rebalancing decision. This thesis focuses on

operator-based rebalancing. There is a growing body of literature on user-based relocation via incentive mechanisms (e.g. Pfrommer et al. (2014) and Ströhle et al. (2019)), as well as regulation of the spatio-temporal distribution of supply by dynamic pricing (e.g., Afeche et al. (2018), Banerjee et al. (2017), Bimpikis et al. (2019), and Guda and Subramanian (2019)). However, regardless of how "good" incentives are set, they cannot completely replace operator-based rebalancing, since operator-based rebalancing can tackle larger imbalances, and can be integrated with cleaning and maintenance operations (Weikl and Bogenberger, 2013). Even if vehicles can move autonomously (and can therefore pickup customers from further away), proactive rebalancing is necessary: Yang et al. (2020) find that the matching radius (i.e., the maximum distance between a customer and the vehicle she is matched with) should not be too large. Most literature on the operational rebalancing decision focuses on bikesharing (e.g., Bruck et al. (2019) and Datner et al. (2019)). The bikesharing relocation problem, however, differs from the rebalancing problem in carsharing systems in the capacity of a relocation truck as well as the average costs of a relocation operation. Thus, detailed insights generated for bikesharing cannot be directly applied to carsharing, but we obviously expect similar results in the sense that relocation improves the profitability of operations.

2.1.1 Static Rebalancing Considering Vehicle and Staff Movements

In carsharing, one must not only consider how vehicles are rebalanced, but also how rebalancing workers continue to the next vehicle. Nourinejad et al. (2015) explicitly address this joint vehicle- and staff-rebalancing problem, but during the day rather than pooled relocation operations when demand is low. They show that if demand increases, fleets grow faster than the rebalancing activities. Table 2.1 contains related literature on static vehicle and staff rebalancing. We list the rebalancing mode, as well as if the models specifically address electric vehicles, and the objective. Bruglieri et al. (2014a) study the Electric Vehicle Relocation Problem as a variant of the 1-skip relocation problem and the rollon-rolloff problem. Workers drive the vehicles, and then use bikes to continue to the next vehicle. For their Milan case study, they report that two workers are sufficient to fulfill on average 86% of 30 relocation requests. When maximizing the profit rather than a service level, they observe that a "fixed revenue component" (future revenue of satisfied customers) per served customer must be at least 15€ to ensure that enough rebalancing occurs to serve most of the customers (Bruglieri et al., 2017). Bruglieri et al. (2018) extend the above model to study the interplay of the different objectives that are minimizing the number of workers, maximizing the number of satisfied requests, and

Paper	Car	Bike	Publ. Transit	Truck	Walking	Electric	Objective
Bruglieri et al. (2014a)	✗	✓	✗	✗	✗	✓	service level
Bruglieri et al. (2017)	✗	✓	✗	✗	✗	✓	profit
Bruglieri et al. (2018)	✗	✓	✗	✗	✗	✓	multiple
Gambella et al. (2018)	✓	✗	✗	✗	✗	✓	battery depletion
Kypriadis et al. (2018)	✗	✗	✗	✗	✓	✗	walking distance
Dror et al. (1998)	✗	✗	✗	✓	✗	✗	costs
Fink and Reiners (2006)	✗	✗	✗	✓	✗	✗	costs
You and Hsieh (2014)	✗	✗	✗	✓	✗	✗	profit

Table 2.1: Comparison of Literature on Static Fleet Balancing

minimizing the maximum tour length of each worker. Gambella et al. (2018) find that an electric carsharing fleet can benefit from rebalancing, even if the demand on average is balanced. Their static routing model assumes that workers give each other lifts to the next vehicle (without explicitly mentioning this rebalancing mode).

Kypriadis et al. (2018) model carsharing rebalancing with walking from the delivery location to the pickup location as a staff-rebalancing mode. The primary objective is to minimize walking distance, minimizing the overall tour length is only a secondary objective. They apply their model to carsharing services in two Italian cities, and state that the walking distance remains substantial, even when explicitly minimizing it. However, they do not compare the walking distance or total relocation time/costs to current operations, other modes or other objectives (e.g., minimization of overall tour duration as primary objective).

Carsharing and traditional car rental share some similarities with respect to optimal operations (Oliveira et al., 2017), we therefore also consider closely related problems in this field. Dror et al. (1998) use a truck that has less capacity than the largest stations to rebalance a station-based carsharing system. They are the first to simply split stations into "cliques" consisting of single vehicles (and thus have unit demand). The cost between two nodes in the same clique is 0, and the cost between two nodes of different cliques is equal to the cost between those two cliques. The focus of their paper is model formulations rather than managerial implications. They therefore do not provide an assessment of the benefits of using trucks, or even relocating at all. Fink and Reiners (2006) consider the problem of rebalancing a car rental fleet consisting of different vehicles. Their problem considers traditional car rental rather than carsharing,

but the modelling of both business models is similar. For this fleet, they model the rebalancing decision as a rolling-horizon problem over a period of one week, and vehicles can be relocated daily using trucks. They report that up to 20% of the fleet size can be saved whilst maintaining a very high service level due to optimal fleet balancing. Krumke et al. (2013) study the k-convoy pickup and delivery problem, in which vehicles are relocated by k drivers. They study the problem in both a static and a dynamic setting and present approximation algorithms for both cases.

2.1.2 Dynamic Rebalancing

Once vehicles can move autonomously, they can also be rebalanced during the course of the day without substantial excess cost. For other types of vehicles, dynamic rebalancing is also possible, but more expensive (however, some literature focuses on these aspects, e.g., Schuijbroek et al. (2017) and Shu et al. (2013) for bikesharing as well as Smith et al. (2013) and Zhang et al. (2018) for non-autonomous mobility on demand integrated with a taxi service).

In this literature review, we restrict ourselves to those papers with a model formulation that can be applied to (partially) driverless vehicle sharing systems with the key characteristics of stochastic demand and on-demand operator-based rebalancing during operating hours rather than pooled rebalancing operations at night, or papers in which multiple vehicle types are explicitly considered. Those problems can be classified along several axes (in addition to the classification by sources of uncertainty, see Tang et al. (2020)): We focus on papers which model rebalancing using Markov chains/decision processes or (closed) queueing networks (and omit data-driven, robust and stochastic programming approaches such as Freund et al. (2020), He et al. (2020), and Tang et al. (2020)), and that can be used to model driverless fleets or fleets of driverless and human-driven vehicles. Table 2.2 gives an overview over related rebalancing problems, and lists if they consider a Markovian approach or a queueing-based approach, if operators can react at any point in time, if vehicles are driverless/driven by employees of the operator, or if fleets are mixed, if customers leave the system immediately ("no waiting"), if the operator has full control over all vehicles (i.e., can force vehicles/drivers to move to a different location) as well as the objective. Repoux et al. (2019) consider a system in which both the number of parking lots and the fleet size are restricted, and customers announce their destination. They use a Markovian model to estimate the probability and associated cost of outages at each location (either no vehicle or no parking spot), and heuristically rebalance vehicles from origin to destination such that

Paper	Markov	Queue	Flex. Intervals	Driverless	Mixed	No Waiting	Full Control	Objective
Repoux et al. (2019)	(✓)	✗	✗	✓	✗	(✓)	✓	served demand
Benjaafar et al. (2018)	✓	✗	✗	(✓)	✗	✓	✓	cost
Braverman et al. (2019)	✗	✓	✓	✓	✗	✓	✓	served demand
Zhang and Pavone (2016)	✗	✓	✓	✓	✗	✗	✓	waiting time
Zhang et al. (2018)	✗	✓	✓	✗	✓	✗	✓	waiting time
Smith et al. (2013)	✗	✓	✓	✗	✓	✗	✓	# of rebalancing vehicles & # of drivers
Wei et al. (2020)	✗	(✓)	✗	✓	✓	(✓)	✗	profit

Table 2.2: Comparison of Literature on Fleet Balancing

cost reduction over rebalancing costs are minimized. They find that dynamic relocations have a clear positive impact on profitability, and the effect of additional relocation operations is marginally decreasing. Benjaafar et al. (2018) consider a product rental network which they apply to DVD rental and free-floating carsharing. They use a Markov decision process (MDP) to set optimal rebalancing amounts between different branch offices minimizing total costs (rebalancing and lost sales). Provably, the optimal rebalancing policy in all decision epochs is to rebalance as little vehicles as possible to reach a convex subset of the state-space called "no-rebalancing region". The disadvantage of both Markovian approaches is that decisions can only be taken after fixed time intervals.

Braverman et al. (2019) use a BCMP network queueing model to find the asymptotically revenue-maximizing/service-level-maximizing routing decision for empty vehicles assuming that customers leave the system if no vehicle is available. They prove that their policy is asympotically optimal for infinite fleet sizes. Zhang and Pavone (2016) model the required relocation operations in an Autonomous Mobility-on-Demand (AMoD) system by introducing optimal virtual customer streams in a Jackson network model. Their case study that reveals that roughly 70% of the current taxi fleet were sufficient if all taxis were replaced by driverless vehicles.

The literature on rebalancing with multiple vehicle types (e.g., driverless and human-driven vehicles) is limited. Zhang et al. (2018) extend the model of Zhang and Pavone (2016) to independently track driverless and human-driven vehicles, assuming that only driverless vehicles can be rebalanced. Smith et al. (2013) use a similar approach in which vehicles and drivers are tracked independently, and vehicles can only be rebal-

anced if there is a driver (in a figurative manner, the presence of drivers makes vehicles autonomous). Smith et al. (2013) and Zhang et al. (2018) show that most vehicles (2/3-3/4 of the fleet) in their New York City case studies can be driven only by customers, and only 1/4-1/3 of the fleet shall be driverless, or driven by an employee. Wei et al. (2020) extend the model of Bimpikis et al. (2019) to spatial pricing in presence of driverless vehicles (owned or rented by the operator) and human-driven vehicles (owned by independent drivers who join the platform if the expected revenue exceeds some threshold).

2.1.3 Competitive Routing and Rebalancing Models

Only few studies take a competitive view on rebalancing carsharing operations, and those take a tactical rather than an operational point of view or consider pricing instead of routing. Albiński and Minner (2020) calculate the number of vehicles that are to be relocated to reach an equilibrium under demand uncertainty (tactical inventory transshipment problem). However, they do not focus on the actual routes of workers who relocate the vehicles (operational routing problem). Balac et al. (2019) use an agent-based model to derive optimal prices for two carsharing operators and other options such as walking or public transport. They provide insights into whether or not it is advisable to additionally rebalance vehicles during the course of the day. They observe that charging the same (comparably high) price is most profitable for both operators. However, high prices are unstable as operators benefit from offering lower prices than their competitors. They state that relocations during the course of the day are unprofitable, and further observe that relocation in presence of competition primarily benefits the competitor who does not rebalance. The origin of this "free-rider" phenomenon, however, is demand during the relocation operations and the selection of a simple policy-based heuristic for relocation. Both simplifications do not apply in our model. Both aforementioned papers are restricted to the two-operator case. For the ridesharing sector with a variable number of operators, Pandey et al. (2019) argue that competition decreases the efficiency as well as the quality of service, and show that even little cooperation can substantially increase the service quality. Their main focus, however, is not on rebalancing during the night, but on assigning customers to routes and the subsequent online re-routing of vehicles.

Also for the ridesharing sector (with competition between traditional taxis and ridesharing services such as DiDi), Yu et al. (2020) study optimal regulations of the ridesharing service. They observe that the ridesharing sector has the potential to drive taxis out

of service by under-cutting prices, and later raise prices, and suggest that regulations can circumvent this. Lanzetti et al. (2020) study the competition between a ridesharing operator and public transportation authorities. They find that these different services are in fact competing, and that the extent of cannibalization depends on local environment. Bernstein et al. (2020) find that ride-hailing drivers operating on two different platforms ("multihome") improves their payoff individually, but the system reaches a prisoners' dilemma if all drivers multihome. Thus, Bernstein et al. (2020) call for incentive mechanisms that make it individually rational to offer one's service on only a single platform.

2.2 Fleet Composition and Technology Choice

While the literature on fleet rebalancing has grown significantly in the last decade, the literature on strategic fleet sizing and composition problems, an application of technology choice, has only recently gained currency.

George and Xia (2011) study the fleet sizing problem in vehicle sharing systems. They use a queueing-theoretical model to calculate the cost-minimal fleet procurement strategy to satisfy all customer demand, assuming that the flows in the system are balanced. Several papers integrate the fleet sizing decision with station location (Freund et al., 2018; Nair and Miller-Hooks, 2014), service region design (He et al., 2017), procurement of parking permits (Lu et al., 2018), or pricing (Al-Kanj et al., 2020). The fleet composition problem has been studied less frequently. Bellos et al. (2017) find that highly fuel-efficient vehicles in a carsharing fleet can help original equipment manufacturers to meet environmental regulations. Dandl and Bogenberger (2018) and Feng et al. (2020) compare fleet sizes between carsharing with human-driven vehicles and driverless ride-hailing, and street-hailing and ride-hailing, respectively. Dandl and Bogenberger (2018) find that ability to pick up customers (associated with waiting) allows operators to reduce the fleet size. Feng et al. (2020) state that full knowledge about the spatial distribution of current customer demand (given by reservations through an app) allows operators to reduce the fleet size. For peer-to-peer vehicle sharing, Baron et al. (2018) find that society may benefit from full, null or partial automation, depending on the extent of additional traffic and the improvements in travel flow.

Technology choice has frequently been studied in the production context, in which decision makers either choose a flexible technology that can produce multiple products, or dedicated machines for each product. Flexible machines allow to hedge against un-

certainty in the demand (Fine and Freund, 1990; Van Mieghem, 1998), unless demands are perfectly correlated and no cost differentials can be exploited. Mixing flexible and dedicated technology can further improve technology choice (Jordan and Graves, 1995). The technology choice problem we consider differs from the traditional technology choice in one key aspect: Demand for both technologies is indistinguishable, but flexibility by mixing technologies is beneficial due to lower operational costs and higher contribution margins. Power producers often face similar tradeoffs: While renewable energy sources are beneficial, their yield is uncertain and intermittent. Thus, power producers also invest in other (flexible or inflexible) power plants. For a Texas, USA, case study, Kök et al. (2020) find that it is in fact beneficial to mix inflexible (high constant yield, medium production cost), renewable (intermittent yield, neglible production cost) and flexible (variable yield, high production cost) energy sources, balancing investment cost and operational flexibility. Moccia and Laporte (2016) study a technology choice problem for mass transit: While buses can only increase frequency, operators can also dispatch additional carriages for light rail. Then, the more expensive mode light rail is preferred if demand is sufficiently large for some origin-destination pair, and both modes may be mixed in the network.

Chapter 3

The Competitive Pickup and Delivery Orienteering Problem for Balancing Carsharing Systems

Competition between one-way carsharing operators is currently increasing. Fleet relocation as a means to compensate demand imbalances constitutes a major cost factor in a business with low profit margins. Existing decision support models have so far ignored the aspect of a competitor when the fleet is rebalanced for better availability. We present mixed-integer linear programming formulations for a Pickup and Delivery Orienteering Problem under different business models with multiple (competing) operators. Structural solution properties, including existence of equilibria and bounds on losses due to competition, of the Competitive Pickup and Delivery Problem under the restrictions of unit-demand stations, homogeneous payoffs and indifferent customers based on results for congestion games are derived. Two algorithms to find a Nash equilibrium for real-life instances are proposed. One can find equilibria in the most general case, the other can only be applied if the game can be represented as a congestion game, that is under the restrictions of homogeneous payoffs, unit-demand stations, and indifferent customers. In a numerical study, we compare different business models for carsharing operations, including a merger between operators and outsourcing relocation operations to a common service provider (coopetition). Gross profit improvements achieved by explicitly incorporating competitor decisions are substantial; and the presence of competition decreases gross profits for all operators (compared to a merger). Using a Munich, Germany, case study, we quantify the gross profit gains due to considering competition as approx. 35% (over assuming absence of competition) and 12% (over assuming omnipresence of competition), and the losses due to presence of competition to be approx. 10%.

3.1 Introduction

Carsharing is an economically and environmentally sustainable alternative to private car ownership and a supplement to public transportation (Bellos et al., 2017). With a larger number of carsharing offers around the world, competition increases. In several cities, more than one carsharing operator offers its service to customers (Kortum et al., 2016), and Mobility-as-a-Service apps allow customers to book the closest vehicle, regardless of the operator. Some operators are starting to merge their companies and fleets, most recently Car2Go and DriveNow (now called "ShareNow") (BMW Group, 2018). Soon after, the former DriveNow shareholder Sixt launched a carsharing service "Sixt Share" that competes with ShareNow (Sixt SE, 2019). Perboli et al. (2018) report frequent changes in the Turin market, with BlueTurino entering the market and competing with Car2Go and Enjoy, while IoGuido withdrew service. Although the body of literature on the optimization of carsharing operations is growing, it mostly has ignored the choice of business and operational models under competition so far (with exceptions of Albiński and Minner (2020) and Balac et al. (2019)). In addition to the merger (and, thus, a monopoly) as currently pursued by DriveNow and Car2Go, and direct competition such as between Sixt Share and ShareNow, operators can cooperate in parts of their operations. For example, they can hire a third-party that relocates vehicles for them, or both operators relocate vehicles such that the overall gross profit is maximized (but still use different workers for rebalancing). We call these modes "coopetition" and "welfare-maximization". Ghosh et al. (2017) report outsourced relocation operations in bikesharing. Brook (2004) reports a collaboration between different carsharing operators concerning other aspects such as billing. In practice, operators frequently ignore competition with respect to relocation. Mostly, they do not even consider the current location of vehicles of the competitor which can be accessed using webscraping techniques, or accessing publicly available application programming interfaces (Kortum et al., 2016; Trentini and Losacco, 2017), let alone foresee how the competitor rebalances those vehicles (which could be derived from historic data). Incorporating the reaction of a competitor on one's own routing and servicing decision is paramount for profitable operations, in particular if fleets are small and margins are low. Operators may service the same customer, thereby reducing the individual payoff, but still incurring relocation costs.

A particular focus of this work is rebalancing which shows significant potential for cost reduction in carsharing systems (Jorge and Correia, 2013). One-way carsharing comes

at the cost of unevenly distributed fleets, as some locations (such as shopping malls) are more frequently the point of origin of a carsharing trip than of its destination (or vice versa). To cause as little customer dissatisfaction as possible and, thus, as few lost sales as possible, vehicles are relocated (predominantly by pooling multiple rebalancing operations during the night Weikl and Bogenberger (2015)). Vasconcelos et al. (2017) report a positive impact of relocation on the profitability of a carsharing service. They state that a carsharing service is only profitable if relocations are performed. Consequently, non-optimal carsharing relocations (or relocations that do not consider competition), can easily result in negative profits.

We answer three research questions and contribute algorithmically to solution methods for carsharing relocation under competition.

RQ 1 How much can operators gain from considering the presence of competition in their rebalancing operations with regards to gross profits? Put differently, what is the price of ignoring the presence of competition?

RQ 2 How much is lost by competing in comparison to jointly optimizing fleet rebalancing with regard to gross profits, and how do alternative business models under competition compare to each other?

RQ 3 Which features drive the gains from considering competition, and the losses due to the presence of competition?

We address the research questions analytically under some technical assumptions, and study the full generality of the model in an extensive numerical study.

Methodologically, we present models and solution algorithms to tackle the aforementioned research questions. We give complexity as well as average case algorithmic performance results. Advantages and drawbacks that the solution algorithms and model formulations entail are studied. We devise a simplified model for carsharing relocation that incorporates the movements of vehicles, as well as the movements of the workers who relocate them. This model is called Pickup and Delivery Orienteering Problem (PDOP). A feasible solution to the PDOP takes two decisions simultaneously: First, it determines which vehicles will be moved and which stations will be serviced with how many vehicles (customer demand is given in a previous step, e.g., by using a demand prediction model as reviewed in Vosooghi et al. (2017)). Second, it specifies which route relocation workers take to perform those relocations. We then create variations of this model to facilitate the special features of different business models under competition: Operators can either

compete directly (Competitive Pickup and Delivery Orienteering Problem (C-PDOP)), jointly optimize their rebalancing (Welfare-Maximizing Pickup and Delivery Orienteering Problem (W-PDOP)), merge their fleets completely (Monopoly Pickup and Delivery Orienteering Problem (M-PDOP)), or outsource their relocation operations to a third party (Coopetitive Pickup and Delivery Orienteering Problem (Coop-PDOP)).

We choose pure strategy Nash equilibria as the solution concept (as opposed to mixed strategy Nash equilibria where players randomize over a set of strategies) which is a single routing and servicing choice rather than a combination of multiple choices. Mixed strategy Nash equilibria provably exist in the C-PDOP since the number of pure strategies is finite (Nash, 1951). Pure strategy Nash equilibria represent a more intuitive, easier to implement solution for the carsharing operators. They represent less controversial concept of competition, and are more likely to be adopted by carsharing operators. Pure strategy Nash equilibria provably exist subject to three assumptions: player-homogeneous payoffs, unit-demand stations, and indifferent customer choice. For the C-PDOP with the above assumptions, we present structural properties and performance guarantees (algorithmic and w.r.t. profitability), and show average case performances in a numerical study. The numerical study also considers the influence of relaxing the assumptions.

To solve the C-PDOP, we present two different algorithms: Iterated Best Response (IBR) and Potential Function Optimizer (PFO). While IBR is faster on most instances and can be used even for the general model where pure strategy Nash equilibria are not guaranteed to exist, PFO returns higher-profit Nash equilibria on average and has a higher degree of fairness with respect to how profits are distributed between operators. Both algorithms draw upon the fact that the C-PDOP with player-homogeneous payoffs, unit-demand stations, and indifferent customer choice belongs to the class of congestion games/potential games, introduced in Monderer and Shapley (1996) and Rosenthal (1973).

The remainder of this Chapter is structured as follows. Section 3.2 presents the model for the PDOP and extensions for the various business models. It also investigates properties of these models. In Section 3.3, we present the IBR and the PFO to find a pure strategy Nash equilibrium of the C-PDOP if it exists. In Section 3.4, we test the aforementioned algorithms on several different, artificially generated instances and a Munich, Germany, case study, before concluding the Chapter in Section 3.5. All proofs can be found in the Appendix 3.A.

3.2 Model

In the following, we first present a master model for carsharing relocation. We then instantiate this model to a competitive variant (C-PDOP), a monopoly variant (M-PDOP) and a coopetition variant (Coop-PDOP), and discuss properties of these problems.

3.2.1 Pickup and Delivery Orienteering Problem

In our carsharing model, each of N operators knows her current fleet distribution and wants to rebalance it, assuming that no customers will be requesting vehicles during the process. We use index $n = 1, 2, \ldots, N$ to refer to each of the operators. In carsharing systems, such relocations are usually performed periodically during each night, and operators can "block" vehicles from being reserved by customers (Weikl and Bogenberger, 2015). To rebalance her fleet, operator n sends workers who drive the vehicles from their current location to another location where they are expected to incur a higher payoff. The overall goal of an operator is to maximize her gross profit (profit of a given fleet), defined as payoffs minus operational costs. The payoffs depend on the presence of competition. While the strategy of the competitors need not be known, all operators must be able to establish the payoff functions of all competitors, and assume that their competitors act rationally and payoff-maximizing.

Stations, Locations and Payoffs

Each operator has a set of stations, D_n, where she can place vehicles. The sets of stations of two different operators are not necessarily disjoint, as operators may place stations very close to each other (e.g., at public transit stops, or the trade fair). Each station is then addressed as $\iota \in D, D = \bigcup_n D_n$. We assume that the operators employ forecasting models and, thus, estimate the expected customer demand in each station, allowing them to predict if vehicles shall be moved to or removed from this station (but not both simultaneously). Each operator n chooses, for each station ι, the number q_n^ι of vehicles to place there. For every station ι, operator n has an estimate of the payoff function $\pi_n^\iota(\mathbf{q}^\iota)$ where $\mathbf{q}^\iota = (q_1^\iota, q_2^\iota, \ldots, q_N^\iota)$ is the vector that describes the number of vehicles placed by each operator n at station ι. The quantity $\pi_n^\iota(\mathbf{q}^\iota)$ captures the direct (expected) revenue from setting a vehicle at this location, minus the direct (expected) costs such as fuel and wear; but it does not include rebalancing costs which will be modeled using a different term.

We use the standard game-theoretic notation of $\mathbf{q}^{\iota}_{-n}$ to denote the vector of the decisions of the players other than n, so that we can write $\mathbf{q}^{\iota} = (q^{\iota}_n, \mathbf{q}^{\iota}_{-n})$. The model allows for a very generic formulation and choice of these payoff functions which adhere to the following:

- keeping the values $\mathbf{q}^{\iota}_{-n}$ of the operators other than n fixed, the function $\pi^{\iota}_n\left(q^{\iota}_n, \mathbf{q}^{\iota}_{-n}\right)$ is *concave*, *non-decreasing* in q_n; the monotonicity suggests that we can only increase revenue by placing more vehicles at the same station, whereas concavity suggests that the marginal payoff (by placing one extra vehicle) decreases with the number of vehicles placed in the station;

- for $m \neq n$, keeping fixed the values $\mathbf{q}^{\iota}_{-m}$ of the operators other than m, the function $\pi^{n}_{\iota}\left(q^{\iota}_m, \mathbf{q}^{\iota}_{-m}\right)$ is *non-increasing* in q_m; this suggests that the revenue of an operator can only decrease if competing operators place more vehicles into that station.

We deliberately do not restrict $\pi^{\iota}_n(\mathbf{q}^{\iota})$ further. This allows us to formulate different types of customer choice. For example, consider a station ι where a single customer is expected who strictly prefers operator 1 (but would take operator 2 if operator 1 is unavailable). This could be modelled by $\pi^{\iota}_1(1,1) = \pi^{\iota}_1(1,0) = \pi^{\iota}_2(0,1) > 0$ and $\pi^{\iota}_2(1,1) = 0$. For another example, consider a station ι where a single customer is expected who is indifferent between operators 1 and 2. This could be modeled by $\pi^{\iota}_1(1,1) = \pi^{\iota}_2(1,1) = 0.5\pi^{\iota}_1(1,0) = 0.5\pi^{\iota}_2(0,1)$.

We further assume that there is a maximum number of vehicles that can be moved to a station profitably; i.e. q^{ι}_n is upper-bounded by some constant $\hat{q}^{\iota}_n$ that represents the maximum demand at station ι. In practice, operators usually employ some "filtering" upfront (Weikl and Bogenberger, 2015), resulting in a reasonably small number of vehicles that can be profitably placed in a station.

In this section we optimize for an operator n, considering that the other operators are "frozen" in their strategies. Thus, the payoff values $\pi^{\iota}_n\left(q^{\iota}_n, \mathbf{q}^{\iota}_{-n}\right)$ vary in q^{ι}_n only. We formulate the problem as an integer program, and as such, "break" the nonlinearity of the payoff functions π^{ι}_n into a linear objective function. The simplest way to do so is to split a station ι into a set $\mathcal{Z}^{\iota}$ of *locations* labeled $i = 1, \ldots, \hat{q}^{\iota}_n$. Each station ι can be split into at most $\hat{q}^{\iota}_n$ locations. We separate the station-based payoff function into a sum of location-based payoff values, assigning to the i-th location the marginal gain of placing the i-th vehicle. These marginal gains can be computed by consecutive differences, that is, $\pi^{i}_n = \pi^{\iota}_n\left(i, \mathbf{q}^{\iota}_{-n}\right) - \pi^{\iota}_n\left(i-1, \mathbf{q}^{\iota}_{-n}\right)$. From our assumption of concavity, the marginal gains are non-increasing, which results in an intrinsic ordering of locations, i.e., without

loss of generality an operator would choose to visit the i-th location associated to a station ι only after visiting all lower-indexed locations $1, \ldots, i-1$.

Locations are represented as nodes in a rebalancing graph. We denote the full set of nodes available to operator n as $\mathcal{Z}_n$. The operator already has vehicles present at some locations but not others; we refer to locations that do not have a vehicle as delivery locations ($\mathcal{Z}_n^-$) and to those with a vehicle present as pickup locations ($\mathcal{Z}_n^+$). Thus, we can write $\mathcal{Z}_n = \mathcal{Z}_n^+ \cup \mathcal{Z}_n^-$. Finally, we also include a depot node which is modeled as both a pickup and delivery location.

In analogy to Bruglieri et al. (2017), we formulate the problem as a variant of a TSP with profits where not all locations need to be visited (similar to prize-collecting TSPs). This allows the integration of two decisions: the decision on which locations to service, and the actual routing decision. Additionally, one can specify subsets of locations $S_n^- \subseteq \mathcal{Z}_n^-, S_n^+ \subseteq \mathcal{Z}_n^+$ that must be visited. In this way, we cater for mandatory customers (e.g., to fulfill a contract), and also enforce that vehicles that are parked illegally are picked up.

Importantly, we assume that nearby stations do not influence each other. This is arguably a simplifying assumption as users might walk to the next available vehicle in a nearby station, even if this increases the walking distance. This assumption is common in the literature on carsharing rebalancing (Bruglieri et al. (2018) for example randomly assigns customers who belong to the catchment area of multiple stations). Considering interdependent demand processes would significantly increase the problem complexity and is beyond the scope of this Chapter.

Arcs and Costs

For each pair of locations $i \in \mathcal{Z}_n^-$, $j \in \mathcal{Z}_n^+$, we associate costs $c_{\langle ij \rangle}, c_{\langle ji \rangle}$ and travel times $\tau_{\langle ij \rangle}, \tau_{\langle ji \rangle}$ between these locations. The costs serve as an estimate for the costs of the amount of fuel, as well as the payment for the time it takes to move between the locations (either proportionate wages or payments to a service provider). Note that the costs and travel times are not necessarily symmetric, and are always equal for pairs of locations from the same station (e.g., $c_{ij} = c_{kl}$ if i and k are locations for the same station ι_1 and j and l are locations for the same station ι_2). Operators drive from pickup to delivery locations, but the travel from a delivery to a pickup location may be done in a lot of different ways, such as using a foldable bike, walking, public transit, or using a second car driven by a colleague ("double-driving"). We can then describe the problem in terms of its underlying network, or directed bipartite graph G, whose nodes correspond to the

locations $\mathcal{Z}_n$, and with arcs A_n between each pickup location and each delivery location, in either direction.

Workers and Tours

Operator n employs W_n workers who relocate vehicles by starting at a depot, visiting pickup and delivery locations in alternating order, and returning to the depot at the end of their shift of at most T units of time. A worker can only relocate one vehicle at a time. As already mentioned, we assume that there is a central depot where workers can collect and return equipment, cleaning supplies, and lost items. Although assuming this rebalancing mode poses a simplified model, we will later show that the game theoretical formulation can be extended to more complex cost models, such as rebalancing using a truck (that can relocate multiple vehicles), or assistance using a minibus (that can transport multiple workers).

Our problem is thus specified by the bipartite graph G (i.e. the set of locations $\mathcal{Z}_n = \mathcal{Z}_n^+ \cup \mathcal{Z}_n^-$ including two copies of the depot $d_n \in \mathcal{Z}_n^-$, $p_n \in \mathcal{Z}_n^+$ and the set of arcs A_n); for each location i (associated to a station ι), an expected revenue π_n^i (corresponding to the marginal gains of placing an i-th vehicle in the corresponding station); for each arc e, a travel cost c_e and travel time τ_e; sets of enforced visits $S_n^+ \subseteq \mathcal{Z}_n^+$, $S_n^- \subseteq \mathcal{Z}_n^-$; a number of workers W_n; and a maximum shift time T. We set up an integer programming formulation to search for a profit-maximizing tour. The decision variables are:

- for each location i, a_i^n denotes whether operator n has a vehicle at location i after rebalancing; in other words, $a_i^n = 1$ iff: either i is a pickup location where n chooses not to remove the vehicle, or i is a delivery location where n chooses to serve with a vehicle;
- for each arc e, x_e^n denotes whether operator n chooses arc e on one of her routes;
- for each location i, t_i denotes the point in time at which location i is visited.

$$\max\ \Pi_n = \sum_{i \in \mathcal{Z}_n \setminus \{d_n, p_n\}} \pi_i^n a_i^n - \sum_{e \in A_n} c_e x_e^n \tag{3.1a}$$

$$\text{s.t.}\quad a_i^n = \sum_{j \in \mathcal{Z}_n^+} x_{\langle ij \rangle}^n = \sum_{j \in \mathcal{Z}_n^+} x_{\langle ji \rangle}^n \qquad \forall i \in \mathcal{Z}_n^- \setminus \{d_n\} \tag{3.1b}$$

$$1 - a_i^n = \sum_{j \in \mathcal{Z}_n^-} x_{\langle ij \rangle}^n = \sum_{j \in \mathcal{Z}_n^-} x_{\langle ji \rangle}^n \qquad \forall i \in \mathcal{Z}_n^+ \setminus \{p_n\} \tag{3.1c}$$

$$\sum_{i \in \mathcal{Z}_n^+ \setminus \{p_n\}} x^n_{\langle d_n i \rangle} \leq W_n \tag{3.1d}$$

$$\sum_{i \in \mathcal{Z}_n^- \setminus \{d_n\}} x^n_{\langle i p_n \rangle} \leq W_n \tag{3.1e}$$

$$t_i + x_{ij}\,(\tau_{ij} + T) \leq t_j + T \qquad \forall \langle i, j \rangle \in A_n \tag{3.1f}$$

$$\sum_{e \in A(C)} x^n_e \leq |C| - 1 \qquad \forall C \subseteq \mathcal{Z}_n \setminus \{d_n, p_n\} \tag{3.1g}$$

$$a^n_i = 1 \qquad \forall i \in S_n^- \tag{3.1h}$$

$$a^n_i = 0 \qquad \forall i \in S_n^+ \tag{3.1i}$$

$$x^n_e, a^n_i \in \{0, 1\} \qquad \forall e \in A_n, i \in \mathcal{Z}_n \tag{3.1j}$$

$$0 \leq t_i \leq T \qquad \forall i \in \mathcal{Z}_n \tag{3.1k}$$

The above model is an alternative notation of a multi-vehicle profitable tour problem on a bipartite graph and differs from the 2-index notation of the VRP in three key components: First, the operator maximizes her gross profit; second, visiting a location is not mandatory – instead there are only incoming and outgoing arcs if a location is visited; and third, the graph G is bipartite, which restricts the routing options and necessitates a split depot. The model can also be interpreted as a pickup and delivery problem or as a dial-a-ride problem with unit capacity (see Parragh et al. (2008) for reviews on these problem classes). (3.1a) maximizes the gross profit given by payoffs minus costs (assuming fleet procurement costs are sunk). (3.1b) and (3.1c) are assignment constraints that link availability and visits of locations, ensuring that each location has the same number of incoming and outgoing arcs (flow conservation), at most one of each. (3.1d) and (3.1e) guarantee that no more than W_n workers leave the depot and return there (either one would be sufficient and directly imply the other). Constraints (3.1f) ensure that all workers return to the depot within the shift length T. Constraints (3.1g) are used for subtour elimination; although formally they are redundant because of constraints (3.1f), including these constraints helps improve the runtime in practice. (3.1h) and (3.1i) specify that all locations in the sets S_n^- and S_n^+ must be visited, that is, a vehicle must either be left there or removed from there.

3.2.2 Competitive Pickup and Delivery Orienteering Problem

In the PDOP model, each operator implicitly assumes that the number of vehicles at a station is known. However, competitors will also react on how many vehicles this oper-

ator deploys at some station. With the rise of Mobility-as-a-Service (MaaS) solutions, the number of customers who are registered with multiple operators will grow which increases the relevance of considering competition in the optimization models.

We define competitive stations as those that can be accessed by multiple players. For example, if shared customers are expected at station ι, and if before relocation two or more operators have a vehicle there, it can be beneficial for one of them to pick up the car and service another station. In the Competitive Pickup and Delivery Problem (C-PDOP), each operator relocates her fleet with the goal of maximizing the gross profit, while considering that the other operators are relocating their fleets and strategizing accordingly. Thus, we consider Nash equilibria as the desired solution concept; i.e. we search for a strategy profile where no operator can benefit from unilateral deviation. A strategy defines the number of vehicles operator n has at a station ι after rebalancing, i.e. the vector $\mathbf{q}_n$. We refer to $\mathbf{q}_n$ as a compact strategy, as it is sufficient to represent the entire solution. In other words, the routing decisions follow directly from the availability at the locations of player n by calculating an optimal solution to (3.1a)-(3.1k).

The model in equations (3.1a)-(3.1k) is further extended by introducing the competitive profit functions Π_n, for each player n. Let $\mathbf{q} = (\mathbf{q}_1, \ldots, \mathbf{q}_N)$ denote the joint profile of (compact) strategies; $\mathbf{q}_{-n}$ denotes the joint profile of all players except n. From $\mathbf{q}_{-n}$ and for a given player n, the competitive profit is defined as the difference between station-separable payoff and costs, $\Pi_n(\mathbf{q}) = R_n(\mathbf{q}) - C_n(\mathbf{q})$. The cost term can be written as $C_n(\mathbf{q}) = C_n(\mathbf{q}_n) = \sum_{e \in A_n} c_e x_e^n$, where x_e^n is an optimal (i.e. cost-minimizing) choice of routing decisions. Notice that the cost terms depend only on the strategy of operator n; henceforth, we shall write $C_n(\mathbf{q}_n)$ instead.

The competitive profit functions can be written as

$$\Pi_n(\mathbf{q}) = R_n(\mathbf{q}) - C_n(\mathbf{q}_n) = \sum_{\iota \in D_n} \pi_n^\iota\left(\mathbf{q}_{-n}^\iota\right) - C_n(\mathbf{q}_n). \tag{3.2}$$

Given the definition of the strategies and the associated gross profits, a strategy profile $\mathbf{q}$ forms a Nash equilibrium if and only if, for every player n,

$$\Pi_n(\mathbf{q}) = \max_{\mathbf{q}_n'} \Pi_n(\mathbf{q}_n', \mathbf{q}_{-n}) \tag{3.3}$$

that is, the Nash equilibrium strategies maximize the gross profit of the players, given the other operators' strategies as input. Thus, no player can profit from deviating unilaterally.

We should state here that our model allows for different profit functions π_n^ι, and as a consequence, such games may in general not have pure strategy Nash equilibria (we provide an example in this section). From a theoretical perspective, we can study conditions under which pure strategy Nash equilibria are guaranteed to exist.

Unit-demand Stations

By this we mean that at most one customer is expected at station ι, and accordingly, each operator places at most one vehicle there. Under this assumption, our notation can be simplified, since the notions of stations and locations become equivalent, and in particular we have $q_n^\iota = a_n^i \in \{0, 1\}$, where i is the unique location associated with station ι. Also under this assumption, we consider competitive locations rather than competitive stations; we let $\mathcal{Z} = \cup_n \mathcal{Z}_n$ denote the set of all locations and $\mathcal{Z}^C$ denote the set of competitive locations.

Indifferent Customer Choice

That is, customers do not have a preference among different operators, and thus select vehicles at random and with equal probability. In particular, under this assumption, all customers are indifferent between all operators they are subscribed to (although different customers may still have different subscriptions). This assumption is realistic if customers with multiple memberships choose the closest vehicle regardless of the operator. Note that, if we assume both unit-demand stations as well as indifferent customer choice, the payoff function $\pi_n^i(\mathbf{q}^i)$ can be described rather succinctly: if $\bar{\pi}_n^i = \pi_n^i(1, 0_{-n})$ is the payoff that player n could extract by being the only operator at location i, and $m = |\{n' \neq n : q_{n'}^i = 1\}|$ is the number of operators at location i (excluding n herself), then

$$\pi_n^i(\mathbf{s}^i) = \frac{\bar{\pi}_n^i}{m+1}.$$

In such a situation, we can efficiently specify the payoff functions with a single parameter $\bar{\pi}_n^i$ for each (location, operator) pair.

Player-Homogeneous Payoffs

We assume that $\pi_n^\iota = \pi^\iota$ does not depend on n, for all stations ι, if n offers her service at location ι. This assumption is justified by evidence that margins are driven by very similar revenues (see Balac et al. (2019)). When combined with the first two assumptions,

this further restricts the dimensionality of the problem: we only need to specify, for each location i, a value π^i for the payoff that any player could extract by being the sole operator at i.

Most of our theoretical results will be proven under the combination of all three assumptions. From now on, we use the terminology of *restricted C-PDOP model* whenever we refer to an instance satisfying all three assumptions (unit-demand stations, indifferent customer choice, and player-homogeneous payoffs). The restricted C-PDOP model has theoretical advantages: it allows for a congestion game formulation (see Lemma 3.1); thus, we can guarantee the existence of at least one pure strategy Nash equilibrium, and that such equilibria can be reached via best-response dynamics (see Corollary 3.1). If the above assumptions do not hold, cases without pure strategy Nash equilibria can exist, meaning that stability is not guaranteed. In Section 3.4 we will indeed experiment with the generality of the model, since the assumptions above can be very restrictive for some carsharing systems.

3.2.3 Examples of Games without Nash Equilibria

Dropping either of the three assumptions (even while keeping the other two), we can construct games without a pure strategy Nash equilibrium.

Multi-Demand Stations

In the example of Figure 3.1, we assume indifferent customer choice and player-homogeneous payoffs, but allow for multi-demand stations. If a total of q vehicles are placed at a station ι, then a total revenue of $\pi^\iota(q)$ is extracted from this station. The revenue is split in proportion to the amount of vehicles by each player; thus, if player 1 and 2 put q_1 and q_2 vehicles respectively (so that $q = q_1 + q_2$), then they receive payoffs of $\frac{q_1}{q}\pi^\iota(q)$ and $\frac{q_2}{q}\pi^\iota(q)$, respectively. Player 1 routes on the left half of the graph, and has essentially two undominated strategies: put two vehicles in each of the stations A and B, or put two vehicles in each of the stations C and D. Similarly, player 2 routes on the right half of the graph, and has essentially two undominated strategies: put a vehicle in each of the stations A and C, or put a vehicle in each of the stations B and D. For simplicity, we assume that the routing costs of each of these strategies is normalized to 0, so we only need to worry about the way payoffs are split.

Finally, we choose the station payoffs as in Figure 3.1. For illustration purposes, let us compute the gross profits if player 1 takes the 'AABB' tour, while player 2 takes the

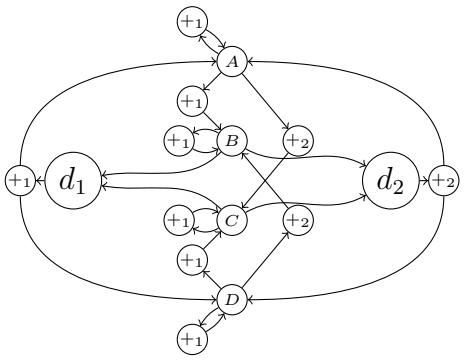

q	1	2	3
$\pi^A(q)$	4	6	7
$\pi^B(q)$	4	7	8
$\pi^C(q)$	4	7	8
$\pi^D(q)$	4	6	7

	0	AC	DB
0	$0, 0$	$0, \frac{24}{3}$	$0, \frac{24}{3}$
AABB	$\frac{39}{3}, 0$	$\frac{35}{3}, \frac{19}{3}$	$\frac{34}{3}, \frac{20}{3}$
DDCC	$\frac{39}{3}, 0$	$\frac{34}{3}, \frac{20}{3}$	$\frac{35}{3}, \frac{19}{3}$

Figure 3.1: Example of a C-PDOP instance with no pure Nash equilibria; there are four delivery districts, labeled A, B, C, D; each player has a depot (d_1 or d_2) and vehicles at distinct pickup locations ($+_1$ and $+_2$ respectively) with zero payoff; player 1 can put at most two vehicles in a district, and has essentially three strategies ('null', 'AABB' tour and 'DDCC' tour); player 2 can put at most two vehicles in a district, and has essentially three strategies ('null', 'AC' tour and 'DB' tour); the concave payoffs are player-homogeneous and follow indifferent customer choice. For a specific choice of location payoffs (center), and payoff matrix (right), no Nash equilibrium exists.

'AC' tour. In this situation, player 1 extracts two-thirds of the revenue from station A and full revenue from station B, while player 2 extracts one-third of the revenue from station A and full revenue from station C. We thus have

$$\Pi_1 = \frac{2}{3}\pi^A(3) + \pi^B(2) = \frac{14}{3} + 7 = \frac{35}{3};$$
$$\Pi_2 = \frac{1}{3}\pi^B(3) + \pi^C(1) = \frac{7}{3} + 4 = \frac{19}{3}.$$

Note that player 2 would rather deviate to the 'DB' tour and increase her gross profit. By doing these calculations for all possible pairs of strategies, we see that no Nash equilibrium exists; in particular, any sequence of iterated best responses gets stuck in the loop $(\text{AABB}, \text{AC}) \to (\text{AABB}, \text{DB}) \to (\text{DDCC}, \text{DB}) \to (\text{DDCC}, \text{AC}) \to (\text{AABB}, \text{AC})$.

Differentiated Customer Choices

In the example of Figure 3.2, we assume unit-demand stations and player-homogeneous payoffs, but allow for differentiated customer choice. Our example is built with two players. Each station ι has an associated unit revenue $\pi^\iota = 1$ that can be extracted. If both operators serve this station, the revenue is split unequally among the operators. In stations A and C, the revenue is split $75 - 25$ in favor of player 1, but in stations

B and D the revenue is split $75 - 25$ in favor of player 2. Thus, if for example both players service station A, then they receive payoffs of 3/4 and 1/4 respectively. Player 1 routes on the left half of the graph, and has essentially two undominated strategies: service stations A and B, or service stations C and D. Similarly, player 2 routes on the right half of the graph, and has essentially two undominated strategies: service stations A and C, or service stations B and D. For simplicity, we assume that the routing costs of each of these strategies is normalized to 0, so we only need to worry about the way payoffs are split.

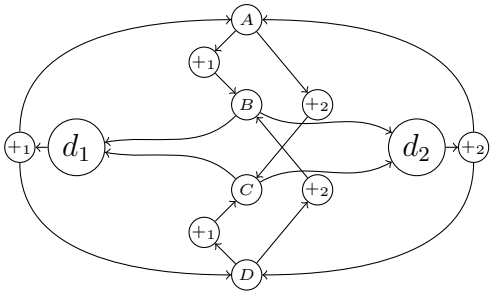

Split	1	2
A	3/4	1/4
B	1/4	3/4
C	1/4	3/4
D	3/4	1/4

	0	AC	DB
0	$0, 0$	$0, 2$	$0, 2$
AB	$2, 0$	$\frac{7}{4}, \frac{5}{4}$	$\frac{5}{4}, \frac{7}{4}$
DC	$2, 0$	$\frac{5}{4}, \frac{7}{4}$	$\frac{7}{4}, \frac{5}{4}$

Figure 3.2: Example of a C-PDOP instance with no pure Nash equilibria; there are four delivery unit-capacity districts, labeled A, B, C, D; each player has a depot (d_1 or d_2) and vehicles at distinct pickup locations ($+_1$ and $+_2$ respectively) with zero payoff; player 1 has essentially three strategies ('null', 'AB' tour and 'DC' tour); player 2 has essentially three strategies ('null', 'AC' tour and 'DB' tour); the payoffs are player-homogeneous with differentiated customer choice (center). Looking at the payoff matrix (right), we can see that no Nash equilibrium exists.

For illustration purposes, let us compute the gross profits if player 1 takes the 'AB' tour, while player 2 takes the 'AC' tour. In this situation, player 1 extracts 75% of the revenue from station A and full revenue from station B, while player 2 extracts 25% of the revenue from station A and full revenue from station C. We thus have

$$\Pi_1 = \frac{3}{4} + 1 = \frac{7}{4};$$
$$\Pi_2 = \frac{1}{4} + 1 = \frac{5}{4}.$$

Note that player 2 would rather deviate to the 'DB' tour and increase her gross profit. By doing these calculations for all possible pairs of strategies, we see that no Nash equilibrium exists; in particular, any sequence of iterated best responses gets stuck in the loop $(\text{AB}, \text{AC}) \to (\text{AB}, \text{DB}) \to (\text{DC}, \text{DB}) \to (\text{DC}, \text{AC}) \to (\text{AB}, \text{AC})$.

Player-Heterogenous Payoffs

In the example of Figure 3.3, we assume unit-demand stations, indifferent customer choice, but not player-homogeneous payoffs. There are two players and five competitive delivery stations; each of the two players has five vehicles at pickup stations. All arcs that appear in the network have equal travel cost of 2. The arcs that do not appear in the network have a travel cost given by the induced directed graph metric (so, for example, the distance from d_1 to location 3 equals 12, since there is a 6-arc path from d_1 to 3 in the original graph). Thus, the problem instance even satisfies the triangular inequality. Note that player 1 routes on the left half of the graph, and has essentially three strategies: do nothing, service locations (1,2,3), or service locations (1,4,5). Similarly, player 2 routes on the right half of the graph, and has essentially three strategies: do nothing, service locations (5,4,3), or service locations (5,2,1). By our choice of arc costs, each of the non-trivial strategies has a travel cost of 14.

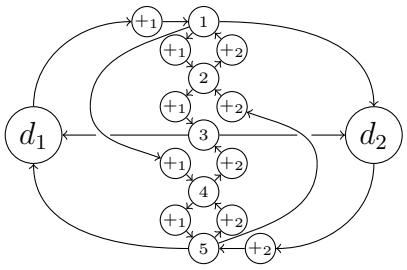

i	1	2	3	4	5
π_i^1	13	2	7	3	9
π_i^2	9	5	7	4	11

	0	S	L
0	$0, 0$	$0, 8$	$0, 11$
S	$8, 0$	$4.5, 4.5$	$0.5, 4$
L	$11, 0$	$5, 0.5$	$0, 1$

Figure 3.3: Example of a C-PDOP instance with no pure Nash equilibria; there are five delivery locations, labeled 1 to 5, which are all competitive; each player has a depot (d_1 or d_2) and five vehicles at distinct pickup locations ($+_1$ and $+_2$ respectively) with zero payoff; each player has essentially three strategies ('null', 'short' tour and 'long' tour); for a specific choice of location payoffs (center), and payoff matrix (right), no Nash equilibrium exists.

Finally, we choose the station payoffs as in Figure 3.3. Since we assume indifferent customer choice, this means that the payoff from a player at a station is halved if the other player also places a vehicle there. For illustration purposes, let us compute the gross profit if player 1 takes the "short" tour, servicing (1,2,3) and player 2 takes the "long" tour, servicing (5,2,1). In this situation, player 1 extracts full revenue from station 3; player 2 extracts full revenue from station 5; and both players extract half revenue from station 1 and 2 each. We thus have

$$\Pi_1 = \frac{13}{2} + \frac{2}{2} + 9 - 14 = \frac{1}{2};$$

$$\Pi_2 = \frac{9}{2} + \frac{5}{2} + 11 - 14 = 4.$$

Note that player 2 would rather deviate to the "short" tour and increase her gross profit. By doing these calculations for all possible pairs of strategies, we see that no Nash equilibrium exists; in particular, any sequence of iterated best responses gets stuck in the loop $(S, S) \to (L, S) \to (L, L) \to (S, L) \to (S, S)$.

The counter-examples presented here are handcrafted to show that pure strategy Nash equilibria may not exist. For these examples, the structure of competitive locations and tours exhibits many symmetries, and revenues vary vastly between the operators and between different locations. We could argue that this situation is in some sense "atypical" and would not arise in "realistic" networks. From a practical perspective, we will observe in Sections 3.4.5–3.4.7 that pure strategy Nash equilibria can be found in most instances.

3.2.4 Properties of the Competitive Pickup and Delivery Orienteering Problem

We first observe the computational hardness of solving the PDOP (and, thus, calculating optimal routes of the C-PDOP):

Theorem 3.1 (NP-Hardness)**.** *The following problems are all NP-hard, even for the special case of two players.*

1. *Given a PDOP instance, compute a tour that maximizes the gross profit.*
2. *Given a C-PDOP instance, a (compact) joint strategy* $\mathbf{q}$*, and a player* n*, compute* $\Pi_n(\mathbf{q})$*.*
3. *Given a C-PDOP instance, a player* n*, and a joint strategy* $\mathbf{q}_{-n}$ *of the remaining players, compute an optimal response for player* n*, i.e. a strategy* $\mathbf{q}_n$ *maximizing* $\Pi_n(\mathbf{q}_n, \mathbf{q}_{-n})$*.*
4. *Given a C-PDOP instance and a joint strategy* $\mathbf{q}$*, determine whether* $\mathbf{q}$ *is a Nash equilibrium.*

In the following, we will focus on Nash equilibria, i.e., we want to establish whether they exist and if they can be reached via "simple" dynamics. Since the restricted C-PDOP is a congestion game (Lemma 3.1), the existence of *pure strategy* Nash equilibria

can be guaranteed. Congestion games are formally described by: a set of players (in our model, the different operators); and a set of resources (in our restricted model, a resource for each location and a resource for each cost-minimizing tour). Each resource has an associated payoff function which depends only on *the number of* players accessing it (in our restricted model, the payoff achieved at a given location decreases with the number of players servicing it). Finally, each player has a set of feasible strategies, or subsets of resources (in our model, this corresponds to the set of feasible tours) and the payoff incurred by a player at a joint strategy is just the sum of the payoffs of the strategies she is using.

Lemma 3.1 (Congestion Game)**.** *Any restricted C-PDOP instance can be transformed into a congestion game. This transformation induces a one-to-one correspondence between (compact) strategies in the C-PDOP instance and strategies in the congestion game, that preserves improving deviations and hence the structure of Nash equilibria.*

Due to the generality of the above definition of the congestion game and the proof (see Appendix 3.A.2), this construction of the congestion game is valid for many different routing problems, as long as the cost term $C_n(\mathbf{q}_n)$ depends only on the strategy of operator n. In particular, our construction can be adapted to integrate maintenance and refueling/recharging operations, or to address other rebalancing modes, such as using a truck or minibus. From this congestion game formulation, the existence of pure strategy Nash equilibria follows directly (due to Rosenthal (1973) and Monderer and Shapley (1996)):

Corollary 3.1 (Existence of Pure Strategy Nash Equilibria)**.** *Any restricted C-PDOP instance has at least one pure strategy Nash equilibrium. Moreover, for any starting strategy* $\mathbf{q}^{(0)}$*, any sequence* $\mathbf{q}^{(0)} \to \mathbf{q}^{(1)} \to \dots$ *obtained by improving deviations (i.e. where* $\mathbf{q}^{(i+1)}$ *is obtained from* $\mathbf{q}^{(i)}$ *by an improving deviation of any one player) must eventually terminate at a Nash equilibrium.*

We can prove not only the existence of pure strategy Nash equilibria, but also the convergence of best-response dynamics to Nash equilibria. We leverage this property in the Iterated Best Response algorithm in Section 3.3. The Nash equilibrium is not necessarily unique; examples are discussed in Section 3.4.

3.2.5 Models for Comparison

To measure the impact of considering competition in the servicing and routing decision, we compare the optimal gross profit of the C-PDOP instance to problem variations in

which (i) competitors collaborate to maximize the overall gross profit, but still only serve their own customers (W-PDOP), (ii) operators merge (M-PDOP) or cooperate in their fleet relocation operations (e.g. by outsourcing to the same third party provider; Coop-PDOP), (iii) competitors ignore each other (Optimistic Quasi-Monopolistic Pickup and Delivery Orienteering Problem (QMO); Pessimistic Quasi-Monopolistic Pickup and Delivery Orienteering Problem (QMP)).

Welfare-Maximizing PDOP

We can measure the losses in gross profit due to competition by comparing the results of the C-PDOP to the system optimum, that is, the welfare-maximizing strategy. Intuitively, this corresponds to a situation in which the operators relocate their fleets by themselves, but only the operator who benefits most from a competitive location serves it. We merge all constraints of model (3.1a)–(3.1k) and add up the objective functions of each player:

$$\max \Pi(\mathbf{q}) = \sum_{n} \Pi_n(\mathbf{q}) \tag{3.4}$$

Regarding the gross profits, we observe:

Corollary 3.2 (Monotonicity of Profits (Welfare-Maximizing Solution))**.** *The optimal gross profit of the welfare-maximizing solution is never less than the joint gross profit of the competitive solution.*

The welfare-maximizizing gross profit can therefore – similar to the costs in the shared customer collaboration vehicle routing problem (Fernández et al., 2018) – serve as an upper bound for the gross profit attainable in a Nash equilibrium.

For the restricted C-PDOP model, we can analytically quantify the loss of efficiency, interpreted in terms of the price of anarchy (Koutsoupias and Papadimitriou, 1999; Roughgarden and Tardos, 2002), and price of stability (Anshelevich et al., 2008; Schulz and Stier Moses, 2003), which is defined as the worst-case bound on the ratio between the worst (resp. best) joint gross profit of a pure strategy Nash equilibrium and the joint gross profit of a welfare-maximizing solution.

Lemma 3.2 (Price of Anarchy and Price of Stability)**.** *The Price of Anarchy and Price of Stability of the C-PDOP can be arbitrarily large. For the restricted C-PDOP model with two players, we have the following results.*

1. *If* $(\mathbf{q}_1, \mathbf{q}_2)$ *is a pure strategy Nash equilibrium and* $(\mathbf{q}'_1, \mathbf{q}'_2)$ *is any strategy, then* $\Pi(\mathbf{q}'_1, \mathbf{q}'_2) \leq \Pi(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i \in \mathcal{Z}^C} \pi^i$.

2. *The absolute difference in welfare between any two Nash equilibria is at most* $\frac{1}{2}\sum_{i \in \mathcal{Z}^C} \pi^i$.

3. *The absolute difference in welfare between any Nash equilibrium and any welfare-maximizing strategy is at most* $\frac{1}{2}\sum_{i \in \mathcal{Z}^C} \pi^i$.

Moreover, all bounds in 1, 2, 3 are tight; and if $\pi^i = \pi \ \forall i \in \mathcal{Z}^C$, *then all bounds in 1, 2, 3 can be replaced by* $\frac{1}{2}\pi|\mathcal{Z}^C|$.

The proof of points 1-3 in the above Lemma relies on a simple expression for the player gross profits in the restricted C-PDOP model; as such, it cannot trivially be extended to the more general model.

Monopoly PDOP and Coopetition PDOP

In the PDOP master model, we assume that only one operator is present, whereas in the competitive C-PDOP model, each operator relocates her fleet separately, while taking into account the strategy of the other operators. We now consider two alternative business models: the Monopoly Pickup and Delivery Orienteering Problem (M-PDOP) in which the competitors merge their fleets with the objective of reducing travel costs and gross profit losses due to competition; and the Coopetition Pickup and Delivery Orienteering Problem (Coop-PDOP) in which the competitors combine their relocation efforts. In order to avoid the combinatorial explosion associated with all possible merge combinations between operators, in this section we focus exclusively on the C-PDOP model with unit-demand stations, indifferent customer choice (but not necessarily homogeneous payoffs), and two operators which consider merging/cooperating.

In the M-PDOP model, vehicles become indistinguishable. To construct an M-PDOP instance from a C-PDOP instance, we merge all locations for all operators in a station ι, and devise a joint payoff function π^ι under the assumption that every customer returns the highest payoff over all operators ($\pi^\iota(\mathbf{q}) = \max_{\mathbf{q}'}\left(\sum_n \pi^\iota_n(\mathbf{q}') \,|\, \sum_n q^\iota_n = \sum_n q^{\iota\prime}_n\right)$). In the case of unit-demand stations, servicing a location i contributes a payoff of $\pi_i = \max(\pi^1_i, \pi^2_i)$ in a monopoly. In general, the monopolist can keep both depots. In the numerical experiment, however, we assume that the depots are at the same geographical location and can, therefore, be merged.

In the Coop-PDOP, operators can collaboratively relocate their fleets (entering coopetition). This may be considered an alternative model when merging the fleets is not an option (due to strategic considerations or cartel law). One of the companies or a third party relocates the vehicles of both fleets, thereby maximizing the sum over both profit functions (with the assumption that a cost or profit sharing mechanism is implemented at a later stage). The two key differences between this and the monopoly solution are that the payoff achieved in a competitive station depends on which operator(s) serve it, and non-competitive delivery locations can only be reached from pickup locations of the same player (so, for example, one cannot generate revenue by moving a vehicle of player 1 to a location where only a customer of player 2 is expected). Similarly to the M-PDOP, we model the Coop-PDOP by devising a joint payoff function π^ι, which returns the sum of payoffs achievable in this station, given the number of vehicles each operator places there ($\pi^\iota(\mathbf{q}) = \sum_n \pi_n^\iota(\mathbf{q})$). We also exclude arcs that connect a pickup location of one player with a delivery location of another player.

Note that it is possible (in both Coop-PDOP and M-PDOP models) that a vehicle of operator 1 is moved to some competitive station ι, and at the same time a vehicle of operator 2 is removed from there. This can occur both due to inhomogeneous payoffs (the location is more attractive to player 1 than player 2), and due to a highly profitable station of player 2 elsewhere.

In practice, the number of workers in the monopoly, W_M, or in coopetition, W_O, often equals the number of workers under competition, but we do not restrict the model as such. As the M-PDOP and Coop-PDOP contain the PDOP as special cases, we immediately have the following results concerning their computational hardness.

Corollary 3.3 (NP-hardness of M-PDOP and Coop-PDOP)**.** *The following problems are NP-hard.*

1. *Given an M-PDOP instance, compute a tour that maximizes gross profit.*
2. *Given a Coop-PDOP instance, compute a tour that maximizes gross profit.*

Intuitively, one would assume that the gross profits of monopoly and coopetition solutions consistently exceed the gross profit of Nash solutions. Although this is not always true (for example, if the players serve distant operating areas with disjoint depots), we provide some assumptions which guarantee the validity of this intuition.

Lemma 3.3 (Monotonicity of Profits (Monopoly or Coopetition Solution))**.** *The following are true for the C-PDOP model, and its monopoly and coopetition variants.*

1. *The optimal gross profit of a monopoly solution is not less than the optimal gross profit of the coopetition solution.*

2. *If the number of workers is at least the sum of people working for the first and second operator ($W_M, W_O \geq W_1 + W_2$), the optimal gross profit of the monopoly or coopetition solution is not less than the optimal gross profit attainable in any pair of strategies (which includes all pure strategy Nash equilibria as well as the welfare-maximizing solution).*

Thus, in realistic settings, competing is inferior to cooperation with respect to (short-term operational) gross profits.

Quasi-Monopolistic PDOP

We identify two slightly different strategies which model that either competition or the rationality of the competitor is ignored: First, we solve the PDOP assuming that the other operators have no vehicles available in any station and does not rebalance to these locations either (optimize against $\mathbf{q}_n = \langle 0, \ldots, 0 \rangle$ for each competitor n). We call this model the "optimistic quasi-monopolistic strategy" (QMO). Second, we solve the PDOP assuming that the other operator has vehicles available at all locations after rebalancing (optimize against $\mathbf{q}_n = \langle \hat{q}_n^1, \ldots, \hat{q}_n^{|D_n|} \rangle$ for each competitor n, where $\hat{q}_n^\iota$ denotes the maximum number of vehicles that operator n can move to station $\iota \in D_n$). This model is called the "pessimistic quasi-monopolistic strategy" (QMP). While both QMO and QMP seem to be good candidates for serving as lower bounds on the gross profit of either player as well as on the total gross profit, it is possible to generate instances in which these problems can result in higher payoffs.

3.3 Algorithms for the Nash Equilibrium Calculation

The most basic approach for finding pure strategy Nash equilibria is Full Enumeration (FE): iterate through all strategy profiles, that is combinations of strategies of all operators, and test if no player has incentive to unilaterally deviate. If this is true, the strategy profile constitutes a Nash equilibrium. Although this approach obviously finds the best pure strategy Nash equilibrium (with respect to social welfare or any other metric) whenever an equilibrium exists, it takes exponentially long in the number of competitive locations. We, therefore, consider alternative approaches: Iterated Best

Response (IBR) and Potential Function Optimizer (PFO). The two algorithms represent two approaches for finding a Nash equilibrium in congestion games: utilizing the improvement dynamics of alternately improving players to find a local optimum of the potential (IBR) or (centrally) finding the global optimum of the potential (PFO). Both algorithms find a Nash equilibrium in congestion games, but they do not necessarily find a welfare-maximizing Nash equilibrium. If the congestion game property is violated, the potential is undefined, but IBR might still be able to find pure strategy Nash equilibria if they exist.

3.3.1 Iterated Best Response Algorithm

Using the IBR, we locally search for a pure strategy Nash equilibrium. We first calculate the optimal strategy for one of the players (say, player 1) against a pre-defined strategy for her competitors, e.g. assuming the competitors play the empty strategy $\mathbf{q}_n = \langle 0, \ldots, 0 \rangle$ (and hence, do not have any vehicles at any competitive location). We then use the strategy of player 1 as input for calculating the optimal strategy of player 2, then player 3 .... We continue with our calculations of best responses until the strategies no longer change. Although the best response iterations cannot be implemented in the field, we assume that operators would calculate the Nash equilibrium theoretically, and implement their equilibrium strategy.

Even if the IBR terminates, it may not necessarily return the best Nash equilibrium for one of the players or a welfare-maximizing Nash equilibrium. Yet, due to Lemma 3.2 we know that any two Nash equilibria do not differ by more than $\frac{1}{2}\sum_{i\in\mathcal{Z}^C}\pi^i$ for the restricted C-PDOP model with two players. Thus, implementing the IBR does not result in arbitrarily bad Nash solutions in such cases. Though operators do not necessarily find the "best" Nash equilibrium, we guarantee that, for two operators, the Nash equilibrium found using IBR is at least as good (for either player) than the optimistic quasi-monopolistic strategy.

Theorem 3.2 (Monotonicity of Profits (IBR vs. QMO))**.** *For the restricted C-PDOP model with two players, the gross profit of each player at the Nash equilibrium reached by IBR, starting from the* $\langle 0, \ldots, 0 \rangle$ *strategy, is at least as high as her gross profit at the optimistic quasi-monopolistic strategy.*

Thus, operators who currently calculate their routes using the optimistic quasi-monopolistic strategy can only benefit from calculating the Nash equilibrium. Conversely, how-

ever, we can generate instances in which the pessimistic quasi-monopolistic strategy beats the IBR solution.

Since the restricted C-PDOP is a congestion game, IBR terminates in a finite number of iterations. If one aims at implementing the IBR in practice, the number of iterations required to reach an equilibrium is critical. General congestion games with an arbitrary number of players belong to the class of PLS-hard games. Thus, even if there exists a polynomial time algorithm for finding pure strategy Nash equilibria, the solution cannot in general be found in polynomial time by myopic players (Ackermann et al., 2008; Fabrikant et al., 2004). The C-PDOP, however, differs from general congestion games: the number of players is low. By presenting an upper bound on the number of iterations, we show that, assuming homogeneous payoffs, unit-demand stations and indifferent customer choice, the IBR for the C-PDOP does not require full enumeration of all strategies.

Theorem 3.3 (Termination of the Iterated Best Response Algorithm)**.** *For the two-player restricted C-PDOP model, the IBR terminates after one player played at most half of her strategies (thus, the maximum number of required recalculations is* $|S| + 2$ *instead of the* $|S|^2$ *we have in the FE algorithm).*

However, the upper bound on the runtime remains substantial, as the number of strategies is exponential in the number of competitive stations, and obviously, no bound can be given in the general model, since pure strategy Nash equilibria might not exist.

The above results all require exact responses, that is, finding an optimal solution to the PDOP, rather than a near-optimal feasible solution. However, some results also will still hold when we relax the notion of optimality.

Theorem 3.4 (Iterated Best Response Algorithm for Approximate Nash Equilibria)**.** *Let* $s_n = (\boldsymbol{a}^n, \boldsymbol{x}^n)$ *denote a* full *strategy of player* n*, i.e., describing the servicing and routing decisions for all arcs and locations. For any* $\epsilon > 0$*, consider the following version of* ϵ*-approximate Iterated Best Response (*ϵ*-IBR): Given a player* n *and a full strategy* $\mathbf{s}_{-n}$ *of the other players, player* n *can compute a full strategy* s_n *with the property that*

$$(1 + \epsilon)\Pi_n(s_n, \mathbf{s}_{-n}) \geq \max_{s'_n} \Pi_n(s'_n, \mathbf{s}_{-n});$$

note that this amounts to finding an approximate solution of the PDOP problem in equations (3.1a)-(3.1k)*. Player* n *can then choose to deviate to* s_n *if this improves her gross profit.*

For the restricted C-PDOP model, ϵ-IBR always terminates after a finite number of iterations, and the final joint strategy $\mathbf{s}$ *is an ϵ-Nash equilibrium, in the sense that for any player n,*

$$(1+\epsilon)\,\Pi_n\left(s_n, \mathbf{s}_{-n}\right) \geq \max_{s'_n}\ \Pi_n(s'_n, \mathbf{s}_{-n}).$$

Though no polynomial-time approximation of the PDOP is known, we can use this result to get *a-posteriori* bounds on the quality of approximate equilibria. For example, if a commercial solver that implements branch-and-cut procedures obtains a solution with a provable optimality gap of $(1+\epsilon)$, this will immediately imply an $(1+\epsilon)$ guarantee on the quality of the equilibrium.

3.3.2 Potential Function Optimizer

Another approach towards finding a Nash equilibrium in congestion games is optimizing the potential function which can be optimized using a standard solver. This is a global function that captures the local incentives for players to change their strategies, and therefore a useful tool for analyzing equilibria. In particular, if a game admits a potential function, the Nash equilibria of the game coincide with the local optima of the potential function.

Since, in the general C-PDOP model, Nash equilibria are not guaranteed to exist, a potential function cannot be defined. However, for the restricted C-PDOP, it is possible to define the potential Φ as the sum

$$\Phi(\mathbf{q}) = \sum_{i \in \mathcal{Z}} H_{y^i} \pi^i - \sum_n C_n(\mathbf{q}_n), \tag{3.5}$$

where y^i is the number of operators servicing location i, π^i is the revenue that can be extracted from location i, and $H_k = 1 + 1/2 + \cdots + 1/k$ is the k-th harmonic number. Notice that the potential is not the same as the social welfare (3.4): The potential function associates a payoff of $H_{y^i}\pi^i$ with a competitive locations where y^i players are available, while the social welfare more realistically assumes that only one of them will be able to service the customer. Intuitively, this means that the PFO tends to select equilibria in which multiple operators have a vehicle available at competitive locations.

Although a global optimum of the potential function is always a pure strategy Nash equilibrium, it is not necessarily a welfare-maximizing one. However, we can characterize instances in which both coincide.

Lemma 3.4 (Optimality of the Potential Function Optimizer)**.** *For the restricted C-PDOP model (which has a well-defined potential function* Φ*),*

1. *Any potential function maximizer* $\max_{\mathbf{s}} \Phi(\mathbf{q})$ *is a Nash equilibrium.*
2. *If the PFO returns a Nash equilibrium in which at most one operator has a vehicle available at any competitive location* ($\sum_n q_n^i \leq 1 \; \forall i \in \mathcal{Z}^C$)*, this Nash equilibrium is welfare-maximizing.*

Thus, the PFO is likely to return the welfare-maximizing Nash equilibrium if revenues are low, costs are high, and margins are tight, as Nash equilibria mostly do not to contain locations where both operators are present.

3.4 Computational Study

In the following, we quantify the average-case gross profit gains and losses, not only for the restricted C-PDOP, but also for the generality of the model. Unless stated otherwise, we focus on player-homogeneous payoffs, unit-demand stations, indifferent customer choice, and two operators. Further, we conduct a sensitivity analysis if the number of operators increases, if payoffs are not player-homogeneous, if stations are multi-demand with decreasing marginal returns, and if customers are not strictly indifferent between operators. To quantify gains and losses, we present a case study featuring the competition between two major carsharing operators in Munich, Germany.

3.4.1 Experimental Design

We conduct our experiments on a Windows 10 computer restricted to a single 2.60GHz core of an Intel Xeon E7-4860 CPU with 4GB of RAM. We implement both algorithms in Java 10, using CPLEX 12.8 for solving the PDOPs. We start IBR against two different strategies: Assuming that the competitors are absent (IBR-0), or starting against the welfare-maximizing strategy (IBR-WP).

To study the effects of competition and the various business models, we randomly generate 100 data sets for different combinations of the parameters mentioned in Table 3.1. When studying multiplayer settings, inhomogeneous payoffs, multi-demand stations with marginal returns, various different customer choice models, some of the parameters have to be defined slightly different. These changes are introduced at a later stage. The parameter levels are motivated by the Munich carsharing market. For each instance,

we randomly sample locations on a square with an edge length of 15km, and use Euclidean symmetric costs with a weight of 0.62 (25km/h traffic speed, 12,5€/h worker wages (Wittenbrink, 2014), 0.12€/km for fuel). Substitution refers to the share of com-

Parameter	Level 1		Level 2	
Substitution	F	Full Subst.	P	Partial Subst. (25%)
Margin	H	$\pi = 8$	L	$\pi = 4$
Density	H	$\lvert\mathcal{V}_n^+\rvert = 12$	L	$\lvert\mathcal{V}_n^+\rvert = 8$

Table 3.1: Parameters for the Experimental Design in the Base Case

petitive locations, that is if all delivery locations are shared (F), or if shared locations are randomly sampled (P). Albiński and Minner (2020) report that approx. 25% of all customers in Munich have multiple memberships; in the P level, we set 25% of all delivery locations as shared.

To quantify the impact of changes in the payoff/cost structure, we thus vary the payoff. In the high-margin scenario, we set the margin of all delivery and competitive pickup locations to $\pi = 8$ (1 otherwise), and in the low-margin scenario, we set the margin to $\pi = 4$ (0.5 otherwise). The payoffs that can be collected at delivery locations if no vehicle is available constitute the baseline profit all relative improvements are measured towards. Pickup locations are associated with a small (but positive) payoff, due to a low (but non-zero) probability that a customer will rent a car from these locations.

The customer density refers to the number of locations which enter the model (possibly contingent to prior filtering). We generate instances with a high density ($|\mathcal{V}_n^+| = |\mathcal{V}_n^-| = 12$) and instances with a low density ($|\mathcal{V}_n^+| = |\mathcal{V}_n^-| = 8$). The customer density can vary between the operators. In all instances, we use one worker. Further, we assume that all locations are optional, i.e., $S_n^1 = S_n^0 = \emptyset$.

3.4.2 Profit Increase due to Considering the Presence of Competition

Table 3.2 lists relative gross profits (towards the baseline, i.e. no rebalancing) for operator 1 and operator 2 for all parameter combinations if margins of both operators coincide, i.e., are either high or low for both operators. More extensive results can be found in the Appendix in Table 3.9. There, instances are addressed by four consecutive letters referring to substitution, margin (same for both operators), and the density of the first

and second operator. For example, *F_H_L_L* refers to full substitution, high margins, and low densities for either operator. In particular, this allows us to study if the larger or smaller operator benefits more. Averaging over all full competition instances, operators

	Operator 1					Operator 2				
Setting	IBR-0	IBR-WP	PFO	QMO	QMP	IBR-0	IBR-WP	PFO	QMO	QMP
avg (full comp.)	83.6	72.4	71.2	-2.77	40.0	52.2	67.9	69.8	-2.11	39.7
min (full comp.)	5.51	4.88	4.88	-63.2	1.04	3.66	3.93	4.62	-61.7	1.15
max (full comp.)	225	204	206	85.9	180	191	207	215	103	199
avg (part. comp.)	111	110	109	88.3	101	109	111	113	92.2	104
min (part. comp.)	6.74	6.74	6.46	3.45	4.89	8.28	8.28	8.92	4.49	5.54
max (part. comp.)	243	243	240	210	231	246	246	251	226	243

Table 3.2: Average, Minimum and Maximum Percentage Profit Increase towards Baseline (No rebalancing) under Various Models and Algorithms as well as Different Experimental Settings for Either Operator (Operator 1 in Left Block, Operator 2 in Right Block)

are even better off not to rebalance at all than to ignore the competitor (QMO). QMO even generates losses in 5 of 8 instances per operator with full competition, while all other approaches result in non-negative gross profit gains (in all instances in which the other does not have lower margins). Thus, it makes sense to incorporate competition in the routing and servicing decision. However, Nash solutions outperform both quasi-monopolistic solutions. As to be expected, the improvement over QMO/QMP increases in the level of competition, since more locations are shared. The maximum attainable gross profit gain is 251% under partial competition and if the second operator is larger than the first mover (vs. no rebalancing operations). These very high relative values stem from the fact that the baseline and all absolute values are rather low, and that (in particular under full competition) quasi-monopolistic solutions often involve losses. As visible from Table 3.3, the player with more competitive locations can generate higher gross profit gains, when operators have different numbers of vehicles to relocate. This is since the larger operator has more non-shared locations (and can thus build a more efficient route). This effect partially alleviates the disadvantage of the second mover. The smaller operator, however, has the larger benefit of considering competition. This is

	Operator 1					Operator 2				
Setting	IBR-0	IBR-WP	PFO	QMO	QMP	IBR-0	IBR-WP	PFO	QMO	QMP
H_H	113	105	103	47.6	81.1	89.5	103	103	48.4	79.9
H_L	130	125	125	84.6	109	49.9	55.3	59.8	13.9	40.5
L_H	65.6	61.3	60.7	15.3	44.3	121	126	130	91.8	117
L_L	80.8	73	71.7	23.6	47.7	60.9	73.9	73.6	25.9	51.1

Table 3.3: Average Percentage Profit Increase towards Baseline (No Rebalancing) under Various Models and Algorithms if Operators have Different Sizes (Operator 1 in Left Block, Operator 2 in Right Block)

since the larger operator serves most competitive locations, and the small operator benefits from moving her vehicles to the few remaining locations. It also becomes apparent that even though gross profits compared to the baseline increase if the network becomes more dense, the relative benefit of considering competition declines as QMO becomes more profitable. Averaging over all settings with two large operators, gross profit gains over the baseline increase from 79.9% (QMP) to 103% (IBR-WP/PFO) for operator 2 (28.9% increase), whilst for two small operators, the gross profit gain increases from 51.1% to 73.9% (44.6% increase). These high relative values are due to low absolute values, for example, the gross profit gain for operator 1 increases from 0.26 (QMP) to 0.43 (IBR-0) over the baseline in the setting with low density and low margins for both operators. In absolute numbers, however, the benefit of considering competition continues to increase.

We observe that IBR-0 privileges operator 1 over operator 2, while the other algorithms (IBR-WP, PFO, QMO, QMP) do not give a clear advantage to either player. This makes IBR-0 the best algorithm for player 1. For the second player, IBR-0 is outperformed by IBR-WP and PFO in almost all instances. IBR-WP tends to return higher gross profits than PFO for player 2 if both operators have the same size, while PFO tends to return higher gross profits if one player is larger than the other. A similar pattern can be observed with respect to welfare (sum over gross profits of both players). In most instances, QMO returns lower gross profits than QMP, but exceptions exist if the gross profit is low.

3.4.3 Profit Loss due to Presence of Competition

Table 3.4 lists the gross profits of the best found Nash equilibrium (profit-maximizing among IBR-0, IBR-WP and PFO) compared to the welfare-maximizing solution, the coopetition solution, and the monopoly solution. Extended results can be found in Appendix 3.B in Table 3.10. There, the substitution level is addressed in the column header, and every row refers to margin (same for both operators), density of operator 1 and density of operator 2. In general, obviously, all instances follow similar tenden-

	Full Comp.				Part. Comp.			
Setting	NE	W-PDOP	Coop-PDOP	M-PDOP	NE	W-PDOP	Coop-PDOP	M-PDOP
avg.	75.6	82.7	107	110	114	115	171	199
min.	6.89	6.89	11.2	11.2	10.1	10.1	37.7	59.6
max.	155	169	204	209	237	240	304	331

Table 3.4: Average, Minimum and Maximum Percentage Profit Increase towards Baseline (No Rebalancing) for Different Experimental Settings for either Player (Full Substitution in Left Block, Partial Substitution in Right Block)

cies: gross profits increase in the number of vehicles that shall be rebalanced (either due to an increasing customer demand or increasing demand imbalance) and with increasing margins, but decrease if competition increases. While joint fleet management (monopoly or coopetition) results in a substantial gross profit increase, the benefit of welfare-maximization is little (consistently less than 2% under partial competition), and does not justify the additional coordination requirement. We observe a tendency towards the closing of the percentage gap between the Nash solution and the other approaches with increasing instance sizes while absolute gaps continue to grow. This effect is less pronounced in the full competition case, since pooling effects do not improve as much as in the partial competition case. This is mainly driven by better routing decisions due to larger pooling effects in the Nash solution. For Coop-PDOP and M-PDOP, we observe that full competition results in a lower improvement than partial competition. This might seem counter-intuitive at first, but can be explained as follows: In the partial competition case, the benefits of pooling increase as the total number of vehicles is higher. Thus, Coop-PDOP and M-PDOP are more efficient with respect to routing.

In some cases, the routes of the M-PDOP/Coop-PDOP and the Nash equilibrium even coincide. High margins decrease the relative gross profit loss from competing, since many customers are served in the competitive solution, and profit differences must thus be attributed to improved routing (and improved pooling does not contribute as much gross profit as serving additional customers).

3.4.4 Impact of an Increasing Number of Players

All previous experiments consider only two operators. While this is sufficient to model competition in some cities, there are markets with more carsharing operators.

Following the numerical design outlined in Table 3.1, we report average gross profit

		2 Op.			3 Op.			4 Op.			5 Op.		
Setting	Baseline	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP
F_H_H	12	118	-2	61	77	-127	1	57	-171	3	48	-178	2
F_H_L	8	92	-34	34	67	-166	0	43	-212	0	44	-245	0
F_L_H	6	12	-2	2	19	-53	1	19	-81	0	11	-84	0
F_L_L	4	13	-3	5	5	-5	0	4	-1	0	5	-2	0
P1_H_H	12	216	164	203	237	191	221	237	211	231	232	196	221
P1_H_L	8	144	110	135	174	130	161	168	140	158	158	120	150
P1_L_H	6	17	11	12	23	9	19	32	28	30	20	16	13
P1_L_L	4	8	5	4	7	1	5	7	7	5	9	7	6
P2_H_H	12	225	173	213	217	155	205	218	151	208	201	137	195
P2_H_L	8	159	114	147	144	89	137	131	73	121	145	89	135
P2_L_H	6	28	13	19	18	2	12	20	-2	13	21	4	13
P2_L_L	4	10	10	7	5	5	4	6	0	3	4	0	2

Table 3.5: Average Percentage Profit Increase over all Operators towards Baseline (No Rebalancing) for an Increasing Number of Operators

increases over all players for IBR-0, QMO and QMP for an increasing number of operators (column title in Table 3.5), different levels of substitution, different densities, and different margins (all operators have the same density and margin to ensure comparability across different numbers of operators). We observe that effects studied for the

two-operator case get more pronounced, as the number of operators increases, but tendencies remain the same. The gross profit increase towards the baseline under all three models decreases if the number of operators increases and competition is either full, or locations are shared among a subset of operators. This is to be expected since each of the operators services less customers on average. If the number of operators increases, it becomes even more critical to consider competition, as ignoring competition results in significant losses (up to 245% for full competition and 5 operators), and assuming that the competitors service all locations results in refraining from any rebalancing if the number of operators increases. When considering the gain towards assuming that all competitive locations are serviced by the competitors, the improvement of considering competition slightly decreases if the number of operators increases, but remains substantial in all instances.

3.4.5 Impact of Inhomogeneous Payoffs

If payoffs for players are inhomogeneous, i.e., differ between players and locations, pure strategy Nash equilibria do not provably exist. However, in many instances, equilibria appear nonetheless, and can be found using IBR. For two players, we consider full and partial substitution, and either player can have high or low location density, following the numerical design outlined in Table 3.1. We alter the definition of margins, since the case $\pi_i^1 = k \cdot \pi_i^2$ is a special case of homogeneous payoffs with provable existence of pure strategy Nash equilibria. Instead, margins for either player are randomly drawn from a high ($\pi \in [6, 10]$) or low ($\pi \in [2, 6]$) interval. With 100 repetitions of 32 instances, Nash equilibria existed in all cases, which is partially due to the full graph with Euclidean costs. The most central results for inhomogeneous payoffs are depicted in Table 3.6, and all results for 32 different instances are reported in Table 3.11 of Appendix 3.B. Compared to the case of homogeneous payoffs, trends become more pronounced: Gross profit increases over the quasi-monopolistic solutions become larger for full competition (on average, the improvement over the baseline is twice as high for the Nash equilibrium than for QMP), whilst under partial competition, considering competition never improves the solution (compared to very small improvements in the homogeneous payoffs case). The operator with lower contribution margins benefits even more from considering competition, as she can circumvent or at least partially counteract gross profit losses.

Setting	IBR-0 1	IBR-0 2	QMO 1	QMO 2	QMP 1	QMP 2
avg. (full comp.)	167	111	-41	-44	88	83
min. (full comp.)	-26	-47	-237	-239	-77	-66
max. (full comp.)	404	371	236	226	358	365
avg. (part. comp.)	311	289	311	289	311	289
min. (part. comp.)	157	122	157	122	157	122
max. (part. comp.)	434	433	434	433	434	433

Table 3.6: Average Profit Increase towards Baseline (No Rebalancing) Considering Inhomogeneous Payoffs

3.4.6 Impact of Stations with Diminishing Marginal Payoffs

To study the impact of diminishing payoffs in larger stations, we fix the density at a high value (12), and assign the vehicles to stations with varying size $\omega \in [1, 2, 3]$, where 1 is the case without diminishing payoffs. We consider full and partial substitution, and assume that every station is either competitive or non-competitive, but there are no stations with some competitive and some non-competitive locations. In absence of any competition, margins are given by

$$\pi_\iota \left(0, q_n^\iota\right) = \sum_{i=1}^{q_n^\iota} \pi_\iota^\star \cdot \lambda^{(i-1)}$$

and the average margin over all locations in a station is 8 (high, H) or 4 (low, L), which results in a maximum margin of $\pi_\iota^\star = \frac{\pi}{\omega} \sum_{i=1}^{\omega} \lambda^{(i-1)}$ where π is the average margin. If multiple operators have vehicles at station ι, gross profits are split fairly depending on the number of vehicles either operator has at this station. $\lambda \in \{0.5, 0.7, 0.9\}$ is the deterioration rate. Instances are then addressed by substitution (full or partial), margin (high or low), deterioration rate, and station size. For example $F\_H\_0.7\_3$ refers to the instance with full substitution, high margins, medium deterioration rate (0.7), and 3 locations per station. In Table 3.12 (in Appendix 3.B), the number of locations per station is moved to the column head. In total, no Nash equilibrium was found in two cases (out of 100 repetitions of 36 different instances). Both affected instances have 3 locations per station and high margins, but different levels of substitution and deterioration rates. With an increasing number of locations per station, the benefit of considering

competition decreases slightly, as the improvement over the baseline decreases for the Nash equilibrium, but increases for the quasi-monopolistic solutions. The former can be explained by moving not too many vehicles to the same station to achieve high payoffs per location, and the latter occurs, since operators gain some payoff from stations, even if the other operator is also having vehicles there. This effect is not very strong except that QMO does not result in a negative gross profit increase over the baseline if the number of locations exceeds 1. Interestingly, the gap between the two operators' gross profits closes with an increasing number of of locations per station (under full competition it decreases from a factor of two difference to approx. 10% difference), under partial competition and a 3 locations per station, operator 2 can even achieve a higher gross profit than operator 1. This is mostly due to the second operator no longer omitting high payoff stations, and correlates with a higher number of iterations of the IBR.

3.4.7 Impact of Other Customer Choice Behaviors

To establish how much gross profit can be gained if customers do not strictly choose vehicles at random and with equal probability, we generate instances with varying competition, margins, and densities for either operator, and address them in analogy to Section 3.4.2. If both operators are available at a competitive location i, a customer chooses operator 1 with probability $\alpha \in \{0.5, 0.75, 1\}$ where $\alpha = 0.5$ is strict availability-based substitution. 36 different instances are repeated 100 times, and pure strategy Nash equilibria are found in all cases, even though they do not provably exist. In Table 3.13, instances are addressed by substitution (in the column header), margin (same for both), density operator 1, density operator 2, and α. For example, $H\_L\_H\_0.75$ in the column F refers to the instance with full substitution, high margins, small operator 1, large operator 2, and customers preferring operator 1 over 2 (selecting it in 3 out of 4 cases). Table 3.7 gives high-level insights into the trends for other customer choice behaviors. Unsurprisingly, the higher the preference for operator 1, the more will the gross profits of the two operators diverge. If customers have a strict preference for operator 1 ($\alpha = 1$), IBR-0 and QMO coincide for operator 1, and operator 2 always collects at least as high payoffs in the equilibrium as in any of the quasi-monopolistic solutions (she might be able to collect additional revenue at locations, that "do not fit" into operator 1's tour). For operator 1 (the "preferred" operator), the benefit of considering competition, thus, decreases if α increases, whilst for operator 2, the benefit of considering competition increases.

	Operator 1			Operator 2		
Setting	IBR-0	QMO	QMP	IBR-0	QMO	QMP
avg. ($\alpha = 0.5$)	102	58	33	85	57	36
min ($\alpha = 0.5$)	6	-52	0	5	-53	0
max ($\alpha = 0.5$)	254	254	121	260	260	136
avg. ($\alpha = 0.75$)	117	89	105	71	27	1
min ($\alpha = 0.75$)	5	2	0	3	-174	0
max ($\alpha = 0.75$)	262	262	252	255	255	3
avg. ($\alpha = 1$)	119	119	119	70	-3	0
min ($\alpha = 1$)	6	6	6	1	-302	0
max ($\alpha = 1$)	258	258	258	257	257	0

Table 3.7: Average Profit Increase towards Baseline (No Rebalancing) Considering Different Customer Choice

3.4.8 Case Study for Munich Carsharing

To quantify profit gains and losses, we consider a Munich, Germany, case study. We use publicly available data from two Munich carsharing providers containing start and end locations and times for carsharing trips, collected in August 2019. Since the data set does not contain any data about the customers, we assume that all customers have both memberships, and have no preference for one operator over the other. Having both memberships is realistic for frequent users, and thus most trips. The carsharing operators have large fleets of ≈ 500 and ≈ 700 vehicles, respectively. We aggregate trips by assigning them to a start and end district. Districts are hexagons with a radius of ≈ 500m which is commonly assumed to be a reasonable walking distance and provides sufficient flexibility to operators (Ströhle et al., 2019). Districts are approximated by stations at the center of the district. We focus on 21 stations with the highest demand during the observation period (16 days (Mon-Thu) during August 2019). The average demand of operator 1 is 185 trips, and the average demand of operator 2 is 153 trips. First, we count the number of trips starting and ending in every station. The differences between arrivals and departures ("demand imbalance") can be described by a cumulative arrival probability $P\left(\hat{i} \geq k\right)$, i.e., the probability that the k$^{\text{th}}$ vehicle moved to station ι is used. $P\left(\hat{i} \geq k\right)$ is derived from the available data. During this interval, external

influences on demand and supply are sufficiently little, and all remaining differences in demand can be attributed to randomness. $P\left(\hat{i} \geq k\right)$ is independent of n since all customers are shared and have no preference for one operator over the other. We define a (joint) payoff function

$$\pi_{\iota}^{n}\left(q_1^{\iota}, q_2^{\iota}\right) = \frac{q_n^{\iota}}{q_1^{\iota} + q_2^{\iota}} \sum_{k=1}^{q_1^{\iota}+q_2^{\iota}} P\left(\hat{i} \geq k\right) \cdot \pi$$

where π is the contribution margin associated with serving additional customers due to rebalancing a vehicle. We set $\pi = 15$ to account not only for direct revenues of the first customer, but also all future users of that car until it must be rebalanced again, as well as the benefit of preventing customer dissatisfaction (due to a low level of service). We chose this value since the data suggests that vehicles are rebalanced after approx. 10 trips, and customers pay approx. 0.35€ per minute (minus direct costs) and trips often take at least 15 minutes.

The rebalancing costs are calculated by the travel time between the stations (given a velocity of 20km/h), and 5 minutes for additional tasks (e.g., loading/unloading the foldable bike, searching for a parking spot), at an hourly wage of 10€/h; and vehicle cost of 0.3€/km. Thus, the minimum rebalancing cost is $\approx$ 3.4€ (moving back and forth between two stations). We use this minimum rebalancing cost as a bound on the maximum number of vehicles that can be profitably moved to a location. Then, there are 28 delivery locations in 9 stations (with 1-5 locations per station). Of the 26 pickup locations distributed across 10 stations, 13 belong to operator 1 and 13 belong to operator 2. The remaining two stations inherently have balanced demand.

Table 3.8 reports the gross profits of either operator when using IBR, QMO, QMP and M-PDOP to find routes. Since M-PDOP and Coop-PDOP coincide if all delivery

	IBR	QMO	QMP	M-PDOP	$\frac{IBR}{QMO}$	$\frac{IBR}{QMP}$	$\frac{IBR}{M-PDOP}$
Operator 1	35.61	25.39	31.43		1.40	1.13	
Operator 2	35.28	26.96	31.59		1.31	1.12	
Operators 1 & 2	70.89	52.35	63.02	79.18	1.35	1.12	0.90

Table 3.8: Results for the Munich Case Study (Absolute Profits in the Left Block, Relative Gaps in the Right Block)

locations are competitive, we do not report this additionally. The gross profit gains due to considering competition range between 31% and 40% (35% on average over both

operators), compared to assuming that the competitor does not move any vehicles to locations with demand for additional vehicles. The high gains stem from the fact that both operators move their vehicles to the same locations, leaving a demand imbalance of 8 vehicles, even though almost all demand could be served (in equilibrium, no vehicle is moved to three potential customers). However, operators can also reach high profits by assuming their competitor serves all locations of all stations. On average, the operators improve their profits by 12% by considering competition compared to assuming that their competitor is omnipresent. Due to the large fleet and the large stations/districts, the benefit of merging or outsourcing the rebalancing operations is small (≈10%). In conclusion, the price of ignoring competition is very high with 35%, while the price of ignoring the rationality of the competitor is lower with 12% (most likely, operators currently assume some strategy between these two extrema, and the gain due to considering competition will most likely be closer to 12% than to 35%). The price of competition is not too high with 10%.

3.4.9 Algorithmic Performance Results

With respect to performance of the algorithms, we now focus on the following aspects: How large are the instances for carsharing relocation that we can solve, in reasonable time, under the various business models; in particular, can IBR and PFO solve the C-PDOP on real-life-sized instances? How large are the differences in gross profits between Nash equilibria found by the different algorithms?

Size of Solvable Instances

All rebalancing problems under consideration are NP-hard problems (Theorem 3.1). However, we can solve medium-sized instances with up to 50 vehicles that shall be rebalanced. Weikl and Bogenberger (2015) record 36 relocations in Munich during one night for one operator. Munich's one-way, free-floating carsharing fleet is among the largest in the world, thus, the size of solvable instances is most likely sufficient in other cities as well. Further, operators are now reducing the number of necessary relocations by offering incentives for user-based relocation (e.g. Ströhle et al. (2019)), and by increasing the fleet size (e.g. George and Xia (2011)).

In Figure 3.4, we show the average runtimes of 10 instances of increasing size on a logarithmic scale for full and partial substitution (the latter with different substitution rates). The instances are solved with the IBR-0 and the PFO. As the relocation problem

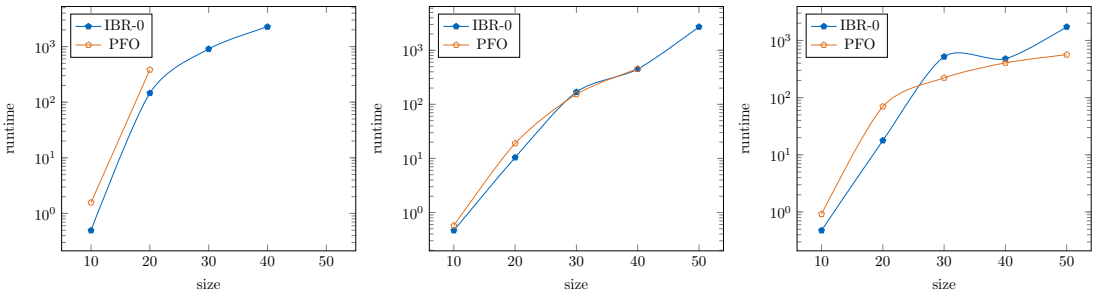

Figure 3.4: Runtime (in seconds) on a Logarithmic Scale for Full Substitution, Partial Substitution (25%), and Partial Substitution (10%)

in carsharing is an operational problem repeated every night, it should not run for more than 30 minutes in the average case (or 10 minutes per iteration in the IBR, since IBR solves most instances in 3-4 iterations). This is possible for instances with up to 50 locations with both algorithms if there is only little substitution. For the IBR, it is still computationally feasible to solve instances with 50 pickup and delivery locations per operator under realistic substitution (25%), but in these instances, we already observe the computational advantage of the IBR-0 over PFO. Under full substitution, both algorithms perform substantially worse, but the IBR-0 can solve instances that are roughly twice as large. Thus, with respect to runtime, the Iterated Best Response is the method of choice, with both algorithms performing roughly the same for low substitution (10%).

Both algorithms only find exact equilibria reliably on comparably small instances. Leveraging Theorem 3.4, we know that provable optimality of the PDOP is not necessary to reach a "sufficiently stable" solution. Thus, Figure 3.5 reports the fraction of instances

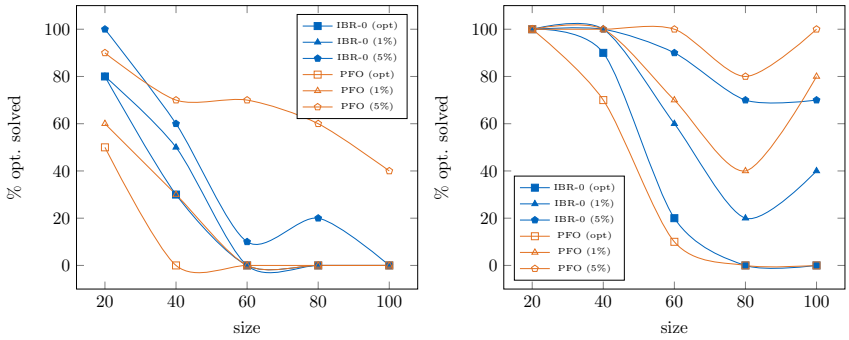

Figure 3.5: Fraction of Instances which Solved to Provable/1%/5% Optimality for Full Substitution and Partial Substitution (25%) using IBR-0 and PFO

of given size which solve (i) to optimality, (ii) to 1% optimality, or (iii) to 5% optimality using IBR-0 and PFO within a time limit of 10 minutes per iteration (IBR-0) or 30

minutes in total (PFO). Similar to the results from Figure 3.4, we observe that IBR-0 and

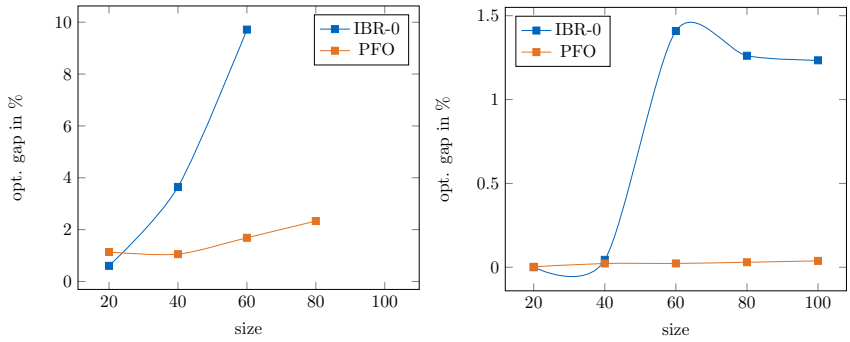

Figure 3.6: Average Optimality Gap for Full and Partial Substitution (25%) using IBR-0 and PFO (if incumbent is found for at least 50% of instances)

PFO solve a similar percentage of small and medium-sized instances. Surprisingly, PFO solves 80% of the instances with 100 customers and vehicles, and partial competition to 1% optimality whilst IBR-0 fails to solve almost any instances with 100 customers and vehicles to 1% optimality (even though the average runtime to optimality is higher for PFO than for IBR). A similar effect can be observed for full competition. The reason for this is that with IBR, terminating with a high optimality gap in any iteration results in a weak approximation guarantee. This is also observable from Figure 3.6, which reports the average gap for PFO, and the average over the worst gap of all iterations using IBR-0 which gets as high as 9.7% for full substitution and 60 locations (only if at least 50% of all instances provided a feasible solution). Also, a deeper look into the branch-and-bound behavior for partial competition reveals that already the first found integer solution often provides a reasonably good bound. Thus, if any feasible solution is found, it frequently already provides a reasonably tight approximation guarantee.

Trade-off due to Equilibrium Selection

Both of the algorithms we presented for the C-PDOP come at a price: Neither of the algorithms provably returns the best Nash equilibrium. In Figure 3.7, we denote the actual gaps between different Nash equilibria. If the number of Nash equilibria increases (derived using full enumeration), we empirically observe the following ordering for the welfare (sum over all gross profits): the average Nash equilibrium (derived using FE) results in lower gross profits than IBR-0, which in turn has lower gross profits than PFO. The profit of PFO is exceeded by IBR-WP, which has a lower gross profit than the best Nash equilibrium (derived using FE). This ordering is inverse to the ordering

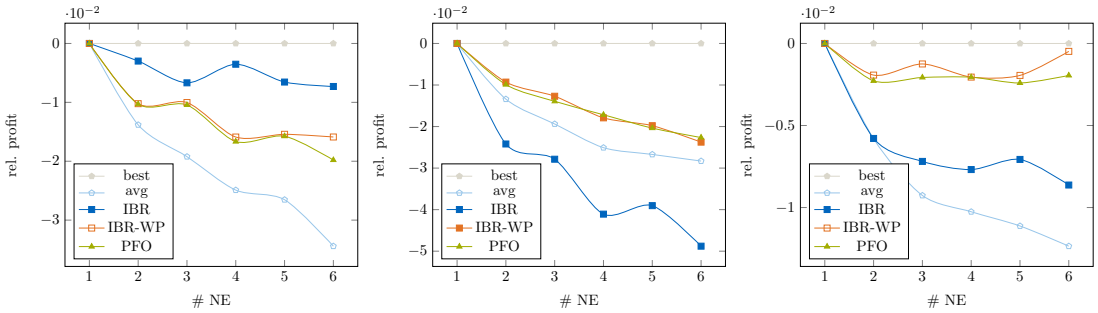

Figure 3.7: Quality of Nash Equilibria (for operators 1, 2 and with respect to welfare respectively)

by runtime. Further, we can see that player 1 (the operator who moves first) profits more from IBR-0 than player 2 (which makes it the second-best and worst algorithm respectively). Thus, there is a tradeoff between solution quality and runtime. If the number of expected Nash equilibria is low, if computation time is a scarce resource, or if any stable solution rather than the best solution is sufficient, IBR-0 is the preferred method, while PFO and IBR-WP (or even FE) are preferred, if margins are low and if solution quality is paramount.

3.5 Conclusion

In this Chapter we study the profitability of the relocation operations under competition in a station-based carsharing system. We present a mathematical model for the relocation problem that arises in carsharing which we call Pickup and Delivery Orienteering Problem (PDOP). We, further, present variations of the PDOP which capture different business models under competition: The C-PDOP models direct competition, while the M-PDOP and the Coop-PDOP model merger/monopoly and coopetition/outsourcing relocation operations, respectively. In the C-PDOP, we introduce competitor who also optimize their fleet and then solve the problem for Nash equilibria, i.e. stable states of the system. The C-PDOP assumes that each operator plans their tour before executing any relocations. In a future line of research, one could investigate a dynamic setting, or multi-stage game, in which operators can change their decision during relocation, as they observe the competitors' moves.

We present two algorithms to find pure strategy Nash equilibria, namely Iterated Best Response (IBR) and Profit Function Optimizer (PFO), both of which are considerably faster than out-of-the-box and brute-force algorithms. Pure strategy Nash equilibria

provably exist if both operators receive the same revenue from servicing a customer ("player-homogeneous payoffs"), stations hold at most one vehicle or customer and demand processes are independent ("unit-demand stations"), and if customers choose a vehicle at random if multiple operators have a vehicle available ("indifferent customer choice"). Even if these assumptions do not hold, we find that equilibria do exist and are reached by IBR in most settings.

We use those model variations to quantify the profitability of the associated business models in competitive carsharing markets. We give two minor conditions which ensure the intuition that monopoly and coopetition result in higher gross profits than direct competition, and show numerically that this gain is substantial. In a Munich, Germany, case study, operators can improve by 10% by merging, or outsourcing to the same third-party service provider. This value is lower than the maximum improvements in the numerical experiments with artificial data, but still substantial. The benefit stems from pooling benefits and the associated improved routing, and indirectly affects the number of serviced customers. Thus, these business models are also preferential with respect to society, due to a lowered environmental impact and increased service levels.

Though there are examples in which considering competition is worse than ignoring the presence of competition for some operators, we show numerically that profitability for all operators increases in realistic settings. If margins are low, this improvement (or vice versa, the *cost of ignorance*) can be up to several orders of magnitude. The main drivers for a high cost of ignorance or benefit of considering competition are fierce competition, a high number of operators, inhomogeneous payoffs, not too large stations, and customer preferences for one operator. In a case study, the gross profit improvement due to considering competition is 35% over assuming that no competition exists, and 12% over assuming that competition is omnipresent.

The more candidate locations there are, the more important relocation becomes, as routes become more efficient. The more of these locations are shared, the more important it becomes to consider competition. We observe that operators might be worse off by ignoring the presence of competition in their routing decision than not relocating any vehicles, and might even lose money, in particular if competition is fierce. For each of the three assumptions (player-homogeneous payoffs, unit-demand stations, and indifferent customer choice), we show that lifting the assumption results in similar tendencies if equilibria exist. Equilibria exist in many realistic instances. Though the studied algorithms do not necessarily find the best equilibrium, we show numerically that they still yield higher gross profits than solutions which do not consider competition. Hence,

operators have an incentive to adopt and implement game-theoretic strategies in their relocation decisions.

This model can also be applied to free-floating systems if multiple nearby locations are clustered into a district, and demand prediction is performed for these districts. For the model to be applicable as is, distances, and thus travel time and cost, must be measured between these districts. That is only realistic if districts are sufficiently small.

Appendix 3.A Proofs

3.A.1 Proof of Theorem 3.1 (NP-Hardness)

We shall prove hardness even for the restricted C-PDOP model, which implies hardness for the general model as well. To prove point 1, we reduce from the traveling salesman problem (TSP), which is well-known to be NP-hard (Karp, 1972). The NP-hardness of TSP on bipartite instances follows from Krishnamoorthy, 1975, whereas the NP-hardness of the prize collecting TSP (on general graphs) follows from Feillet et al., 2005. Since our setting combines both modifications, we provide for completeness an NP-hardness proof, adapting the reduction techniques in those papers as well as in Volgenant and Jonker, 1987.

Given an instance of TSP comprised of a complete graph K_n with vertices labeled $1, \ldots, n$, and arc routing costs c_e for $e \in E_n$, construct a PDOP instance G with $4n+2$ locations, one worker, and $S^0 = S^1 = \emptyset$ (all locations are optional) as follows. For each node $i \in K_n$ we include four copies i^0, i^1, i^2, i^3 in G, such that $i^0, i^2 \in \mathcal{Z}^-$ and $i^1, i^3 \in \mathcal{Z}^+$ (if the costs are known to satisfy the triangular inequality, for example in the Euclidean TSP, then the TSP remains NP-hard, and a somewhat simpler reduction can be used where only two copies i, i' of each node are required.). We further include additional depot nodes d and p. Fix a node $1 \in K_n$; then arc $\langle dj^3 \rangle$ has cost $c_{\langle 1j \rangle}$ for each $j \neq 1$; further, arc $\langle 1^0 p \rangle$ has cost 0. For each arc $\langle ij \rangle \in A$, the corresponding arc $\langle i^0 j^3 \rangle$ has cost $c_{\langle ij \rangle}$. For each node $i \in K_n$, the arcs $\langle i^0 i^1 \rangle$, $\langle i^2 i^1 \rangle$ and $\langle i^2 i^3 \rangle$ have cost 0. Any other arc has cost Cn where $C = \max_{e \in E_n} c_e$. As far as profits go, for each node $i \in K_n$, the corresponding node i^2 has profit Cn. Any other node has profit 0. Finally, set all travel times to be 0 and an arbitrary positive time window T, so that we can drop the restrictions on the travel time of the worker.

To finish the proof of the reduction, simply observe that a tour $(i_1 i_2 \ldots i_n i_1)$ in K_n having cost L, with $i_1 = 1$, can be lifted into a d-p-path $(d\,[i_2] \ldots [i_n]\,[i_1]\,p)$ in G, where $[i]$ denotes the sequence $i^3 i^2 i^1 i^0$. This path has profit $Cn^2 - L \in [Cn(n-1), Cn^2]$. Moreover, any d-p-path not of this form has profit at most $Cn(n-1)$. Thus K_n admits a tour of cost at most L if and only if G admits a d-p-path of profit at least $Cn^2 - L$.

To prove point 2, we observe that a PDOP instance (for one player) is a special case of a C-PDOP instance (with two players) in which the other player routes on a trivial graph and $\mathcal{Z}^C = \emptyset$. To prove points 3 and 4, we observe that a PDOP instance is a special case of a C-PDOP instance in which all locations are competitive but the distances from depot nodes d_2, p_2 to the rest of the graph are prohibitively large (so that player 2's best

strategy is to play the empty strategy $\langle 0, \ldots, 0 \rangle$ and the Nash equilibria correspond to player 1's best responses to this strategy).

3.A.2 Proof of Lemma 3.1 (Congestion Game)

Given a restricted C-PDOP instance with c locations (and n players), we construct a congestion game with at most $c+n2^c$ resources (and also with n players). The congestion game is described by the following components: a set of common resources R; a payoff function p_r for each resource $r \in R$; and a set of valid strategies ξ^n for each player, where each strategy in ξ^n is a subset of R.

Recall that, from the discussion on Section 3.2.2, we can specify the payoff function at each location i by a single value π^i that represents the payoff that any player could extract by being the sole operator at i. For each location $i \in \mathcal{Z}$ we include a resource, also denoted $i \in R$, with payoff function

$$p_i(y^i) = \frac{\pi^i}{y^i},$$

where y^i is the number of players having a vehicle available at location i. Moreover, for each player n and each (compact) strategy $\mathbf{q}$ of that player in the original C-PDOP instance, our congestion game includes a resource, which is denoted $(\mathbf{q}, n) \in R$, with constant, negative payoff function

$$p_{(\mathbf{q},n)} = -C_n(\mathbf{q});$$

these can be in theory obtained by solving the PDOP instance described in equations (3.1a)-(3.1k), while setting fixed the variables corresponding to locations.

A valid strategy of player n consists of exactly one resource of the second type, $(\mathbf{q}, n)$, and all associated locations i such that $q^i = 1$; that is, the set ξ_n of valid strategies is given by

$$\xi_n = \{\{(\mathbf{q}, n)\} \cup \{i : q^i = 1\} : \mathbf{q} \text{ is a compact strategy of player} n\}.$$

In this way, we obtain a valid formulation of a congestion game, as laid out by Rosenthal, 1973. The profit of a player n playing strategy $x_n = \{(\mathbf{q}_n, n)\} \cup \{i : q^i_n = 1\}$ is defined as $\sum_{r \in \{(\mathbf{q}_n, n)\} \cup \{i : q^i_n = 1\}} p_r(y^r)$, where y^r is the number of players accessing resource r. There are however two key differences from the usual formulation: players are maxi-

mizing payoffs instead of minimizing costs; and payoffs may assume positive or negative values. These differences are without loss of generality since the standard potential argument can still be applied, as we shall show in the proof of Corollary 3.1.

The (compact) strategies $\mathbf{q}$ in the standard C-PDOP instance are in one-to-one correspondence to the valid strategies $\mathbf{x}$ in the congestion game, where $x_n = \{(\mathbf{q}_n, n)\} \cup \{i : q_n^i = 1\}$. This correspondence preserves profits; if P_n is the profit function of player n in the congestion game, then

$$\begin{aligned} P_n(\mathbf{x}) &= \sum_{r \in \{(\mathbf{q}_n, n)\} \cup \{i : q_n^i = 1\}} p_r(y^r) = \sum_{i : q_n^i = 1} \frac{\pi^i}{y^i} + \pi^{(\mathbf{q}_n, n)} \\ &= \sum_{i : q_n^i = 1} \pi^i(\mathbf{q}^i) - C_n(\mathbf{q}_n) = R_n(\mathbf{q}) - C_n(\mathbf{q}) = \Pi_n(\mathbf{q}). \end{aligned}$$

Thus, a deviation is improving in the C-PDOP instance if and only if it is improving in the congestion game. It follows that $\mathbf{q}$ is a Nash equilibrium for the C-PDOP instance if and only if the corresponding strategy $\mathbf{x}$ is a Nash equilibrium for the congestion game.

3.A.3 Proof of Corollary 3.1 (Existence of Pure Strategy Nash Equilibria)

We apply Rosenthal, 1973's potential argument. For a given strategy profile $\mathbf{q}$, let $y^i = y^i(\mathbf{q})$ denote the number of operators placing a vehicle at location i. Define the potential function

$$\Phi(\mathbf{s}) = \sum_{i \in \mathcal{Z}} \sum_{q=1}^{y^i} p_i(q) - \sum_n C_n(\mathbf{q}_n) = \sum_{i \in \mathcal{Z}} H_{y^i} \pi^i - \sum_n C_n(\mathbf{q}_n),$$

where $H_{y^i} = 1 + \frac{1}{2} + \cdots + \frac{1}{y^i}$ denotes the harmonic number of order y^i.

Next observe that the potential function keeps track of the changes in profit when a player deviates. For example, suppose player n deviates from strategy $\mathbf{q}_n$ to $\tilde{\mathbf{q}}_n$ while the other players keep to strategy $\mathbf{q}_{-n}$, and let $\tilde{y}^i$ denote the number of vehicles at location i in the strategy profile $(\tilde{\mathbf{q}}_n, \mathbf{q}_{-n})$. We will prove that the change in potential equals the change in the profit of player n:

$$\Phi(\tilde{\mathbf{q}}_n, \mathbf{q}_{-n}) - \Phi(\mathbf{q}_n, \mathbf{q}_{-n})$$

$$
\begin{aligned}
&= \sum_{i\in\mathcal{Z}} \sum_{q=1}^{\tilde{y}^i} p_i(q) - C_n(\tilde{\mathbf{q}}_n) - \sum_{n'\neq n} C_{n'}(\mathbf{q}_{n'}) - \left(\sum_{i\in\mathcal{Z}} \sum_{q=1}^{y^i} p_i(q) - C_n(\mathbf{q}_n) - \sum_{n'\neq n} C_{n'}(\mathbf{q}_{n'}) \right) \\
&= \sum_{i:q_n^i=0,\tilde{q}_n^i=0} \left(\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) \right) + \sum_{i:q_n^i=1,\tilde{q}_n^i=1} \left(\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) \right) \\
&\quad + \sum_{i:q_n^i=0,\tilde{q}_n^i=1} \left(\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) \right) + \sum_{i:q_n^i=1,\tilde{q}_n^i=0} \left(\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) \right) \\
&\quad - C_n(\tilde{\mathbf{q}}_n) + C_n(\mathbf{q}_n) \\
&= \sum_{i:q_n^i=1,\tilde{q}_n^i=1} p_i(\tilde{y}^i) - \sum_{i:q_n^i=1,\tilde{q}_n^i=1} p_i(y^i) + \sum_{i:q_n^i=0,\tilde{q}_n^i=1} p_i(\tilde{y}^i) - \sum_{i:q_n^i=1,\tilde{q}_n^i=0} p_i(y^i) \\
&\quad - C_n(\tilde{\mathbf{q}}_n) + C_n(\mathbf{q}_n) \\
&= \sum_{i:\tilde{q}_n^i=1} p_i(\tilde{y}^i) - C_n(\tilde{\mathbf{q}}_n) - \left(\sum_{i:q_n^i=1} p_i(y^i) - C_n(\mathbf{q}_n) \right) \\
&= \Pi_n(\tilde{\mathbf{q}}_n, \mathbf{q}_{-n}) - \Pi_n(\mathbf{q}_n, \mathbf{q}_{-n}).
\end{aligned}
$$

In the second equality, we split the sums over $i \in \mathcal{Z}$ into four cases, depending on whether each of q_n^i, $\tilde{q}_n^i$ is 0 or 1. In the third equality, we observe that

- if $q_n^i = 0, \tilde{q}_n^i = 0$, then $\tilde{y}^i = y^i$ and $\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) = 0$;
- if $q_n^i = 1, \tilde{q}_n^i = 1$, then $\tilde{y}^i = y^i$ and $\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) = 0 = p_i(\tilde{y}^i) - p_i(y^i)$;
- if $q_n^i = 0, \tilde{q}_n^i = 1$, then $\tilde{y}^i = y^i + 1$ and $\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) = p_i(\tilde{y}^i)$;
- if $q_n^i = 1, \tilde{q}_n^i = 0$, then $\tilde{y}^i = y^i - 1$ and $\sum_{q=1}^{\tilde{y}^i} p_i(q) - \sum_{q=1}^{y^i} p_i(q) = -p_i(y^i)$.

We conclude that a deviation by any player is improving if and only if the potential increases. Since there are finitely many strategies, Φ possesses a global maximum. The corresponding strategy must be a Nash equilibrium since no player's deviation would increase the value of the potential function and thus not increase her profit. Moreover, if $\mathbf{q}^{(0)} \to \mathbf{q}^{(1)} \to \ldots$ is a sequence where $\mathbf{q}^{(n+1)}$ is obtained from $\mathbf{q}^{(n)}$ by an improving deviation from one of the players, then the potential must strictly increase through the sequence. Such a sequence must then terminate at a local maximum, which is a Nash equilibrium.

3.A.4 Proof of Corollary 3.2 (Monotonicity of Profits (vs. Welfare-Maximizing Solution))

By definition, the welfare-maximizing solution is the solution to the PDOP with two competing operators in which the joint profit is maximal. Thus, no other solution, including any Nash equilibrium solution, can be better.

3.A.5 Proof of Lemma 3.2 (Price of Anarchy and Price of Stability)

Consider an instance of the restricted C-PDOP model with two players. We first derive an useful relation between the profits of a player when the other player changes her strategy. Let $\mathbf{q}_1$ be any strategy for player 1 and $\mathbf{q}_2, \mathbf{q}'_2$ be any two strategies for player 2. By definition of the profit function we have

$$\Pi_1(\mathbf{q}_1, \mathbf{q}_2) = \sum_{i:q_1^i=1,q_2^i=0} \pi^i + \sum_{i:q_1^i=1,q_2^i=1} \frac{\pi^i}{2} - C_1(\mathbf{q}_1);$$
$$\Pi_1(\mathbf{q}_1, \mathbf{q}'_2) = \sum_{i:q_1^i=1,q_2'^i=0} \pi^i + \sum_{i:q_1^i=1,q_2'^i=1} \frac{\pi^i}{2} - C_1(\mathbf{q}_1);$$

putting these two equations together we see that

$$\Pi_1(\mathbf{q}_1, \mathbf{q}_2) - \Pi_1(\mathbf{q}_1, \mathbf{q}'_2) = \frac{1}{2} \sum_{i:q_1^i=1,q_2'^i=1,q_2^i=0} \pi_i - \frac{1}{2} \sum_{i:q_1^i=1,q_2'^i=0,q_2^i=1} \pi_i. \tag{3.6}$$

To prove point 1, let $(\mathbf{q}_1, \mathbf{q}_2)$ be any Nash equilibrium and $(\mathbf{q}'_1, \mathbf{q}'_2)$ be any strategy. Applying equation (3.6) and the definition of Nash equilibrium,

$$\begin{aligned}\Pi_1(\mathbf{q}'_1, \mathbf{q}'_2) &= \Pi_1(\mathbf{q}'_1, \mathbf{q}_2) + \frac{1}{2} \sum_{i:q_1'^i=1,q_2^i=1,q_2'^i=0} \pi_i - \frac{1}{2} \sum_{i:q_1'^i=1,q_2^i=0,q_2'^i=1} \pi_i \\ &\leq \Pi_1(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2} \sum_{i:q_1'^i=1,q_2^i=1,q_2'^i=0} \pi_i - \frac{1}{2} \sum_{i:q_1'^i=1,q_2^i=0,q_2'^i=1} \pi_i\end{aligned}$$

Similarly for player 2, we have

$$\Pi_2(\mathbf{q}'_1, \mathbf{q}'_2) \leq \Pi_2(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2} \sum_{i:q_2'^i=1,q_1^i=1,q_1'^i=0} \pi_i - \frac{1}{2} \sum_{i:q_2'^i=1,q_1^i=0,q_1'^i=1} \pi_i;$$

Adding both equations, and observing that the two sets $\{i : q_1'^i = 1, q_2^i = 1, q_2'^i = 0\}$ and

$\{i : q_2'^i = 1, q_1^i = 1, q_1'^i = 0\}$ are disjoint subsets of $\mathcal{Z}^C$, as well as that $\pi_i \geq 0$, we obtain

$$\begin{aligned}&\Pi(\mathbf{q}_1', \mathbf{q}_2') = \Pi_1(\mathbf{q}_1', \mathbf{q}_2') + \Pi_2(\mathbf{q}_1', \mathbf{q}_2') \\ &\leq \Pi_1(\mathbf{q}_1, \mathbf{q}_2) + \Pi_2(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i \in \mathcal{Z}^C} \pi_i = \Pi(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i \in \mathcal{Z}^C} \pi_i.\end{aligned}$$

Since $(\mathbf{q}_1', \mathbf{q}_2')$ was taken to be any strategy, the above equation implies that the difference in welfare between any two Nash equilibria is at most $\frac{1}{2}\sum_{i \in \mathcal{Z}^C} \pi_i$, which proves point 2. Now let $(\mathbf{q}_1, \mathbf{q}_2)$ be any Nash equilibrium and $(\mathbf{q}_1^*, \mathbf{q}_2^*)$ be a welfare maximizing strategy. We get that

$$\Pi(\mathbf{q}_1, \mathbf{q}_2) \leq \Pi(\mathbf{q}_1^*, \mathbf{q}_2^*) \leq \Pi(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2}\sum_{i \in \mathcal{Z}^C} \pi_i,$$

thus the difference between any Nash equilibrium and any welfare maximizing strategy is at most $\frac{1}{2}\sum_{i \in \mathcal{Z}} \pi_i$, proving point 3.

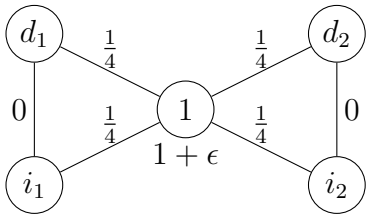

	0	1
0	$0, 0$	$0, \frac{1}{2}+\epsilon$
1	$\frac{1}{2}+\epsilon, 0$	$\frac{\epsilon}{2}, \frac{\epsilon}{2}$

	0	1
0	$0, 0$	$0, \frac{1}{2}$
1	$\frac{1}{2}, 0$	$0, 0$

Figure 3.8: Example of a C-PDOP instance (left); there is only one delivery location, labeled 1, which is competitive; each player has a depot (d_1 or d_2) and one vehicle at a separate location (i_1 or i_2) with null payoff; the corresponding payoff matrix (center) and the special case $\epsilon = 0$ (right).

To prove that the Price of Anarchy and the Price of Stability can be arbitrarily high, consider the game depicted in Figure 3.8. There is only one competitive location with payoff $\pi = 1 + \epsilon$ for both players, for some small positive ϵ. Either player incurs a traveling cost of $c = 1/2$ to relocate a vehicle to the competitive location and return to the depot. Thus each player has only two strategies (1 for 'move' or 0 for 'don't move'). Looking at the payoff matrix, we see that the only Nash equilibrium occurs when both players decide to service the location, for a welfare of ϵ, whereas the maximum possible welfare is $\frac{1}{2} + \epsilon$, occuring when only one player services the location. Thus both the Price of Anarchy and the Price of Stability equal $\frac{1/2+\epsilon}{\epsilon} = 1 + \frac{1}{2\epsilon}$ which is arbitrarily large as ϵ can be arbitrarily small. Since the C-PDOP generalizes the restricted C-PDOP, the

above result also carries over to the C-PDOP.

To show the bounds obtained in points 1, 2, and 3 are tight, consider the same game as before, but take $\epsilon = 0$. Now both strategies $(q_1, q_2) = (1, 0)$ or $(0, 1)$ are welfare-maximizing as well as Nash equilibria, achieving welfare $\frac{1}{2} = \frac{1}{2}\sum_{i \in \mathcal{Z}^C} \pi_i$. Moreover, $(q_1, q_2) = (1, 1)$ is still a Nash equilibrium with welfare exactly equal to 0. Thus the absolute difference between the welfare-maximizing strategy and the worst Nash equilibrium, as well as between best and worst Nash equilibria, is exactly $\frac{1}{2}\sum_{i \in \mathcal{Z}^C} \pi_i$.

3.A.6 Proof of Corollary 3.3 (NP-hardness of M-PDOP and Coop-PDOP)

A PDOP instance is a special case of a Coop-PDOP (resp. M-PDOP) instance (derived from two players) in which the second player routes on a trivial graph and $\mathcal{Z}^C = \emptyset$. Thus, both problems remain NP-hard.

3.A.7 Proof of Lemma 3.3 (Monotonicity of Profits (vs. Monopoly or Coopetition Solution))

To prove points 1 and 2, simply observe that the space of feasible solutions increases as we move from C-PDOP to Coop-PDOP to M-PDOP. In other words, a feasible strategy (x^1, a^1, x^2, a^2) with W_1, W_2 workers in the C-PDOP model can be merged into a feasible strategy $(\bar{x}, \bar{a})$ with $W_1 + W_2$ workers in the Coop-PDOP model; and a feasible strategy (x, a) with W_O workers in the Coop-PDOP model is also a feasible strategy (with the same number of workers) in the M-PDOP model. Moreover, the payoffs associated to competitive locations only increase as we move from C-PDOP to Coop-PDOP to M-PDOP. To see this, let ι be a station and $\mathbf{q}^\iota$ represent the vehicles of each operator at station ι. The joint payoff at location ι is the same for the C-PDOP and Coop-PDOP models, and equals $\sum_n \pi_n^\iota(\mathbf{q}^\iota)$. This, in turn, is less or equal than the joint payoff for the M-PDOP model, which is defined as $\max_{\mathbf{q}'} \left(\sum_n \pi_n^\iota\left(\mathbf{q}'\right) \mid \sum_n q_n^\iota = \sum_n q_n^{\iota\prime}\right)$. Therefore, the optimal profit does not decrease as we move from C-PDOP to Coop-PDOP to M-PDOP (as long as the number of workers is consistent, i.e. $W_M \geq W_O \geq W_1 + W_2$).

3.A.8 Proof of Theorem 3.2 (Monotonicity of Profits (IBR vs. QMO))

We begin by deriving the following relation between the potential and payoff functions, valid for the restricted C-PDOP with two players:

$$\begin{aligned}\Phi(\mathbf{q}_1,\mathbf{q}_2) &= \sum_{i\in\mathcal{Z}} H_{y^i}\pi^i - C_1(\mathbf{q}_1) - C_2(\mathbf{q}_2)\\ &= \sum_{i:q_1^i=1,q_2^i=0}\pi^i + \sum_{i:q_1^i=0,q_2^i=1}\pi^i + \sum_{i:q_1^i=1,q_2^i=1}\frac{3}{2}\pi^i - C_1(\mathbf{q}_1) - C_2(\mathbf{q}_2)\\ &= \left(\sum_{i:q_1^i=1,q_2^i=0}\pi^i + \sum_{i:q_1^i=1,q_2^i=1}\frac{\pi^i}{2} - C_1(\mathbf{q}_1)\right) + \left(\sum_{i:q_2^i=1}\pi^i - C_2(\mathbf{q}_2)\right)\\ &= \Pi_1(\mathbf{q}_1,\mathbf{q}_2) + \Pi_2(\langle 0,\dots,0\rangle,\mathbf{q}_2); \end{aligned} \tag{3.7}$$

similarly, one has $\Phi(\mathbf{q}_1,\mathbf{q}_2) = \Pi_1(\mathbf{q}_1,\langle 0,\dots,0\rangle) + \Pi_2(\mathbf{q}_1,\mathbf{q}_2)$.

Consider a sequence of iterated best responses starting from the $\langle 0,\dots,0\rangle$ strategy and ending at a Nash equilibrium, with player 1 moving first into an optimistic strategy:

$$(\langle 0,\dots,0\rangle,\langle 0,\dots,0\rangle) \xrightarrow{1} (\mathbf{q}_1^O,\langle 0,\dots,0\rangle) \xrightarrow{2} (\mathbf{q}_1^O,\mathbf{q}_2^{(1)}) \to \cdots \to (\mathbf{q}_1^N,\mathbf{q}_2^N).$$

Since $\mathbf{q}_2^{(1)}$ is a best response to player 1 playing $\mathbf{q}_1^O$, the following is also a sequence of iterated best responses:

$$(\mathbf{q}_1^O,\mathbf{q}_2^O) \xrightarrow{2} (\mathbf{q}_1^O,\mathbf{q}_2^{(1)}) \to \cdots \to (\mathbf{q}_1^N,\mathbf{q}_2^N).$$

In particular, this implies that the potential value at the Nash equilibrium retrieved through IBR is at least the potential value at the optimistic quasi-monopolistic strategy, that is,

$$\Phi(\mathbf{q}_1^N,\mathbf{q}_2^N) \geq \Phi(\mathbf{q}_1^O,\mathbf{q}_2^O).$$

Note that the optimistic strategy maximizes a player's profit with respect to the other player playing the empty strategy; in particular, we have $\Pi_1(\mathbf{q}_1^O,\langle 0,\dots,0\rangle) \geq \Pi_1(\mathbf{q}_1^N,\langle 0,\dots,0\rangle)$ and $\Pi_2(\langle 0,\dots,0\rangle,\mathbf{q}_2^O) \geq \Pi_2(\langle 0,\dots,0\rangle,\mathbf{q}_2^N)$. Using (3.7), we get

$$\begin{aligned}\Pi_1(\mathbf{q}_1^N,\mathbf{q}_2^N) &= \Phi(\mathbf{q}_1^N,\mathbf{q}_2^N) - \Pi_2(\langle 0,\dots,0\rangle,\mathbf{q}_2^N)\\ &\geq \Phi(\mathbf{q}_1^O,\mathbf{q}_2^O) - \Pi_2(\langle 0,\dots,0\rangle,\mathbf{q}_2^O)\end{aligned}$$

$$= \Pi_1(\mathbf{q}_1^O, \mathbf{q}_2^O).$$

Using a similar reasoning for player 2, we get that $\Pi_2(\mathbf{q}_1^N, \mathbf{q}_2^N) \geq \Pi_2(\mathbf{q}_1^O, \mathbf{q}_2^O)$. Thus both players are better off if they agree on a Nash equilibrium obtained by iterated best responses from the empty strategy.

3.A.9 Proof of Theorem 3.3 (Termination of the Iterated Best Reponse Algorithm)

We first define the inverse $\neg\mathbf{q}$ of a strategy $\mathbf{q}$ as the strategy in which a vehicle is available at precisely the competitive locations where $\mathbf{q}$ does not have a vehicle available (formally, for every $i \in \mathcal{Z}^C$ we have that $\neg q^i = 1 - q^i$). To prove this theorem we need the following key auxiliary result: for any strategies $\mathbf{q}_1$ and $\mathbf{q}_2$, if $\mathbf{q}_1$ is a best response to $\mathbf{q}_2$, then $\mathbf{q}_1$ is also a best response to its inverse $\neg\mathbf{q}_1$. To see this, let $\mathbf{q}_1'$ be an arbitrary strategy for player 1; applying Equation (3.6) twice,

$$\begin{aligned}
\Pi_1(\mathbf{q}_1', \neg\mathbf{q}_1) =& \Pi_1(\mathbf{q}_1', \mathbf{q}_2) + \frac{1}{2} \sum_{i:q_1'^i=1,q_2^i=1,\neg q_1^i=0} \pi_i - \frac{1}{2} \sum_{i:q_1'^i=1,\neg q_1^i=1,q_2^i=0} \pi_i \\
\leq& \Pi_1(\mathbf{q}_1, \mathbf{q}_2) + \frac{1}{2} \sum_{i:q_1'^i=1,q_2^i=1,q_1^i=1} \pi_i - \frac{1}{2} \sum_{i:q_1'^i=1,q_1^i=0,q_2^i=0} \pi_i \\
=& \Pi_1(\mathbf{q}_1, \neg\mathbf{q}_1) + \frac{1}{2} \sum_{i:q_1^i=1,q_1^i=0,q_2^i=0} \pi_i - \frac{1}{2} \sum_{i:q_1^i=1,q_2^i=1,q_1^i=1} \pi_i \\
&+ \frac{1}{2} \sum_{i:q_1'^i=1,q_2^i=1,q_1^i=1} \pi_i - \frac{1}{2} \sum_{i:q_1'^i=1,q_1^i=0,q_2^i=0} \pi_i \\
=& \Pi_1(\mathbf{q}_1, \neg\mathbf{q}_1) - \frac{1}{2} \sum_{i:q_1^i=1,q_2^i=1,q_1'^i=0} \pi_i - \frac{1}{2} \sum_{i:q_1'^i=1,q_1^i=0,q_2^i=0} \pi_i \leq \Pi_1(\mathbf{q}_1, \neg\mathbf{q}_1).
\end{aligned}$$

Next, let us consider a sequence of iterated best responses,

$$(\mathbf{q}_1^0, \mathbf{q}_2^0) \xrightarrow{1} (\mathbf{q}_1^1, \mathbf{q}_2^0) \xrightarrow{2} (\mathbf{q}_1^1, \mathbf{q}_2^1) \xrightarrow{1} \ldots \xrightarrow{2} (\mathbf{q}_1^N, \mathbf{q}_2^N) \xrightarrow{1} (\mathbf{q}_1^{N+1}, \mathbf{q}_2^N),$$

starting with player 1 and ending at a Nash equilibrium after $2N + 1$ iterations. For ease of exposition we only consider the case that the sequence ends with a movement of player 1. For a fixed $0 < i \leq N$, there are two possibilities:

- if $\mathbf{q}_1^{i+1} = \neg\mathbf{q}_2^i$ then we have reached a Nash equilibrium: since $\mathbf{q}_2^i$ is a best response

to $\mathbf{q}_1^i$, it must be a best response to its inverse $\neg\mathbf{q}_2^i = \mathbf{q}_1^{i+1}$; note that this would imply $i = N$;

- if $\mathbf{q}_1^{i+1} \neq \neg\mathbf{q}_2^i$, then player 1 will never play $\neg\mathbf{q}_2^i$ on subsequent iterations: assume otherwise that $\mathbf{q}_1^{j+1} = \neg\mathbf{q}_2^i$ for some $j > i$. Then $\neg\mathbf{q}_2^i$ would be an (equally) best response to $\mathbf{q}_2^i$ for player 1, so that $(\neg\mathbf{q}_2^i, \mathbf{q}_2^i)$ would be a Nash equilibrium as above. In particular, $\mathbf{q}_2^i$ would be a best response to $\mathbf{q}_1^{j+1} = \neg\mathbf{q}_2^i$ for player 2. Putting all these together, we get a contradiction, as

$$\Phi(\mathbf{q}_1^{i+1}, \mathbf{q}_2^i) = \Phi(\neg\mathbf{q}_2^i, \mathbf{q}_2^i) \geq \Phi(\neg\mathbf{q}_2^i, \mathbf{q}_2^j) = \Phi(\mathbf{q}_1^{j+1}, \mathbf{q}_l^j) > \Phi(\mathbf{q}_1^{i+1}, \mathbf{q}_2^i).$$

By similar principles, we have: if $0 \leq i < j \leq N$, then $\mathbf{q}_2^i \neq \mathbf{q}_2^j$ (i.e. 2 will not repeat strategies); if $0 < i < j \leq N$, then $\mathbf{q}_1^i \neq \mathbf{q}_1^j$ (i.e. 1 will not repeat strategies except possibly for $\mathbf{q}_1^0$); if $0 < i \leq j \leq N$, then $\mathbf{q}_2^j \neq \neg\mathbf{q}_1^i$; and if $0 < i < j \leq N$, then $\mathbf{q}_1^j \neq \neg\mathbf{q}_2^i$.

In other words, we conclude that the strategies $\mathbf{q}_1^1, \neg\mathbf{q}_2^1, \mathbf{q}_1^2, \neg\mathbf{q}_2^2, \ldots, \mathbf{q}_1^N, \neg\mathbf{q}_2^N$ must be all different. As there are only $|S|$ possible strategies, it follows that $2N \leq |S|$; in other words, player 2 can play at most $|S|/2$ different strategies, player 1 can play at most $|S|/2 + 1$ different strategies, and the total number of iterations is at most $|S| + 1$ (for a maximum number of $|S| + 2$ recalculations).

3.A.10 Proof of Theorem 3.4 (Iterated Best Response Algorithm for Approximate Nash Equilibria)

We start by presenting a different conversion from C-PDOP to a congestion game that can handle the description of full strategies. We construct a congestion game having a resource for each location and for each arc in the network. To each arc $e \in A$ we associate a constant negative profit $p_e = -c_e$. Similarly to the proof of Lemma 3.1, for each location $i \in \mathcal{Z}$ we include a resource with profit function

$$p_i(y^i) = \frac{\pi^i}{y^i},$$

where y^i is the number of players having a vehicle available at location i.

The valid strategies for each player correspond to feasible tours that only visit locations/arcs associated with that player, i.e. satisfying constraints (3.1b)-(3.1k) from the PDOP model. In other words, for each valid strategy $s_n = (\mathbf{a}^n, \mathbf{x}^n)$ of player n, we associate a corresponding strategy in the congestion game consisting on those arcs

$e \in A$ and locations $i \in \mathcal{Z}$ for which $a_e^n, x_i^n = 1$. As in the proof of Lemma 3.1, this defines a valid congestion game, albeit in a profit-maximizing instead of cost-minimizing formulation; and where resources may assume positive or negative values. The standard potential argument as in the proof of Corollary 3.1 can then be applied to conclude that any sequence of improving deviations must eventually reach a Nash equilibrium.

Next we consider the ϵ-IBR as described in the statement of the theorem. By definition, a player only deviates if the ϵ-approximately optimal routing found is a strict improvement to that player's profit. Therefore, the ϵ-IBR dynamics still yield a sequence of improving deviations and must terminate after a finite number of iterations. All is left is to prove the quality guarantee of the final strategy, i.e. that $\mathbf{s}$ is an ϵ-Nash equilibrium. Let s'_n be the ϵ-approximate best response to $\mathbf{s}_{-n}$ found by player n. Since player n opts to not deviate from s_n, it follows that

$$(1+\epsilon)\Pi_n(s_n, \mathbf{s}_{-n}) \geq (1+\epsilon)\Pi_n(s'_n, \mathbf{s}_{-n}) \geq \max_{s'_n} \Pi_n(s'_n, \mathbf{s}_{-n});$$

as this relation holds for every player n, the final strategy is an ϵ-Nash equilibrium.

3.A.11 Proof of Lemma 3.4 (Optimality of the Potential Function Optimizer)

Point 1 follows directly from Lemma 3.1, as the maxima of the potential function correspond directly to Nash equilibria of the congestion game. To prove point 2, let $\mathbf{q}$ be a maximum potential Nash equilibrium in which each location is visited by at most one player ($\sum_n q_n^i \leq 1 \; \forall i \in \mathcal{Z}$). Note that this implies, for each location i, that

$$H_{y^i}\pi^i = y^i\pi^i = \sum_n \pi^i(\mathbf{s})q_n^i,$$

and as such $\Phi(\mathbf{q}) = \Pi(\mathbf{q})$. Now, if $\tilde{\mathbf{q}}$ is any other solution, we have

$$\begin{aligned}
\Pi(\tilde{\mathbf{q}}) &= \sum_{i\in\mathcal{Z}}\sum_n \pi^i(\tilde{\mathbf{q}})\tilde{q}_n^i - \sum_n C_n(\tilde{\mathbf{q}}_n) \\
&\leq \sum_{i\in\mathcal{Z}} H_{\tilde{y}^i}\pi^i - \sum_n C_n(\tilde{\mathbf{q}}_n) \\
&= \Phi(\tilde{\mathbf{q}}) \leq \Phi(\mathbf{q}) = \Pi(\mathbf{q}).
\end{aligned}$$

We conclude that **q** is also a welfare-maximizing strategy and thus there cannot be a better Nash equilibrium.

Appendix 3.B Tables

In the following, we give detailed results for Subsections 3.4.2–3.4.7. Table 3.9 shows the percentage increase towards a baseline under different algorithms for either operator for all combinations of parameters.

	Operator 1						Operator 2					
Setting	Baseline	IBR-0	IBR-WP	PFO	QMO	QMP	Baseline	IBR-0	IBR-WP	PFO	QMO	QMP
F_H_H_H	12	160	139	130	-5.46	77.8	12	86.2	125	123	-13.1	66.8
F_H_H_L	12	225	204	206	85.9	180	8	38.5	59.6	67	-61.7	10.1
F_H_L_H	8	79.5	63.6	64.9	-63.2	19.6	12	191	207	215	103	199
F_H_L_L	8	136	110	106	-42.9	30.5	8	59	98.9	97.6	-44.9	32.4
F_L_H_H	6	24.9	21.4	21.4	-12.5	1.9	6	11.1	17.2	17.2	-14.9	1.26
F_L_H_L	6	28.5	28.6	28.7	21.1	8.05	4	3.66	3.93	4.62	-4.4	1.15
F_L_L_H	4	9.59	8.03	7.82	-6.67	1.04	6	21.1	22.9	24.7	13.8	4.25
F_L_L_L	4	5.51	4.88	4.88	1.57	1.07	4	7.04	8.9	8.9	5.31	2.49
P_H_H_H	12	238	234	234	190	223	12	231	240	239	199	228
P_H_H_L	12	243	243	240	210	231	8	149	149	159	117	143
P_H_L_H	8	167	167	164	128	152	12	246	246	251	226	243
P_H_L_L	8	171	166	165	129	152	8	168	178	178	137	164
P_L_H_H	6	28.4	27.4	27.4	18.3	21.2	6	30	32.3	32.3	23.1	23.4
P_L_H_L	6	24.5	24.5	24.2	21.3	17.7	4	8.28	8.28	8.92	4.49	7.25
P_L_L_H	4	6.74	6.74	6.46	3.45	4.89	6	27	27	27.5	24.6	20.9
P_L_L_L	4	10.8	10.6	10.6	6.44	6.66	4	9.29	9.61	9.61	6.04	5.54

Table 3.9: Percentage Profit Increase towards Baseline (Operator 1 in Left Block, Operator 2 in Right Block) under Various Models and Algorithms as well as Different Experimental Settings with Substitution Rates, Margins, and Densities for Either Operator

Analogously, Table 3.10 lists detailed results for the total profit (sum over both operators) towards the baseline, and shows that the profit decreases in all instances when competition is present.

	Full Substitution					Partial Substitution				
Setting	baseline	NE	WP	Coop-PDOP	M-PDOP	baseline	NE	WP	Coop-PDOP	M-PDOP
H_H_H_H	24	132	156	184	184	24	237	240	304	331
H_H_H_L	20	151	160	198	204	20	207	209	276	305
H_H_L_H	20	155	169	204	209	20	216	218	287	316
H_H_L_L	16	104	113	143	143	16	172	175	249	277
L_L_H_H	12	19.3	19.4	42.4	42.5	12	29.9	29.9	91.5	121
L_L_H_L	10	19.1	19.3	38.1	44.6	10	18.1	18.2	57.8	88.7
L_L_L_H	10	17.9	17.9	36.7	41.8	10	19.1	19.1	65.9	93.8
L_L_L_L	8	6.89	6.89	11.2	11.2	8	10.1	10.1	37.7	59.6

Table 3.10: Percentage Profit Increase towards Baseline (Full Substitution in Left Block, Partial Substitution in Right Block) for Different Experimental Settings with Substitution Rates, Margins, and Densities for either Player

Table 3.11 presents detailed information about attainable profit increases if the assumption of homogeneous payoffs does not hold.

Table 3.12 outlines profit increases if more than one vehicle is in demand in a district.

Table 3.13 outlines profit increases if customers do not choose vehicles completely at random, if both operators have a vehicle available at this location.

	Full Substitution								Partial Substitution							
	Operator 1				Operator 2				Operator 1				Operator 2			
Setting	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP
H_H_H_H	12	142	-3	100	12	118	-1	98	12	434	434	434	12	430	430	430
H_H_H_L	12	279	143	251	8	82	-62	46	12	424	424	424	8	355	355	355
H_H_L_H	8	98	-85	35	12	260	141	258	8	358	358	358	12	433	433	433
H_H_L_L	8	152	-59	84	8	71	-72	52	8	354	354	354	8	349	349	349
H_L_H_H	12	321	55	274	6	1	-163	-37	12	423	423	423	6	257	257	257
H_L_H_L	12	404	236	358	4	-46	-215	-61	12	425	425	425	4	147	147	147
H_L_L_H	8	212	-21	108	6	113	-23	96	8	362	362	362	6	244	244	244
H_L_L_L	8	316	39	194	4	-47	-239	-66	8	358	358	358	4	122	122	122
L_H_H_H	6	20	-170	-38	12	273	70	279	6	295	295	295	12	416	416	416
L_H_H_L	6	145	-18	92	8	140	-15	96	6	292	292	292	8	336	336	336
L_H_L_H	4	-26	-237	-77	12	371	226	365	4	169	169	169	12	427	427	427
L_H_L_L	4	0	-222	-67	8	211	39	154	4	170	170	170	8	336	336	336
L_L_H_H	6	162	-110	35	6	57	-117	8	6	294	294	294	6	254	254	254
L_L_H_L	6	257	63	169	4	-17	-187	-65	6	277	277	277	4	133	133	133
L_L_L_H	4	62	-166	-54	6	174	48	155	4	182	182	182	6	245	245	245
L_L_L_L	4	130	-105	-49	4	8	-135	-56	4	157	157	157	4	141	141	141

Table 3.11: Percentage Profit Increase towards Baseline with Inhomogeneous Payoffs (Full Substitution in Left Block, Partial Substitution in Right Block) for Different Experimental Settings with Substitution Rates, Margins, and Densities for either Player

		1 location per district						2 location per district						3 location per district					
		Op. 1			Op. 2			Op. 1			Op. 2			Op. 1			Op. 2		
Setting	Baseline	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP	IBR-0	QMO	QMP
F_H_0.9	12	161	2	74	102	2	88	134	14	84	115	22	92	113	14	63	103	24	85
F_H_0.7	12	161	5	81	101	3	87	135	31	85	113	31	94	128	48	96	118	56	82
F_H_0.5	12	158	8	86	102	5	78	147	56	110	99	40	85	143	78	105	124	72	98
F_L_0.9	6	36	-11	9	17	-14	9	31	10	10	22	8	12	23	11	10	28	19	13
F_L_0.7	6	39	-9	9	16	-16	9	25	12	10	23	14	10	37	29	12	31	24	12
F_L_0.5	6	33	-19	8	22	-15	11	33	3	11	30	8	13	59	27	18	40	14	15
P_H_0.9	12	249	202	232	237	203	232	192	172	185	195	179	191	141	117	132	168	148	160
P_H_0.7	12	244	201	227	237	204	234	208	187	202	200	184	196	181	157	174	174	155	168
P_H_0.5	12	243	197	227	234	202	230	223	206	215	214	201	214	220	203	215	215	200	209
P_L_0.9	6	36	27	27	29	25	24	33	29	29	29	27	27	25	24	22	33	32	28
P_L_0.7	6	35	22	25	36	28	30	33	30	31	32	30	28	38	35	33	38	36	33
P_L_0.5	6	35	26	26	31	24	24	42	39	37	39	37	36	54	44	46	61	52	55

Table 3.12: Average Profit Increase towards Baseline (No Rebalancing) With Diminishing Returns for an Increasing Number of Locations per District

	Full Substitution								Partial Substitution							
	Operator 1				Operator 2				Operator 1				Operator 2			
Setting	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP	Baseline	IBR-0	QMO	QMP
H_H_H_0.5	12	145	-2	65	12	106	3	79	12	254	254	121	12	251	251	125
H_H_H_0.75	12	248	125	242	12	15	-139	0	12	254	254	242	12	253	253	2
H_H_H_1	12	254	254	254	12	14	-266	0	12	258	258	258	12	257	257	0
H_H_L_0.5	12	229	96	113	8	48	-53	19	12	250	250	116	8	179	179	42
H_H_L_0.75	12	251	176	242	8	13	-174	0	12	262	262	252	8	185	185	3
H_H_L_1	12	251	251	251	8	7	-302	0	12	253	253	253	8	194	194	0
H_L_H_0.5	8	93	-52	18	12	173	88	113	8	182	182	33	12	260	260	136
H_L_H_0.75	8	165	67	157	12	111	18	1	8	191	191	171	12	255	255	3
H_L_H_1	8	188	188	188	12	103	-63	0	8	189	189	189	12	249	249	0
H_L_L_0.5	8	149	-43	28	8	48	-49	26	8	186	186	21	8	182	182	33
H_L_L_0.75	8	187	76	169	8	14	-159	0	8	178	178	158	8	182	182	2
H_L_L_1	8	186	186	186	8	7	-271	0	8	182	182	182	8	180	180	0
L_H_H_0.5	6	27	-25	3	6	10	-28	0	6	24	24	1	6	26	26	0
L_H_H_0.75	6	26	2	7	6	6	-51	0	6	24	24	9	6	28	28	0
L_H_H_1	6	24	24	24	6	10	-68	0	6	29	29	29	6	23	23	0
L_H_L_0.5	6	26	10	1	4	5	-15	1	6	33	33	3	4	9	9	1
L_H_L_0.75	6	25	21	10	4	3	-13	0	6	27	27	8	4	4	4	0
L_H_L_1	6	33	33	33	4	1	-28	0	6	34	34	34	4	7	7	0
L_L_H_0.5	4	6	-7	1	6	19	13	2	4	10	10	1	6	30	30	2
L_L_H_0.75	4	10	3	5	6	21	13	0	4	7	7	3	6	26	26	1
L_L_H_1	4	6	6	6	6	19	3	0	4	8	8	8	6	25	25	0
L_L_L_0.5	4	6	0	1	4	6	1	1	4	6	6	0	4	9	9	1
L_L_L_0.75	4	5	2	0	4	5	-1	0	4	10	10	6	4	11	11	0
L_L_L_1	4	6	6	6	4	7	-4	0	4	8	8	8	4	10	10	0

Table 3.13: Percentage Profit Increase towards Baseline with Other Customer Choice Behaviors (Full Substitution in Left Block, Partial Substitution in Right Block) for Different Experimental Settings with Substitution Rates, Margins (Equal), and Densities for either Player, and Variable Customer Preferences

Chapter 4

Feature-Based Selection of Carsharing Relocation Modes

One-way and free-floating carsharing systems must be rebalanced to achieve a high service level, and thus generate benefits for users and society. In practice, vehicles can be relocated with multiple modes (e.g., by truck or by driving them), but a single mode is sufficient in many instances. Obviously, a single mode is preferred from a computational standpoint: The routing problems are less complex since less synchronization is necessary, and thus solve much faster. It remains an open question which features drive the decision on the best mode, and if operators can decide a-priori whether hybridization of several modes is beneficial/necessary, and which modes one should hybridize among. We build a classifier based on linear regression which predicts the costs for all individual modes. The advantage of this approach is that cost estimates (i) can be used as a feature in other approaches, and (ii) allow operators to estimate the necessary budget upfront. However, cost estimates cannot be used directly to determine key drivers for modal choice. We, thus, use logistic regression and decision trees for determining the best mode. These approaches are better at determining relevant features that explain which mode is preferred in an instance. We find that the most important features to decide between modes are vehicle and truck costs per kilometer as well as their velocities, and the average number of vehicles that shall be relocated per day (that is, the imbalance of the system). In most instances, the decision is between driving vehicles to rebalance them and rebalance staff by biking, or loading vehicles onto a truck. Hybridization proves useful in ≈20% of all instances, and a simple rule-based classifier is able to predict correctly that hybridization is necessary in most instances.

4.1 Introduction

Free-floating carsharing can provide benefits to users and society. One carsharing vehicle can replace multiple private vehicles (Shaheen and Cohen, 2013). Since vehicles in carsharing systems are often more efficient than privately owned vehicles, emissions can be reduced (Bellos et al., 2017; Firnkorn and Müller, 2011). Baptista et al. (2014) report that carsharing can even provide mobility to communities with previously little access to mobility. However, operating such a system profitably is challenging: If users frequently take one-directional trips, e.g. to cover the last mile (Liang et al., 2016), fleets become unbalanced, and vehicles have to be relocated. Only by relocating vehicles, operators can achieve sufficiently high service levels. Rebalancing also circumvents accumulating vehicles in some regions, increasing the availability of parking spaces, and may allow for smaller fleet sizes. Obviously, these rebalancing movements pose a substantial cost factor ranging in the same order of magnitude as vehicle procurement costs (Vasconcelos et al., 2017). Rebalancing adds mileage and increases congestion, counteracting environmental and societal benefits. Thus, rebalancing must be optimized to reduce impacts. To that end, rebalancing operations are typically pooled during the night (Almeida Correia and Antunes, 2012), and operators employ different "relocation modes", including bike, public transportation, car, and truck (Weikl and Bogenberger, 2015). The real-world case studies for European and North American free-floating carsharing services by Bruglieri et al. (2014b) and Nourinejad et al. (2015) indicate that operators frequently use (foldable or shared) bike as a relocation mode. In more theoretical contributions, truck, car and public transit are suggested as rebalancing modes (e.g., Gambella et al. (2018) and Herbawi et al. (2016) for rebalancing modes that are similar to "car", Dror et al. (1998) and Krumke et al. (2013) for rebalancing by truck, and Huang et al. (2020) and Santos and Almeida Correia (2019) for public transit). Bikes are beneficial from an environmental standpoint, since not all segments of the rebalancing route result in emissions. Further, bikes are flexible and can often use a more direct (and thus faster) route between any two locations. Public transportation is similarily beneficial from an environmental standpoint, but highly dependent on a dense network with frequently running trains. Since carsharing providers have large fleets of vehicles that they can also use for staff rebalancing, implementing the rebalancing mode "car" is simpler than the other modes. On the other hand, the necessary number of workers increases significantly, making this mode usually more expensive. Rebalancing vehicles using a truck is beneficial if demand is clustered, e.g. if multiple vehicles must be transported from the airport to the trade

fair. However, trucks are highly inflexible and result in low velocities in narrow city centers if they are even permitted to enter. Truck drivers might have to park the truck at some location and walk to retrieve a car. Weikl and Bogenberger (2015) report that modes are mixed in their case study.

The operator's routing problem can be formulated as a Vehicle Routing Problem (VRP)/Pickup and Delivery Problem (PDP) with multiple modes and the consequential synchronization constraints. The synchronization constraints obviously result in long running times, or suboptimal solutions within a reasonable runtime for an operational planning problem. The mode selection problem has tactical and strategic implications: for example, trucks must be bought or rented, employees hired. Thus, operators have an incentive to use the same (subset of) modes for an extended period of time, even if on some days due to the random realization of the demand imbalance, using a different or additional mode would reduce costs slightly.

We focus on the problem of selecting the cost-minimizing mode(s) for rebalancing a carsharing system. As stated before, there are features that hinder or support the adoption of modes. While it is often logical how a feature impacts the total rebalancing costs (e.g., if additional vehicles are to be rebalanced, costs increase), the magnitude of their influence on the costs and their influence on the modal choice, remain unknown. Using different classifiers, we establish the relevance of said features for the modal choice. We predict the costs of each mode individually using linear regression, the probability of each mode using logistic regression, and the best mode using a decision tree (the latter may use the cost estimates as additional features). Using linear regression as an intermediate step is advantageous for two main reasons: (i) cost estimates (and differences between them) are a valuable feature for deciding whether or not to include a mode; and (ii) it allows operators to allocate sufficient budgets a-priori. We establish the decrease of excess cost and increase of accuracy due to cost estimates as features. Linear regression and decision trees have in common that they are easy to understand for decision makers: The cost contribution of each feature is directly visible (in linear regression), and cutoff points between different modes are obvious (decision trees). Logistic regression is particularily well-suited to determine the relevance of the features, since the odds ratio allows one to directly read off the relevance of a feature.

The contribution of this Chapter is two-fold: (i) We investigate which features drive the modal choice, and how strongly they impact the choice. As such, we support carsharing operators in their decision on which modes to use to rebalance their fleet. An operator has to set the feature values for her operating area, and can obtain the optimal

mode or modal combination. Since the decision tool provides intermediate information (cost estimates for the individual modes), she can also easily adapt the decision rule to include more modes, or be more strict. (ii) We contribute to the literature on integrating methods from operational research and analytics. Our approach is applicable in many different domains in which insights of a operational decision impact strategic decisions, and results should carry over between different environments. The feature-based selection algorithms only require information that can be known to carsharing providers prior to starting their service. On a side note, we show that in most of the instances, the optimal rebalancing mode is either bike or truck, whereas other modes hardly ever become relevant. Further, the intermediate results can also be used to steer tactical decisions such as setting the shift length.

Section 4.2 introduces the available modes for rebalancing carsharing systems, and shows how they can be integrated. In Section 4.3, we present feature-based selection algorithms and apply them on synthetic data. Section 4.4 concludes this Chapter.

4.2 Rebalancing Modes and Problem Formulations

We first introduce the individual modes, and subsequently introduce the model formulation for the Multi-Mode Carsharing Relocation Problem (M-CRP). The single-mode Carsharing Relocation Problems (CRPs) can be instantiated from the M-CRP by restricting usage of other modes.

Locations with an excess in vehicles are called "pickup locations", and locations with a demand for additional vehicles are called "delivery locations". Pickup and delivery locations have unit-demand which does not restrict the problem, since locations can be split as explained by Dror et al. (1998). All pickup and delivery demand must be satisfied.

4.2.1 Relocation Problem with Car

A group of workers relocate vehicles by driving them. To get to the next vehicle, workers give each other lifts. We consider two versions of the Carsharing Relocation Problem with Car (CRP-C): In the Carsharing Relocation Problem with Car and Dedicated Helpers (CRP-CD), a dedicated second worker ("helper") follows the first worker ("relocator"), and gives her lifts from each delivery to the next pickup location. In the Carsharing Relocation Problem with Car and Shared Helpers (CRP-CS), all workers can be "helpers"

and "relocators", and roles may switch during the shift. Figure 4.1 depicts an example

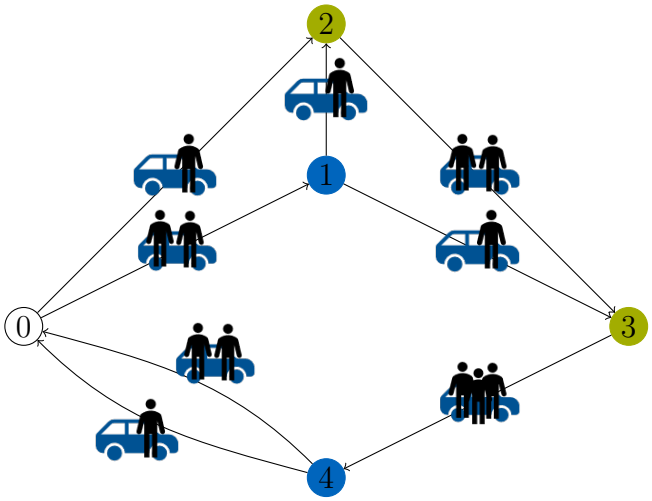

Figure 4.1: Example Routing CRP-C: Three workers leave the depot in two vehicles, jointly relocate the vehicles (from 1 to 2 and from 3 to 4), and return to the depot in two vehicles. The grouping of workers into relocator-helper pairs changes during the route.

route for the CRP-CS.

Using cars to rebalance a carsharing fleet poses a very natural choice, since a large number of vehicles is available to operators. As such, rebalancing by car can always work as a default option if none of the other modes can be used. Also, using cars for rebalancing can easily be integrated with rebalancing by public transportation. However, the number of employees needed for this mode is often higher than with the other modes, since arcs are traversed by two workers. If helpers are "shared", synchronization becomes necessary which impacts the computation time for an optimal tour as well as the tour itself. Sharing helpers is only preferential if the number of vehicles that shall be rebalanced is sufficiently high (the number of vehicles indirectly drives the number of workers, and multiple workers are necessary to make sharing of helpers more profitable than dedicated helpers).

The CRP-CD is modeled as a VRP with maximum tour duration on a bipartite graph (with symmetric travel times but asymmetric routing costs). The CRP-CS is modeled as a VRP with synchronization constraints and a maximum tour duration. Obviously, any solution to the CRP-CD is also feasible for the CRP-CS. Another special case of the CRP-CS is relocation by van. There, only some vehicles can be used for staff rebalancing, and these vehicles may have a larger capacity. We omit this mode, since the number of workers in most instances does not substantially exceed the vehicle capacity, and the

vehicle capacity is almost never utilized.

4.2.2 Relocation Problem with Bike

As an alternative to hitching rides with other relocators for trips from delivery to pickup location, operators can use a (foldable) bike that fits into the trunk of the carsharing car (or use bikesharing if it is reliably available). In urban areas, the advantage of the mode bike is clear: Bike lanes are less congested and often more direct than roads for motorized travel, resulting in similar travel times at lower costs (vehicle costs and salaries). However, bike usage may be restricted by weather (e.g., in the Scandinavian winters) or road conditions (e.g., in North American cities without proper bike lanes). This either results in lower velocities or higher costs to include the possibility that the mode must be switched. Obviously, the worker and her bike must travel the same route, reducing the possibility for integrating relocating by bike with other modes.

The Carsharing Relocation Problem with Bike (CRP-B) is modeled as a VRP with maximum tour duration on a bipartite graph, and both travel times and routing costs are asymmetric. Figure 4.2 depicts an example route for the CRP-B.

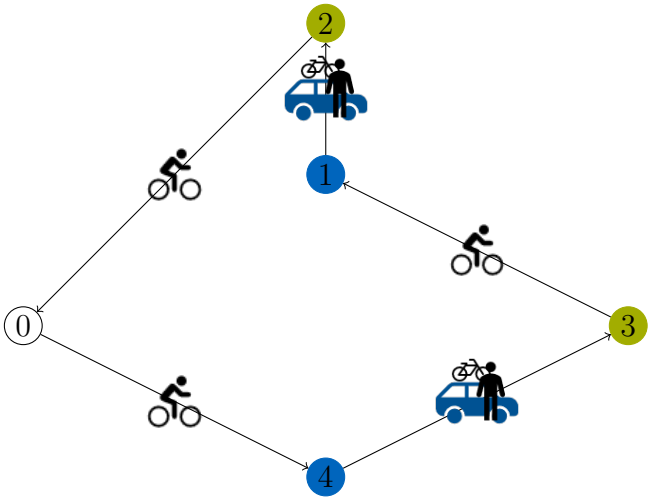

Figure 4.2: Example Routing CRP-B: One worker leaves the depot riding a bike. When relocating vehicles from 4 to 3, and from 1 to 2, she stores the bike in the trunk, to be able to use the bike for riding from 3 to 1 and for returning to the depot.

4.2.3 Relocation Problem with Public Transit

Workers can also use public transit to reach the next pickup location. This implies walking to the next location, using one or more lines (incurring waiting time at every

changeover), and walking to the destination. This mode is cheap (costs for public transportation day passes are negligible compared to vehicle costs or wages), and can be very fast if carsharing vehicles are parked in close proximity to public transportation stations, but can be much slower than other modes if walking distances and waiting times increase. If public transportation is unavailable in parts of the operating area, this mode can become infeasible due to the low walking speed (violating shift length restrictions).

The Carsharing Relocation Problem with Public Transit (CRP-P) is also modeled as a VRP with maximum tour duration on a bipartite graph. Figure 4.3 depicts an example route for the CRP-P. The travel times and routing costs for the route segment from

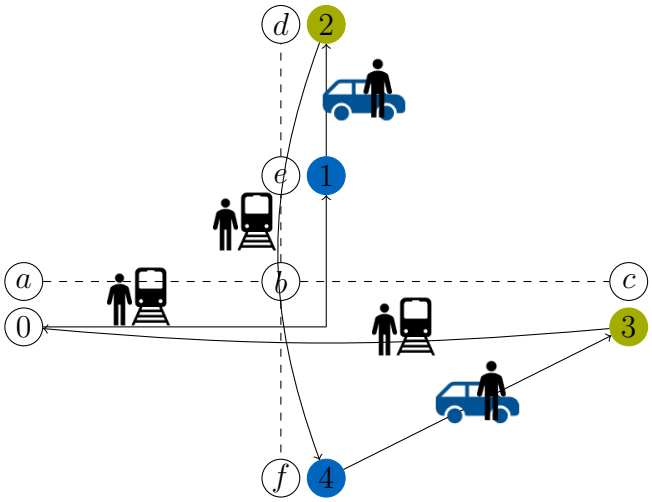

Figure 4.3: Example Routing CRP-P: The routing decision is restricted by the availability of public transportation. One worker leaves the depot using public transport (walking to station a on the horizontal line, transferring at b to the vertical line to get to e). She also uses public transport from 2 to 4 (using the vertical line), and to return to the depot (using the horizontal line).

delivery to pickup location are derived from a shortest-path problem on an auxiliary graph containing expected travel and waiting times. Thus, we must assume that the service frequency is constant during the rebalancing period. While this restricts the applicability of the rebalancing problem in the late evening (when public transit phases down), it is realistic during the night after the public transit system operates on its night schedule.

4.2.4 Relocation Problem with Truck

In particular for longer route segments, it can be beneficial to pool multiple relocation operations by loading vehicles onto a truck rather than driving each vehicle individually.

Inaccessibility poses a challenge for truck routing in city centers with narrow roads: It can be necessary to park the truck at a nearby location (or one of multiple access points), and walk to the pickup location to retrieve the vehicle. Rebalancing by truck is beneficial if the increased capacity can be utilized, i.e., if pickup and delivery demands are clustered. This mode is clearly dominated if walking distances are large, or if trucks face severe issues maneuvering in the city, resulting in low average velocities.

The Carsharing Relocation Problem with Truck (CRP-T) is modeled as VRP with simultaneous pickup and delivery, multiple pickup options, capacity limit and maximum tour duration. Figure 4.4 depicts an example route for the CRP-T.

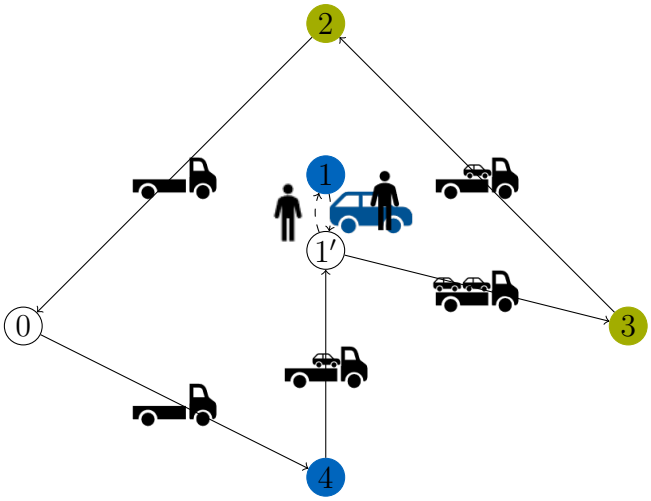

Figure 4.4: Example Routing CRP-T: The truck (driven by one worker) leaves the depot empty. Two vehicles are loaded onto the truck consecutively. The vehicle at 1 is not directly accessible, forcing the truck driver to retrieve the vehicle by foot. Both vehicles are delivered consecutively, before the truck returns to the depot.

4.2.5 Multi-Modal Carsharing Relocation Problem

Obviously, operators can benefit from integrating the above modes (all or a subset of them). Then, a pickup or delivery demand is satisfied if a worker using any of the permitted modes visits this demand node. Workers are no longer restricted to one mode, but might hitch a ride for parts of their tour and use public transit for another route segment.

The M-CRP is modeled as a VRP with multiple synchronization constraints on load, service and timing, as well as constraints particular to the modes.

Notation

The set of all locations $\mathcal{N}$ is comprised of pickup locations ($\mathcal{N}^+$), delivery locations ($\mathcal{N}^-$) with unit capacity, entry points for pickup locations $E_i, i \in \mathcal{N}^+$ that are inaccessible by truck, and the depot ($\mathcal{D}$). If excess supply or demand at some geographical location is higher than 1, this location is split into multiple co-located nodes (similar to Dror et al. (1998)). If a pickup node i is directly accessible, we create a copy i' that represents the access point to i ($i' \in E_i$). There is exactly one vehicle available for relocation at every pickup location. If multiple vehicles are present at the same geographic location, multiple nodes are created (similar, if multiple vehicles are required at the same location). The depot is 0 and $n{+}1$ where 0 marks the start of all tours and $n{+}1$ the end ($\mathcal{D} = \{0, n{+}1\}$).

$\mathcal{W}$ workers and $\mathcal{X}$ trucks leave the depot 0 and return to the duplicate depot $n+1$. Workers can hitch rides with colleagues, use a bike or public transport. The maximum number of workers riding bikes is $\mathcal{Y} \leq \mathcal{W}$. The number of vehicles leaving the depot is limited to $\mathcal{V}$. The capacity of a carsharing vehicle is set to θ workers, and the capacity of a truck is denoted by κ cars.

As previously stated, every node is visited by at most two vehicles. This implies that each arc (i, j) between two locations $i \in \mathcal{N}$ and $j \in \mathcal{N}$ is used by at most two vehicles. We use the variables v_{ij}^a, w_{ij}^a, x_{ij}, y_{ij}^a, z_{ij} to denote the routing of carsharing vehicles (v_{ij}^a), workers in those vehicles (w_{ij}^a), trucks (x_{ij}), bike workers (y_{ij}^a), and public transport workers (z_{ij}), respectively. Here, i refers to the start of an arc, j to its end, and $a \in \mathcal{A} = \{1, 2\}$ to one of the two parallel arcs. v_{ij}^a, x_{ij} and x_{ij} are binary variables as every arc can be used by at most one vehicle, truck, or public transportation worker, whilst w_{ij}^a and y_{ij}^a are Integer, referring to the number of workers on an arc. As variable y_{ij}^a does not distinguish between bike workers riding their bike and bike workers traveling by car, we introduce an auxiliary variable $y_{ij}^{a\prime}$ which tracks bike workers traveling by car (thus, $y_{ij}^a - y_{ij}^{a\prime}$ refers to bike workers traveling by bike). u_{ij} tracks if a truck collects a vehicle at j through entry point i. s_i is the timing variable and tracks when a node $i \in \mathcal{N}$ is visited. q_i represents the load on the truck when the truck leaves node i.

The costs are comprised of routing costs for each type of transport: c_{ij}^v for vehicles, c_{ij}^w for workers in those vehicles, c_{ij}^x for trucks, c_{ij}^y for workers riding bikes, and c_{ij}^z for workers using public transport. We denote the travel time as $t_{ij}^v, t_{ij}^w, t_{ij}^x, t_{ij}^y, t_{ij}^z$, respectively. The cost for loading a vehicle onto a truck and the associated time are denoted as t_{ik}^u, c_{ik}^u where k is the pickup location and i is the entry point.

Model Formulation

We minimize the total routing cost for the entire fleet:

$$\min \sum_{a\in\mathcal{A}} \sum_{i,j\in\mathcal{N}} \left(c^v_{ij} v^a_{ij} + c^w_{ij}\left(w^a_{ij} + y^{a\prime}_{ij}\right) + c^y_{ij}\left(y^a_{ij} - y^{a\prime}_{ij}\right)\right) + \sum_{i,j\in\mathcal{N}} \left(c^x_{ij} x_{ij} + c^z_{ij} z_{ij}\right) + \sum_{i,j\in\mathcal{N}} c^u_{ij} u_{ij}. \tag{4.1}$$

Routing costs consist of vehicle routing costs $\mathbf{c}^v$, worker routing costs for workers in vehicles $\mathbf{c}^w$ (both "bike" workers $\mathbf{y}'$ and regular workers $\mathbf{w}$), routing costs for workers on bikes $\mathbf{c}^y$ $(\mathbf{y} - \mathbf{y}')$ and workers using public transport $\mathbf{c}^z$, as well as truck routing costs $\mathbf{c}^x$, and costs for accessing locations which cannot directly be reached with a truck $\mathbf{c}^u$.

We first consider the routing of the truck:

$$\sum_{j\in\mathcal{N}} x_{ij} = \sum_{j\in\mathcal{N}} x_{ji} \qquad \forall i \in \mathcal{N} \setminus \mathcal{D} \tag{4.2}$$

$$\sum_{j\in\mathcal{N}} x_{ij} = 0 \qquad \forall i \in \mathcal{N}^+ \tag{4.3}$$

$$u_{ik} \leq \sum_{j\in\mathcal{N}} x_{ij} \qquad \forall k \in \mathcal{N}^+, i \in E_k \tag{4.4}$$

$$q_i \geq q_j + \sum_{k\in\mathcal{N}} u_{ik} - (1 - x_{ji})\,\kappa \qquad \forall i \in \bigcup_{l\in\mathcal{N}^+} E_l, j \in \mathcal{N} \tag{4.5}$$

$$q_i \leq q_j + \sum_{k\in\mathcal{N}} u_{ik} + (1 - x_{ji})\,\kappa \qquad \forall i \in \bigcup_{l\in\mathcal{N}^+} E_l, j \in \mathcal{N} \tag{4.6}$$

$$q_i \geq q_j - 1 - (1 - x_{ji})\,\kappa \qquad \forall i \in \mathcal{N}^-, j \in \mathcal{N} \tag{4.7}$$

$$q_i \leq q_j - 1 + (1 - x_{ji})\,\kappa \qquad \forall i \in \mathcal{N}^-, j \in \mathcal{N} \tag{4.8}$$

$$q_i \leq \sum_{j\in\mathcal{N}} x_{ij}\kappa \qquad \forall i \in \mathcal{N} \tag{4.9}$$

$$\sum_{j\in\mathcal{N}} x_{0j} \leq \mathcal{X} \tag{4.10}$$

$$\sum_{j\in\mathcal{N}} x_{j,n+1} \leq \mathcal{X} \tag{4.11}$$

A truck must leave all locations it enters (flow conservation, constraints (4.2)), but cannot directly visit pickup locations (constraints (4.3)). Instead, there exists at least one entry point ($i \in E_k \subset \mathcal{N}$) from where this point k can be visited (constraints (4.4)). Subtour elimination is necessary, but will follow directly from timing which

we introduce later. We ensure that the capacity of the truck (parameter κ) is never exceeded by guaranteeing that the load on the truck when leaving some node i does not exceed κ (constraints (4.5)-(4.9)). When leaving and entering the depot, the load of the truck is the same intrinsically, and we do not restrict this any further at this stage (constraints (4.18) will ensure that sufficient vehicles are available at the depot). The number of trucks leaving and entering the depot is limited to $\mathcal{X}$, and if some trucks are not required, they will not leave the depot (constraints (4.10)-(4.11)).

Next, we consider the routing of vehicles driven by workers:

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} v_{ji}^a - \sum_{j\in E_i} u_{ji} = \sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} v_{ij}^a - 1 \qquad \forall i \in \mathcal{N}^+ \qquad (4.12)$$

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} v_{ji}^a - 1 = \sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} v_{ij}^a - \sum_{j\in\mathcal{N}} x_{ij} \qquad \forall i \in \mathcal{N}^- \qquad (4.13)$$

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} v_{ij}^a \leq 2 \qquad \forall i \in \mathcal{N}^- \qquad (4.14)$$

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} v_{ji}^a \leq 1 \qquad \forall i \in \mathcal{N}^- \qquad (4.15)$$

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} v_{ij}^a \leq 1 \qquad \forall i \in \mathcal{N}^+ \qquad (4.16)$$

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} v_{ji}^a \leq 2 \qquad \forall i \in \mathcal{N}^+ \qquad (4.17)$$

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} v_{0j}^a + q_0 \leq \mathcal{V} \qquad (4.18)$$

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} v_{j,n+1}^a + q_{n+1} \leq \mathcal{V} \qquad (4.19)$$

All vehicles which are not loaded onto a truck must be driven from their pickup locations to a delivery location. If a vehicle is driven from $i \in \mathcal{N}^+$ to $j \in \mathcal{N}^-$ (and thus, not carried on a truck: $\sum_{k\in E_i} u_{ki} = 0$ and $\sum_{k\in\mathcal{N}} x_{kj} = 0$), one more vehicle leaves i than enters i (constraints (4.12)), whilst one more vehicle enters j than leaves j (constraints (4.13)). The number of vehicles at any node is restricted by constraints (4.14)-(4.17). Implicitly, there are two vehicles leaving i and one vehicle entering i, if i is serviced by car, and one vehicle leaving i and no vehicle entering i, if i is serviced by bike or public transportation (no vehicle enters or leaves i, if i is serviced by truck). The reverse holds for delivery locations j. Similar to trucks, subtour elimination will follow from the timing constraints introduced at a later stage. The number of vehicles leaving (and entering) the depot (being driven or loaded onto a truck) is limited to $\mathcal{V}$ (constraints

(4.18)-(4.19)).

Worker routing must reflect that the movement of workers is contingent to vehicles, bikes, or public transport:

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} w_{ij}^a + \sum_{j\in\mathcal{N}} z_{ij} = \sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} w_{ji}^a + \sum_{j\in\mathcal{N}} z_{ji} \qquad \forall i \in \mathcal{N}\setminus\mathcal{D} \tag{4.20}$$

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} y_{ij}^a = \sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} y_{ji}^a \qquad \forall i \in \mathcal{N}\setminus\mathcal{D} \tag{4.21}$$

$$v_{ij}^a \le w_{ij}^a + y_{ij}^{a\prime} \qquad \forall i,j \in \mathcal{N}, a \in \mathcal{A} \tag{4.22}$$

$$w_{ij}^a + y_{ij}^{a\prime} \le \theta \cdot v_{ij}^a \qquad \forall i,j \in \mathcal{N}, a \in \mathcal{A} \tag{4.23}$$

$$y_{ij}^{a\prime} \le y_{ij}^a \qquad \forall i,j \in \mathcal{N}, a \in \mathcal{A} \tag{4.24}$$

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} w_{0j}^a + \sum_{j\in\mathcal{N}} z_{0j} + \sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} y_{0j}^a \le \mathcal{W} \tag{4.25}$$

$$\sum_{a\in\mathcal{A}}\sum_{j\in\mathcal{N}} y_{0j}^a \le \mathcal{Y} \tag{4.26}$$

Workers driving vehicles (or hitching a ride with another vehicle) and workers commuting by public transport can switch roles (constraints (4.20)), whilst workers using bikes are fixed in their role, as they have to bring along their bike. Thus, we model flow conservation for bike workers (constraints (4.21)) and other workers (constraints (4.20)) separately. Subtours will be eliminated via timing constraints. Vehicles cannot operate unless a worker travels along the same arc (constraints (4.22)), and no worker can travel by car, unless a vehicle travels there as well (constraints (4.23)). The latter set of constraints also ensures that the capacity θ for workers per vehicle is not exceeded. Constraints (4.24) are variable assignment constraints that link workers with bikes (regardless of their current mode of transport) and bike workers driving cars. At most $\mathcal{W}$ workers can leave the depot (constraint (4.25)), and at most $\mathcal{Y}$ of these can use bikes (constraint (4.26)).

Worker and vehicle movements are synchronized using timing variables:

$$s_i + v_{ij}^a \left(T_{\max} + t_{ij}^v\right) \le s_j + T_{\max} \qquad \forall i,j \in \mathcal{N}, a \in \mathcal{A} \tag{4.27}$$

$$s_i + z_{ij} \left(T_{\max} + t_{ij}^z\right) \le s_j + T_{\max} \qquad \forall i,j \in \mathcal{N} \tag{4.28}$$

$$s_i + \left(y_{ij}^a - y_{ij}^{a\prime}\right) \left(T_{\max} + t_{ij}^y\right) \le s_j + T_{\max} \qquad \forall i,j \in \mathcal{N}, a \in \mathcal{A} \tag{4.29}$$

$$s_i + x_{ij} \left(T_{\max} + t_{ij}^x\right) + \sum_{g\in\mathcal{N}} t_{ig}^u u_{ig} \le s_j + T_{\max} \qquad \forall i,j \in \mathcal{N} \tag{4.30}$$

For each node, we define a point in time s_i. This is possible, as every node is visited exactly once. All s_i must be lower than or equal to the time limit/maximum shift length $T_{\max}$. Obviously, cars and workers in cars travel at the same velocity, resulting in equal travel times t^v_{ij}. A vehicle which leaves i and travels directly to j, thus, imposes the timing constraint that s_j must be at least t^v_{ij} more than s_i (constraints (4.27)). Similarly, we impose timing constraints for public transportation (constraints (4.28)). If a bike worker is traveling by car, timing is already taken care of by constraints (4.27). If she travels by bike ($y^a_{ij} - y^{a'}_{ij} = 1$), constraints (4.29) ensure feasible routes with respect to timing. For the timing of the truck, we further consider the time required to walk to a pickup location and load the vehicle onto the truck (constraints (4.30)).

The domains are as follows:

$$u^a_{ji} \in \{0,1\} \qquad \forall i \in \mathcal{N}, j \in E_i \tag{4.31}$$
$$v^a_{ij} \in \{0,1\} \qquad \forall i,j \in \mathcal{N}, a \in \mathcal{A} \tag{4.32}$$
$$w^a_{ij} \in \mathbb{N}_0 \qquad \forall i,j \in \mathcal{N}, a \in \mathcal{A} \tag{4.33}$$
$$x_{ij} \in \{0,1\} \qquad \forall i,j \in \mathcal{N} \tag{4.34}$$
$$y^a_{ij}, y^{a'}_{ij} \in \mathbb{N}_0 \qquad \forall i,j \in \mathcal{N}, a \in \mathcal{A} \tag{4.35}$$
$$z_{ij} \in \{0,1\} \qquad \forall i,j \in \mathcal{N} \tag{4.36}$$
$$q_i \geq 0 \qquad \forall i \in \mathcal{N} \tag{4.37}$$
$$s_i \geq 0 \qquad \forall i \in \mathcal{N}. \tag{4.38}$$

4.3 Feature-Based Selection and Numerical Study

To investigate which features drive the modal choice, we employ three different classification algorithms, namely multiple linear regression, multinomial logistic regression, and decision trees. We first introduce the numerical design and the classification algorithms. The classifiers are then used to derive features driving the modal choice. Subsequently, we show how the algorithms can be used to determine the benefit of hybridization, show that misclassified instances are incorrectly classified due to random variances in the tours, and apply the model to several example cities.

4.3.1 Data Generation and Features

There are various factors which can influence the rebalancing cost. These are related to the city/environment (e.g., average velocities for the individual modes and quality of public transportation), as well as the operator and her fleet (e.g. number of vehicles that shall be rebalanced and shift length). A list of features $f \in \mathcal{F}$ can be found in

Feature		High Value	Low Value
No. Vehicles	$\lvert\mathcal{N}^+\rvert$	20	10
Shift Length [min]	$T^{\max}$	480	240
Velocity Car [km/h]	v^v	40	20
Costs Car [€/km]	c^v	0.3	0.1
Capacity Car	θ	4	2
Velocity Bike [km/h]	v^b	20	10
Velocity Walking [km/h]	v^w	6	3
Velocity Publ. [km/h]	v^p	50	25
Distance between lines [km]	d^p	3	1.5
Time Changeover [h]	t^c	0.2	0.1
Costs Employee [€/h]	c^e	20	10
Velocity Truck [km/h]	v^t	30	15
Costs Truck [€/km]	c^t	1.875	0.625
Capacity Truck	κ	4	2
No. Access Pts.	$\lvert A_i\rvert$	4	2
Max. Distance Access Pt. [km]	d^a	0.4	0.0

Table 4.1: List of Potential Features for the Experimental Design

Table 4.1. We randomly generate instances abiding to these features. The number of pickup and delivery locations coincides, and is given by *# vehicles*. Velocities and costs are self-explanatory (but costs must be transformed into a distance-based measure since all costs in the objective function ((4.1)) are distance-based). Thus, worker-related costs must be calculated as *costs employee/velocity*. The velocity of the car is restricted by the speed limit within the city, but not as much subject to congestion as during the day, permitting higher velocities. The vehicle costs either only consist of fuel spent, or refer

to the total cost of ownership. Trucks are slower than cars, but within the city, the speed difference is not too high. The cost of the truck is modeled after Wittenbrink (2014), but excluding costs for the driver, and considering that original equipment manufacturers may be able to utilize idle time of their existing fleet of trucks for rebalancing (in the lower value). This number of vehicles does not correspond to the fleet size, but is a measure for the imbalance of the system.

Every night, only a small number of vehicles is selected for rebalancing (e.g., by employing demand prediction, Vosooghi et al. (2017), or simple filtering rules), and the number of vehicles does not change too much between two periods. If the imbalance is less than 10 vehicles, operators will most likely not rebalance every night (resulting in a larger number of vehicles that shall be rebalanced). If the imbalance exceeds 20, operators either rebalance more frequently, or partition the operating area into districts upfront (analogous to Weikl and Bogenberger (2015)). The public transport network is modeled after the underground network of most major European cities with dense lines resulting in a good coverage. It consists of a grid with stops every 1.5km, and the distance between two horizontal (vertical) lines varies between 1.5 and 3km. Changeover between lines takes on average between 6 and 12 minutes (less frequent than during rush hour, as most public transport services are restricted at night when relocation takes place). Since public transport uses a dedicated (congestion-free) network, the average velocity can be as high as 50km/h (e.g., Munich underground), but may be slower for transportation above ground.

As mentioned before, trucks cannot reach all pickup points directly, but the truck must be parked at some access point. Access points are randomly sampled in a circle around the pickup location, using a uniformly sampled distance between 0 and a given maximum d^a that varies between 0.0 and 0.4km. The minimum value is attainable in many North American cities where even city centers are car-friendly. The maximum value is realistic for historic city centers (e.g., in Europe) where few major roads exist that can be several hundred meters apart.

Some features in Table 4.1 have a non-linear impact on the rebalancing cost (e.g., the marginal cost of an additional vehicle is decreasing due to pooling effects), features may jointly impact the rebalancing cost (e.g., worker-related cost per kilometer of biking), and the decision can be driven by the ratio between any two features. Thus, we generate additional features:

$$\log\left(\#\text{ vehicles}\right) = \ln\left(|\mathcal{N}^+|\right) \qquad \ln\left(|\mathcal{N}^+|\right)$$

$$\text{Costs employee bike [€/km]} = \frac{c^e}{v^b} \qquad c^b$$

$$\text{Costs employee walking [€/km]} = \frac{c^e}{v^w} \qquad c^w$$

$$\text{Avg. costs line change [€]} = c^e \cdot t^c \qquad c^l$$

$$\text{Avg. distance walking to station [km]} = d^p \cdot \int_0^1 \int_0^1 \sqrt{x^2 + y^2} dy dx \qquad d^s$$

$$\text{Avg. costs walking closest station [€]} = d^s \cdot c^e \qquad c^s$$

$$\text{Avg. distance closest acc. pt.} = \frac{1}{1 + |\mathcal{N}^+|} \cdot d^a \qquad \min d^a$$

as well as all ratios of two velocities, such as

$$\text{Velocity Bike vs. Walking} = \frac{v^b}{v^w} \qquad \frac{v^b}{v^w}.$$

For the decision trees, we also use the predicted costs of the linear regression as input, i.e., the predicted cost of the mode "car with dedicated helpers" C^{cd} and "car with shared helpers" C^{cs}, the predicted cost of the mode "bike" C^b, the predicted cost of the mode "public transit" C^p, and the predicted cost of the mode "truck" C^t, as well as minimum costs $C^{\min}$ and cost differences between the selected and the cheapest mode (using predicted costs), Δ^{cd}, Δ^{cs}, Δ^b, Δ^p, and Δ^t.

Instances are sampled using a space-filling design for the experiment. In total, 600 design points are generated using a 16-dimensional Sobol sequence. Every design point contains the cost for all individual modes. To determine the costs, we solve the CRPs outlined in Section 4.2 using IBM ILOG CPLEX v12.10 (managed through Java) on a Windows computer with 16GB of memory and four cores. To show that our classification of the best mode is good, and that misclassification mainly occurs if hybridization improves the solution, we additionally calculate the cost of the hybrid solution for 150 of the above design points. Further, instances that were incorrectly classified are replicated. Details on the latter instances will follow during the course of this Chapter.

Table 4.2 lists general descriptive statistics regarding the training and test data. We list the average cost for rebalancing the entire fleet, and how much the instances (with other features) deviate from this average, given by the relative root-mean-squared error (RRMSE). RRMSE (all) can be interpreted as the total variance in the rebalancing cost, not differentiating between random and feature-driven variance. We also report how frequently any mode is the best, second best, etc. mode. Rows Δ report the cost

increase if this mode were chosen in all instances, compared to the minimum costs of an instance, separately for all instances and those instances in which the best mode differs. For the test data, we additionally report how frequently a single-mode route

	Car (Ded.)	Car (Shared)	Bike	Publ. T.	Truck
Training and Test Data (600 instances)					
Avg. cost	235.18	212.95	143.71	220.06	178.45
RRMSE (all)	34.7%	32.3%	33.6%	34.3%	35.7%
# 1	0	4	408	7	181
# 2	1	138	172	113	176
# 3	81	239	19	201	60
# 4	212	219	0	112	57
# 5	306	0	1	167	126
Δ	104.44	82.22	12.56	89.28	89.97
Δ (misclass.)	104.44	82.78	39.25	90.49	128.53
Test Data (150 instances)					
equal to hybrid	(0)	0	88	0	32
used in hybrid	(0)	2	107	0	41
Comparison (20 instances)					
RRMSE (same)	19.4%	13.5%	19.1%	11.4%	12.7%

Table 4.2: Basic Statistics on Training and Test Data

coincides with the hybrid solution (and is, thus, the single best mode, and the benefit of hybridization is non-existent), and how frequently a mode is used as part of the hybrid solution (regardless of whether or not other modes are present). Since any incorrect cost estimate or misclassification can either be due to random cost differences between two instances with the same features or due to imprecisions of the classifiers, we list the RRMSE for 20 identical instances that have average values for all features. This RRMSE (same) can serve as a proxy for the cost differences between any two days in the same carsharing system (random cost variance), and also provides an upper bound on the classification performance.

On average, we see that bike is the best mode, followed by truck. The other modes result in very similar average costs. This is due to a high dependency on the features,

given by the high RRMSE (all) of $\approx 1/3$ across all modes. Together with the low RRMSE (same), this suggests that costs (and thus the mode selection), are in fact driven by features. The reasonably low RRMSE (same) indicates that costs of similar instances are in fact similar, and thus that it makes sense to predict costs as an intermediate step for the feature-based selection. The added costs of 39.25 among instances in which bike is not the best mode is already relatively low, but leaves some room for improvement. Any feature-based algorithm must be evaluated against this benchmark. Surprisingly, the added costs of always choosing mode truck are higher than those for mode car (shared) which suggests that the instances where mode truck is optimal differ in their structure from other instances. The low RRMSE (same) suggests that carsharing operators can use the costs of a sample day to predict the costs of other days (using the same features), and allocate sufficient budgets accordingly. However, bike and cars with dedicated helpers show much higher random errors of 19.1% and 19.4%, respectively. This can be explained by the bipartiteness of the routing graph, but restricts the predictive power of the cost estimates. Then, restricting oneself to the same mode over a longer period of time may be suboptimal on some days due to random realization of the spatial distribution of demand.

4.3.2 Classification and Prediction Algorithms

To decide whether a mode shall be part of the routing, we predict costs of the individual modes using linear regression, and the best mode is predicted via logistic regression and decision trees.

Cost Estimation per Mode using Multiple Linear Regression

We estimate costs for rebalancing the fleet with each of the modes using separate multiple linear regression models. The assumption of multiple linear regression classifiers is that every describing variable (*feature* $f \in \mathcal{F}$ with value v_f^i for instance $i \in \mathcal{I}$) has linear influence on the objective. The total costs C_m of mode m are then approximated by

$$C^m = \sum_{f \in \mathcal{F}} \beta_f^m \cdot v_f^i + \beta_0^m \tag{4.39}$$

where β_f^m is a weight on each feature f of mode m, and β_0^m is the intercept of the linear function. Given a set of numeric or boolean features, the regression outputs β values that minimize the error, given by ordinary least squares (OLS), i.e., by the squared

difference between observed and predicted feature value (on the training data). We exclude features which cannot have an influence on a mode (e.g., velocity of bike for truck-based relocation) to reduce the effects of overfitting to the training data.

Table 4.3 lists the β values for each feature. Features which we excluded a-priori are marked with "−", the value $\beta_f^m = 0$ indicates that the linear regression ignored this feature since it does not improve the cost estimate (when minimizing the Akaike criterion). The β values reflect basic intuition: The more cars an operator has to relocate, the more expensive is the entire process. Due to the combination of linear and logarithmic components in the number of vehicles, ranking the modes by influence of additional vehicles is not straight-forward, but in general, total costs of the mode truck grow slower in the number of locations as well as velocities than costs of the other modes. This indicates that in a more imbalanced system (with more rebalancing necessary), operators resort to a less environmentally friendly mode (truck). Alternatively, a higher imbalance can also be a result of a larger operating area, providing mobility to a larger share of the population. Public authorities can counteract the usage of trucks in presence of larger imbalances by encouraging more frequent rebalancing, and operators can reduce the imbalance by incentivizing round trips and user-based rebalancing. The hourly wage (c^e) and costs for the car (c^v) have the highest impact on the mode car, followed by bike and public transit (both of which also have combined features containing c^e), and are lowest for the mode truck. A longer shift length slightly decreases the costs (as fewer routes start from the depot on average, increasing the benefits of pooling). Higher velocities decrease the total costs. This explains why operators rebalance during the night when average velocities are higher due to lower congestion. This is beneficial for the society, as rebalancing then does not contribute to congestion during the rush hours (even though the number of vehicles that are rebalanced is not exceedingly high, searching for a parking spot can substantially contribute to congestion). While it seems contradictory at first that an increasing average distance to the closest station (d^p) decreases the costs of mode public transit, this effect is alleviated by the increasing effect of c^s.

Table 4.4 shows the out-of-sample performance of the classifier. A common metric for a classifier is listing the cost of misclassification as well as listing the probability of misclassification. The cost can be represented in two ways: for each linear regression, we can report the average deviation between actual (C_i^m) and predicted ($\hat{C}_i^m$) costs

$$\frac{\sum_{i \in \mathcal{I}} |\hat{C}_i^m - C_i^m|}{|\mathcal{I}|},$$

Feature	Car (D)	Car (S)	Bike	Publ. T.	Truck
$\lvert\mathcal{N}^+\rvert$	5.055	9.8453	0	7.9445	8.39
$T^{\max}$	-0.0584	0	-0.0513	-0.146	-0.1013
v^v	-7.4095	-6.4485	-1.9165	-1.8213	0
c^v	425.0377	336.6039	164.9144	160.9452	79.7144
θ	-	-3.4498	-	-	-
v^b	-	-	0	-	-
v^w	-	-	-	0	0
v^p	-	-	-	-0.8535	-
d^p	-	-	-	50.453	-
t^c	-	-	-	0	-
c^e	14.3235	12.5612	3.6474	5.7231	0.9454
v^t	-	-	-	-	-1.4956
c^t	-	-	-	-	90.9407
κ	-	-	-	-	-12.6297
$\lvert A_i\rvert$	-	-	-	-	0
d^a	-	-	-	-	0
$\ln(\lvert\mathcal{N}^+\rvert)$	67.243	0	84.9794	40.2385	-18.8213
c^b	-	-	86.7523	-	-
c^w	-	-	-	5.9394	0
c^l	-	-	-	0	-
d^p	-	-	-	-133.4148	-
c^s	-	-	-	37.6697	-
$\min d^a$	-	-	-	-	0
intercept	-56.674	29.8112	-177.9893	-109.6515	2.4168

Table 4.3: β Values of Linear Regression Models per Mode

and the relative root-mean-squared error (RRMSE)

$$\text{RRMSE} = \sqrt{\frac{\sum_{i\in\mathcal{I}}\left(\hat{C}_i^m - C_i^m\right)^2}{|\mathcal{I}|}} \cdot \frac{1}{\bar{C}^m},$$

where $|\mathcal{I}|$ is the number of instances, and $\bar{C}^m$ is the average cost over all instances. Another measure for the quality of an estimator is the coefficient of determination (R^2)

$$R^2 = 1 - \frac{\sum_{i\in\mathcal{I}} \left(\hat{C}_i^m - C_i^m\right)^2}{\sum_{i\in\mathcal{I}} \left(\bar{C}^m - C_i^m\right)^2}.$$

We also report the costs arising from incorrectly choosing one mode if another would be preferred. In this case, the cost of misclassification only corresponds to the best mode, rather than the entire ordering. We list how frequently any mode was incorrectly chosen as best mode (false positives (FP)) or incorrectly chosen as inferior (false negatives (FN)). The difference between actual and predicted costs increases with the actual costs

	Car (D)	Car (S)	Bike	Publ. T.	Truck	Overall
Cost deviation	37.26	27.92	22.50	26.83	19.57	
RRMSE	21.0%	16.5%	20.8%	16.3%	14.8%	
R^2	44.8%	65.4%	34.5%	68.1%	78.5%	
FP	0	2	2	2	6	8%
FN	0	1	9	0	2	8%
Δ cost (#1)	-	20.64	14.39	9.40	11.64	13.84

Table 4.4: Out-of-Sample Accuracy Metrics Linear Regression

(which is to be expected). This entails that the absolute cost of misclassification, that is the excess cost operators have to cover when choosing a suboptimal mode, increases in the instance size (the relative cost of misclassification remains approx. constant). The number of misclassified instances is very low with 12/150 instances (8.0%), and the impact of the misclassification is comparatively low. If a misclassification occurs, this adds costs of 13.84 on average, and is slightly lower if the mode "truck" is involved. This means that on average, operators spend less than 14€ daily – or less than 10% of the total rebalancing cost – due to misclassification which is a reasonable price for a simple approach. Also, linear regression adds fewer costs than a simple selection of mode "bike" in all instances (compared to Table 4.2). Authorities can encourage a more environmentally friendly mode by covering this cost difference. Further, misclassification occurs more frequently if hybridization among the modes "bike" and"truck" improves the solution, and both modes are part of the hybrid solution (7/12 misclassified instances benefit from hybridization, vs. 20% in the entire test data). Assuming

that the necessity for hybridization is Bernoulli-distributed,this difference is statistically significant for $\alpha = 0.99$ using a two-tailed two-sample z-test. Misclassified instances are commonly "borderline" instances; multiple modes are promising candidates, and hybridization among these modes improves the solution.

Overall, this classifier based on linear regression allows operators to easily establish the best mode. The classifier performs well with respect to accurary and cost of misclassification as metrics. However, linear regression only helps in establishing features that drive rebalancing costs, and features driving modal choice cannot readily be found.

Prediction of the Best Mode using Multinomial Logistic Regression

Using (multinomial) logistic regression, we estimate the probability mode m minimizes the cost for instance i. Logistic regression estimates the weights of a linear regressor, and uses the logit of this linear regressor to predict the odds that a mode is best (probability that mode m is best relative to the probability that the reference mode is best). As such, multinomial logistic regression maximizes the multinomial log-likelihood (the probability of predicting the correct class)

$$L = \sum_{i \in \mathcal{I}} z_{im} \cdot \ln\left(P_m\left(v\right)\right)$$

where $z_{im} = 1$ iff the best mode for instance i is m (0 otherwise). $P_m\left(v\right)$ is the predicted probability for mode m, given the feature vector v_i of instance i

$$P_m\left(v_i\right) = \frac{\exp\left(\sum_{f \in \mathcal{F}} v_f^i b_f^m\right)}{1 - \exp\left(\sum_{f \in \mathcal{F}} v_f^i b_f^m\right)} \qquad \forall m \in b, p, c$$
$$P_t\left(v_i\right) = 1 - P_b\left(v_i\right) - P_p\left(v_i\right) - P_c\left(v_i\right).$$

b refers to weights on features (per mode). The classifier chooses the mode with the highest probability.

Tables 4.5–4.6 list the b values for the modes car (shared), bike, and public transit. The mode "truck" serves as reference category and is, thus, not listed. The out-of-sample accuracy metrics can be found in Table 4.7. Using this classifier, 14/150 instances are misclassified (90.7% accuracy), and the cost deviation is low with 16.09. Thus, the logistic regression classifier performs better than using the same mode "bike" on all instances, but does not perform as good as the linear regression classifier.

Feature	Car (S)	Bike	Publ. T.
$\lvert\mathcal{N}^+\rvert$	-1.4869	2.9413	-2.7651
$T^{\max}$	-0.0247	-0.0039	-0.15
v^v	2.3782	0.2399	-1.7333
c^v	-241.2598	-48.0656	-163.4323
θ	-0.6476	0.1085	6.6672
v^b	-0.9895	0.3481	-1.8377
v^w	-1.0752	-0.9432	17.4514
v^p	1.5861	0.0161	1.8658
d^p	2.9968	-6.9205	-95.0621
t^c	198.52	91.5958	64.3535
c^e	-2.9399	-0.0496	0.6431
v^t	0.6625	0.1045	0.1217
c^t	27.0828	13.766	27.5977
κ	-4.353	-2.05	-3.3037
$\lvert A_i\rvert$	-0.0585	1.2758	10.642
d^a	-43.4125	-32.5253	23.7726

Table 4.5: b Values of Multinomial Logistic Regression (Basic Features)

To establish the importance of the individual features, we report the odds ratio with slight abuse of notation. Since the order of magnitude of the values differs substantially, we instead report the odds ratio if the feature value changes by 10% of the difference between minimum and maximum value, from the 45% percentile to the 55% percentile. Besides, we have multiple interaction terms on different subsets of variables (e.g., costs of an employee per hour influence four higher-order features, and the velocity of the mode bike influences five second-order terms). Rather than reporting the odds ratio for all combinations of minimum and maximum characteristics for all other variables, we assume that those features f' with whom some explanatory variable interacts, are set to their average value. Exemplarily,

$$OR_f^{m\star} = \exp\left(v_f^{\min} - v_f^{\max}\right) b_f^m + \left(v_f^{\min} - v_f^{\max}\right) v_{f'}^{\text{avg}} b_{\hat{f}}^m$$

shows the adapted odds ratio for a feature f that interacts with feature f' as a multi-

Feature	Car (S)	Bike	Publ. T.
$\ln(\lvert\mathcal{N}^+\rvert)$	12.1653	-42.943	14.1244
c^b	24.4473	-12.676	8.5237
c^w	4.78	0.9899	-8.5303
c^l	-7.416	-5.3704	-15.5641
d^p	7.6325	15.6181	82.7444
c^s	4.8882	0.0255	10.0205
$\min d^a$	53.8532	129.3161	17.6215
$\frac{v^b}{v^w}$	-6.3868	1.4339	4.5752
$\frac{v^b}{v^v}$	16.0568	-12.3153	-132.1201
$\frac{v^b}{v^p}$	40.8518	17.8401	23.0729
$\frac{v^b}{v^t}$	-3.4386	-17.8907	-36.734
$\frac{v^w}{v^c}$	82.1028	48.3739	69.8523
$\frac{v^w}{v^p}$	-22.191	-47.3956	-199.2226
$\frac{v^w}{v^t}$	-35.3511	55.155	158.4973
$\frac{v^v}{v^p}$	-3.9882	7.6667	-7.6918
$\frac{v^v}{v^t}$	-0.3903	-1.4769	-0.5563
$\frac{v^p}{v^t}$	6.6265	4.9853	-3.2219
intercept	-226.1646	39.7793	37.299

Table 4.6: b Values of Multinomial Logistic Regression (Advanced Features)

plicative term such that $\hat{f} = f \cdot f'$. Table 4.8 lists the odds ratio for all modes except truck. Since the number of instances for which car or public transit is the preferred mode is very little, the odds ratios grow very large, but remain insignificant at a significance level of 0.05. Thus, the odds ratios have limited power in describing the influence of individual parameters on the probability of choosing public transit and car as modes. We therefore focus on bike vs. truck, and only consider the other modes in a side remark. Increasing wages reduce the probability of mode bike the most, followed by c^v. This suggests that two societal goals are competing: While high wages are beneficial for workers (and increase equality since rebalancing is commonly a low-income job), high wages increase the probability of using the less environmentally friendly mode "truck". Interestingly, an increasing velocity of public transit and an increasing distance between

Metric	Car	Bike	Publ. T.	Truck	Overall
Incorrect # 1 (FP)	0	6	3	7	9.3%
Incorrect not # 1 (FN)	2	10	0	4	9.3%
Δ cost (# 1)	15.17	14.64	20.68	11.26	16.09
Classified as Car	0	0	0	0	
Classified as Bike	2	97	0	4	
Classified as Publ. T.	0	3	0	0	
Classified as Truck	0	7	0	37	

Table 4.7: Out-of-Sample Accuracy Metrics of Mode Usage via Multinomial Logistic Regression

Feature	Car (S)	Bike	Publ. T.
$\lvert\mathcal{N}^+\rvert$	2.56	1.08	2.25
$T^{\max}$	0.03	0.91	0.55
v^v	0.81	1.98	23.37
c^v	0.04	0.38	0.01
θ	3.79	1.02	0.88
v^b	0.001	2.19	0.08
v^w	262.49	1.01	0.24
v^p	400.6	0.53	25.08
d^p	0.0001	0.61	24.74
t^c	0.18	1.12	2.39
c^e	0.32	0.23	0.66
v^t	1.12	0.81	2.50
c^t	31.49	5.59	29.53
κ	0.52	0.66	0.42
$\lvert A_i\rvert$	8.04	0.93	0.86
d^a	3.09	0.99	0.3

Table 4.8: Odds Ratio derived from Multinomial Logistic Regression

lines exert a strong negative influence on the probability of choosing mode bike as well.

The former can be explained by switching to public transit, the latter appears to be random, and the randomness can be explained by using split variables: The distance between lines decreases the probability of mode bike, while the average distance to the closest station has a strong positive influence. An increasing capacity of the truck has a larger impact than the velocity of the truck. This can be interpreted as using the mode truck despite being slow if pooling benefits are sufficient (due to a larger capacity). If a truck is used (regardless of the capacity), the society benefits from improved pooling and shorter routes. Local authorities should therefore permit larger trucks if trucks are used regardless of the capacity (that the operator will most likely use), but can impose restrictions on the size of trucks in the city center to deter operators from switching from bike to truck. If the costs for the truck increase, this increases the probability that the mode bike is chosen. The influence of shift length on modal choice is miniscule, and the capacity of the car, velocity of walking, and the number of access points only exert a small influence.

Prediction of the Best Mode using Decision Trees

A decision tree consists of multiple consecutive decisions which split the data sample along boolean criteria. Every decision increases the similarity of classified variables (increases the probability that an instance in the sample is of class m). The classification objective is to maximize the number of correctly classified instances (that is, correct prediction of the best mode), using the normalized information gain as a proxy. Thus, similarity within any subtree and dissimilarity between different subtrees is maximized.

We consider three different decision trees: (i) only using the features of the city and fleet (including feature combinations), (ii) using features and real routing costs in the training data, and evaluating against cost estimates from linear regression, and (iii) using features and cost estimates. The first approach is the simplest, and can be applied without any additional predictions. As such, it may prove best in determining influencing factors (if the accuracy permits it). Both other approaches are upfront expected to yield higher accuracies, but cost estimates may dominate other features. Approach (ii) is better than approach (iii) if cost estimates from the linear regression are not biased, since the decision tree becomes more precise. Vice versa, if cost estimates are not true (e.g., due to missing some explanatory variables, or since some explanatory variables do not have an approximately linear impact on the objective), approach (iii) is more accurate than approach (ii). Thus, approaches (ii) and (iii) jointly allow us to determine if the cost estimates are biased, and it is unclear upfront which approach performs better.

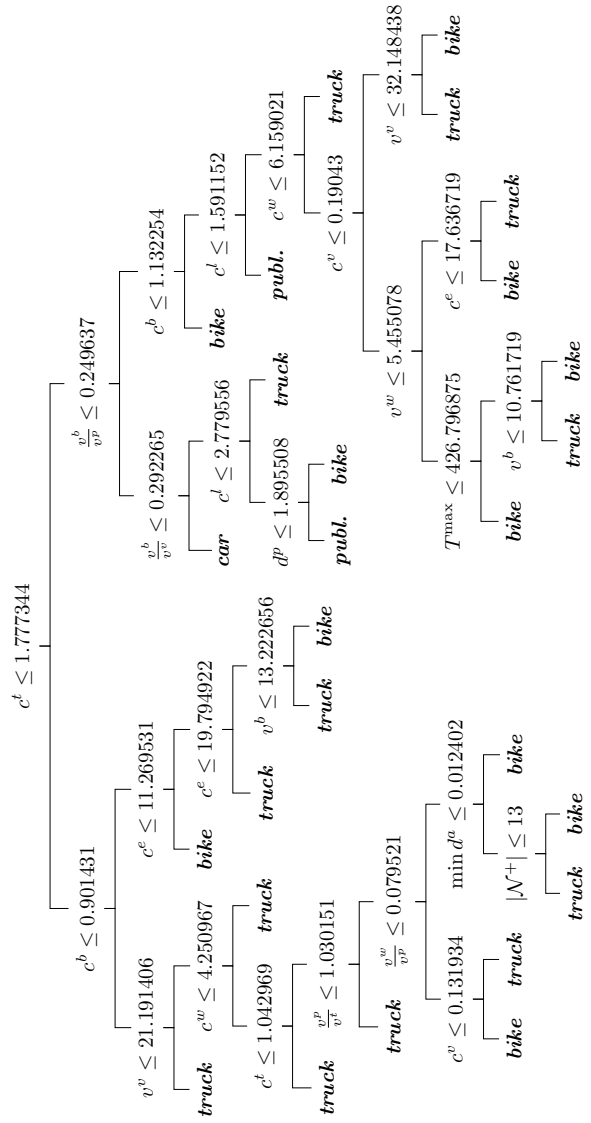

Figure 4.5: Decision Tree with Features only

Figures 4.5–4.7 depict the decision trees. The decision tree is largest if only the basic features are used (with 18 distinct features and 27 leaves, vs. 3 features and 4 leaves in approach (ii) and 13 features and 17 leaves in approach (iii)). This makes it most prone to misclassification due to overfitting which we counteract by increasing the minimum number of instances per leaf. For example, the tree uses the features $\frac{v^p}{v^t}$ and $\frac{v^w}{v^p}$ in a subtree that never decides for the mode public transit. However, having less features (e.g., by increasing the minimum number of items per leaf) reduces the number of classes to two (bike and truck), and thus decreases the classification performance.

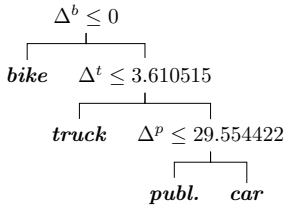

Figure 4.6: Decision Tree with Features and True Costs

Thus, additional cost estimation features are necessary to select the cost-minimizing

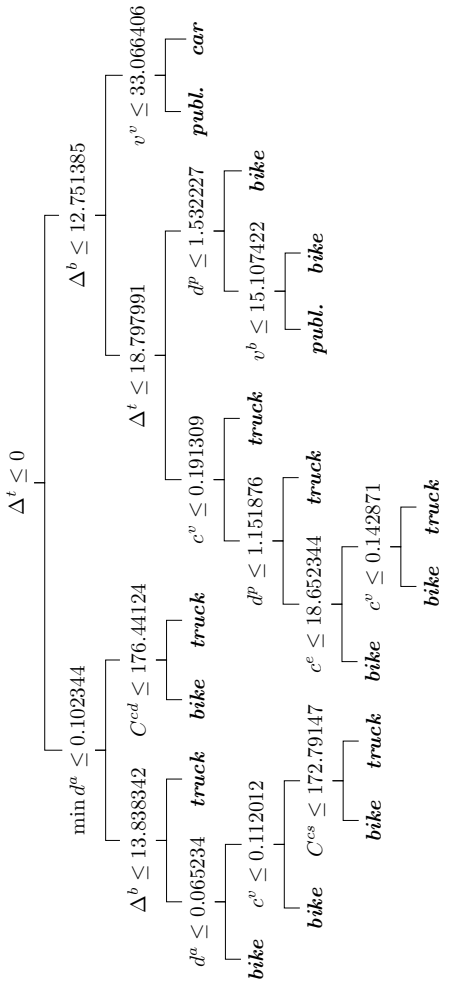

Figure 4.7: Decision Tree with Features and Estimated Costs

mode. Approach (ii) results in a very small tree that only uses the cost differences Δ^m between some mode m and the best mode as features (trained on actual costs and

evaluated against cost estimates) which suggests that the feature selection of the linear regression is already reasonable, and that the features (and higher-order features) have an approximately linear influence on the rebalancing cost (and thus indirectly on the modal choice). Approach (ii) does not provide significant insights into features driving the modal choice. Approach (iii) also overfits, but not as much as approach (i).

Table 4.9 lists the out-of-sample accuracy metrics (how frequently some mode was incorrectly predicted as best mode, or incorrectly as not best mode, cost increase due to misclassification, and confusion matrix) for all three approaches. We see that approach

Appr.	Metric	Car	Bike	Publ. T.	Truck	Overall
(i)	Incorrect # 1 (FP)	0	11	0	14	16.7%
	Incorrect not # 1 (FN)	2	14	0	9	16.7%
	Δ cost (# 1)	15.17	29.49	-	28.25	28.34
	Classified as Car	0	0	0	0	
	Classified as Bike	2	93	0	9	
	Classified as Publ. T.	0	0	0	0	
	Classified as Truck	0	14	0	32	
(ii)	Incorrect # 1 (FP)	0	2	5	7	9.3%
	Incorrect not # 1 (FN)	2	10	0	3	9.3%
	Δ cost (# 1)	11.58	10.71	10.36	11.10	10.84
	Classified as Car	0	0	0	0	
	Classified as Bike	0	97	0	2	
	Classified as Publ. T.	2	3	0	0	
	Classified as Truck	0	7	0	39	
(iii)	Incorrect # 1 (FP)	0	8	3	7	12%
	Incorrect not # 1 (FN)	2	10	0	6	12%
	Δ cost (# 1)	15.17	19.11	19.73	19.58	19.11
	Classified as Car	0	0	0	0	
	Classified as Bike	2	97	0	6	
	Classified as Publ. T.	0	3	0	0	
	Classified as Truck	0	7	0	35	

Table 4.9: Out-of-Sample Accuracy Metrics of Mode Usage via Decision Trees

(i) only uses two classes on the test data (bike and truck), and the other approaches seldomly classify instances as mode "car" or "public transit", but this classification is never correct. Thus, features driving the choice of other modes except bike and truck cannot be determined well using decision trees. The number of incorrectly classified instances is reasonably low for approach (iii) (12%), but higher than just using the linear regression. Approach (ii) only uses the information from the linear regression, and thus results in very similar classification performance.

Approaches (i) and (iii) only classify instances as mode "car" or "public transit" if both bike and truck are disadvantageous due to very low velocities or very high costs. To determine the modal split between truck and bike, we observe that the split is mainly driven by the costs for the truck c^t, Δ^t and the cost for the employee c^e (and higher-order features depending on c^e). Velocities and their relative differences also play a role in determining the best mode, but their impact is less pronounced than the costs. Since velocities are less relevant, using foldable bikes (with a lower velocity) is a reasonable choice for a rebalancing mode.

4.3.3 Decision on Hybridization

Table 4.2 shows that hybridization improves the solution in 20% of all instances in the test data. A closer look at the test data reveals a tendency that misclassified instances have very similar costs for multiple modes (linear regression classifier), or a higher probability for multiple modes in logistic regression.

We use these insights to decide whether or not to use multiple modes. Using linear regression classifier, we decide to use a mode if its predicted cost difference to the best mode $\Delta^m, m \in \{c, b, p, t\}$, is in the range RRMSE (same). In the logistic regression, we decide for hybridization if a suboptimal mode is chosen with a probability of at least 0.1%.

BBoth methods individually result in a false negative rate of 30.0% and 23.3% (missing 9 and 7 of 30 instances that shall be mixed). In the missed instances, bike and truck are mostly mixed with public transit or car, and the second mode only covers a small part of the rebalancing route. They result in false positives (classification as hybrid, but single mode sufficient) in 51.7% (multiple linear regression) and 40.8% (multinomial logistic regression) of all 120 test data instances that do not benefit from hybridization, respectively. Thus, the classifier based on logistic regression is better qualified to predict if an instance shall be rebalanced with multiple modes. However, using both methods combined by deciding for hybridization whenever one of the methods suggests to do

so, only four instances were misclassified, but the false positive rate increases to 59.2% in the combined case with the same cutoff values. In these instances, the very cost-efficient mode bike or truck is combined with a different mode on a small part of the tour. While false positives are less problematic for operators since they do not increase costs, only computation time, operators might aim at reducing the false positive rate at the expense of the false negative rate. Table 4.10 reports the false negative and false positive rates for different cutoff values. Clearly, a very high cutoff value is ideal if a single method is chosen. If both methods are combined, it is possible to reduce the cutoff values and subsequently the false positive rate and still achieve a low false negative rate. Determining individual features that drive the decision for hybridization is not straightforward, since single features do not predict well if costs for different modes are close, or more than one mode has a high probability. We observe that hybridization is more frequently optimal on larger instances, as in larger instances, operators retain pooling benefits in either class, and the average cost difference between bike and truck grows smaller for larger instances.

		Lin. Reg.		Log. Reg.		Both	
Cutoff (lin)	Cutoff (log)	FPR	FNR	FPR	FNR	FPR	FNR
0.6	0.8	43.3	43.3	10.0	66.7	47.5	23.3
0.7	0.9	44.2	36.7	14.2	56.7	49.2	16.7
0.8	0.95	45.0	33.3	21.7	50.0	50.8	13.3
0.8	0.99	49.2	33.3	30.0	33.3	54.2	13.3
0.8	0.999	51.7	30.0	40.8	23.3	59.2	13.3

Table 4.10: False Positive and False Negative Rates in % for Different Cutoff Points

4.3.4 Replicated Test

Whilst the rebalancing itself is an operational problem, the modal choice is of a more strategic nature: The operator must buy trucks and bikes, have sufficient vehicles available at the depot, and hire enough personnel to perform all necessary operations (and the number of workers can depend on the modal choice). Thus, a single replication is not sufficient to make a statement about the optimality of the modal choice. We replicate the incorrectly classified instances 20 times each and record the same metrics as before (listed in Table 4.11). In total, there are 8 different feature combinations that resulted

in misclassification in at least 4 out of 5 algorithms. First, we observe that of those in-

Alg.		Car (S)	Bike	Publ. T.	Truck	Overall
Lin Reg	Incorrect #1	0	0	13	84	60.6%
	Incorrect not #1	1	96	0	0	60.6%
	Δ cost (#1)	25.56	21.93	8.60	24.04	21.97
Log Reg	Incorrect #1	0	8	0	84	57.5%
	Incorrect not #1	1	84	7	0	57.5%
	Δ cost (#1)	23.88	22.75	7.10	24.04	22.75
Tree (i)	Incorrect #1	0	0	0	104	65.0%
	Incorrect not #1	1	96	7	0	65.0%
	Δ cost (#1)	68.39	26.21	44.67	27.85	27.85
Tree (ii)	Incorrect #1	0	0	8	91	61.9%
	Incorrect not #1	1	96	2	0	61.9%
	Δ cost (#1)	68.39	23.38	14.13	25.66	24.18
Tree (iii)	Incorrect #1	0	8	0	84	57.5%
	Incorrect not #1	1	84	7	0	57.5%
	Δ cost (#1)	23.88	22.75	7.10	24.04	22.75

Table 4.11: Out-of-Sample Accuracy Metrics (Replicated)

stances, none is misclassified in all replications, hybridization is necessary in at least one replication in all instances except one, and the best mode differs between replications of the same instance in all but one instances, suggesting that the number of correctly classified instances (56/160 - 68/160 depending on the algorithm) is reasonable. While operators cannot expect that the predictors return the best mode for all days (with a random realization of pickup and delivery demand), it indicates that the modal selection is not biased.

4.3.5 Application to Sample Cities

To show the validity of our approach, we apply it to the sample cities of Milan (Italy), Munich (Germany), and Toronto (Canada). These are cities for which case studies exist that mention a relocation mode.

Parameters

Table 4.12 lists the parameter values used for Milan (Italy), Munich (Germany) and Toronto (Canada). Whenever possible, the values are set similar to Bruglieri et al. (2014b), Nourinejad et al. (2015), and Weikl and Bogenberger (2015), and the remaining open feature values are set to realistic values.

Feature		Milan	Munich	Toronto
No. Vehicles	$\lvert\mathcal{N}^+\rvert$	15	20	12
Shift Length [min]	$T^{\max}$	300	240	240
Velocity Car [km/h]	v^v	25	20	20
Costs Car [€/km]	c^v	0.3	0.3	0.1
Capacity Car	θ	2	4	2
Velocity Bike [km/h]	v^b	18	18	15
Velocity Walking [km/h]	v^w	5	5	5
Velocity Publ. [km/h]	v^p	35	40	30
Distance between lines [km]	d^p	2.2	2	3
Time Changeover [h]	t^c	0.18	0.15	0.15
Costs Employee [€/h]	c^e	13	15	10
Velocity Truck [km/h]	v^t	15	15	20
Costs Truck [€/km]	c^t	1.4	1.5	1.3
Capacity Truck	κ	3	3	3
No. Access Pts.	$\lvert A_i\rvert$	4	4	2
Max. Distance Access Pt. [km]	d^a	0.4	0.4	0.2

Table 4.12: Feature Values for Example Cities (Milan, Munich, Toronto)

Insights

Table 4.13 lists the classification for all classifiers and cities. Using linear regression, we support the modal choice reported in the case studies (bike), but also show that truck is not much more expensive (Table 4.14). Operators may therefore frequently benefit from hybridization. The other algorithms also frequently suggest to use trucks, supporting the decision to use different modes (bike and truck) in the same city.

	Lin. Reg.	Log. Reg.	Dec. (i)	Dec. (ii)	Dec. (iii)	Hybrid (Lin)	Hybrid (Log)
Milan	B	T	B	B	T	✗	✓
Munich	B	T	T	B	T	✓	✓
Toronto	B	B	T	B	T	✓	✓
Milan (+10%)	B	T	B	B	T	✗	✓
Munich (+10%)	B	T	B	B	T	✓	✓
Toronto (+10%)	B	B	B	B	T	✓	✓
Milan (-10%)	B	T	B	T	T	✗	✓
Munich (-10%)	T	T	T	T	T	✗	✓
Toronto (-10%)	B	T	T	T	B	✓	✓

Table 4.13: Predicted Best Mode per Classifier

	Car (D)	Car (S)	Bike	Publ. T.	Truck
Milan	312.2	273.7	148.4	217.9	150.1
Munich	426.0	373.3	202.4	304.2	203.7
Toronto	194.6	171.4	93.3	162.6	99.9

Table 4.14: Predicted Costs (Multiple Linear Regression) for Example Cities

The decision for hybridization is stable against slight input manipulation, while the choice of the optimal mode changes if the parameters are varied by +/-10% (reported in the second and third block in Table 4.13). This also indicates that operators should decide for hybridization if multiple modes result in reasonable costs, or a reasonable choice probability.

4.4 Conclusion

We study how vehicles shall be relocated in a carsharing system: Workers can either drive the car themselves, or load them onto a truck. To make relocation profitable, operators "pool" multiple relocation operations. Then, the worker somehow has to continue to the next vehicle. She can do so by biking, using public transportation, or – if there are

multiple relocation workers – by getting a lift by a colleague.

Not all modes are necessary in all cities, and hybridization adds a substantial computational overhead. We develop algorithms to determine the best mode in a city and for a given fleet, and use these to establish key drivers for the modal choice. (i) Costs per mode are predicted using linear regression, and the cheapest mode is used. (ii) The probability of a mode being cost-minimizing is predicted using logistic regression, and the mode with the highest probability is used. (iii) The best mode is predicted using one of three different decision trees. Among the five classifiers, using the cheapest mode (cost estimate using linear regression) results in the highest accuracy and lowest added cost due to misclassification.

We find that in most instances, the optimal mode is bike or truck. The most cost-efficient mode is bike if wages for workers are not too high, and costs for vehicles (fuel, wear, tear) are low. Thus, the societal goals of low emissions (rebalancing by bike) and high salaries (resulting in truck as cost-minimizing mode), compete, as well as high prices for fuel (and thus tax income) and low emissions. If the high imbalance in the system (and thus the high number of vehicles that shall be rebalanced, and the optimality of mode "truck") is a consequence of a larger operating area, the societal goals of providing mobility to a larger part of the population and environmentally friendly rebalancing may also compete. Since the optimality of the mode "truck" is also affected by the accessibility (given by the average distance between a parking spot and the location of the vehicles, as well as the maximum capacity that still allows the truck to maneuver in the city), rebalancing by bike is more common in dense (often historic) city centers than in sub-urban regions. These competing goals open a new avenue for future research: How can public authorities incentivize operators to rebalance using environmentally friendly modes, while reaching a high availability and paying fair wages to their employees? Public transit and car are only the preferred mode if bike and truck are impeded: If accessibility by truck is very low, and if bikes travel very slowly. We find that in most instances, hybridization is in fact not necessary.

Our analysis is restricted to operator-based rebalancing. In future research, one can use a similar analysis might allow operators to determine to what extent they should incentivize users to rebalance, and which to rebalance using one of the aforementioned modes, and which vehicles should be rebalanced by users. This process can then serve as a preprocessing step to the routing.

Chapter 5

It's All in the Mix: Technology Choice for Vehicle Sharing

With the rise of driverless vehicles, operators of vehicle sharing systems might benefit from introducing them into their fleet mix. Whilst optimization models for fleet sizing and rebalancing exist for both the driverless and human-driven case, no model considers the value of driverless vehicles in sharing systems due to operational benefits despite higher investment costs. This Chapter studies a new technology choice and mix problem with stochastic demands. We show when vehicle sharing operators benefit from introducing driverless vehicles in their fleet mix, and when driverless vehicles are necessary for profitable operations. The problem is formalized as a fleet sizing and composition problem that consists of two stages: The first stage problem (strategic decision) focuses on fleet size and composition. We formulate the second stage problem (rebalancing; operational decision) as a semi-Markov decision problem and as a closed queueing network; both reformulations are solved using linear programming. Many vehicle sharing systems can benefit from driverless vehicles, and frequently from using a mix of driverless and human-driven vehicles. If the contribution margin for driverless vehicles exceeds the contribution margin for human-driven vehicles, driverless vehicles are imperative to obtain profitable vehicle sharing systems. Mixed fleets are beneficial if contribution margins of the two vehicle types are similar (e.g., in carsharing), and if fixed costs vary substantially between vehicle types. In case studies for DiDi (China) and New York City, driverless vehicles are beneficial, while mixed fleets only increase profits in the DiDi case study. Even if demand is relatively balanced, operators shall consider introducing driverless vehicles in their fleet. Further, once driverless vehicles are used, operators can enter new markets that were previously unprofitable.

5.1 Introduction

Mobility systems currently face two paradigm shifts in parallel: the advent of a "sharing economy" including vehicle sharing; and driverless vehicles (Qi and Shen, 2019). Vehicle sharing comprises different types of shared mobility, including one-way carsharing (e.g., ShareNow) and ride-hailing (e.g., Uber). Vehicle sharing constitutes an environmentally friendly alternative to private car ownership, permits one-way trips, is more flexible than public transportation, and cheaper than taxis (Martin and Shaheen, 2011). For many users, in particular in urban areas, renting vehicles for individual trips is more cost-efficient than buying a car (Baptista et al., 2014). However, vehicle sharing systems are so far often unprofitable due to high operational costs and low utilization, and either require subsidies or cease operations (like AutoLib in Paris).

Currently, first vehicle sharing operators experiment with driverless vehicles in urban environments (Lyft, 2020), expecting advantages for both users and operators. For users, renting driverless vehicles is often more beneficial than owning them (Wall Street Journal, 2017), resulting in increasing demand for vehicle sharing. For operators, driverless vehicles come with significant operational benefits with regards to fleet rebalancing. Most user flows are imbalanced in the sense that in a given time interval and at a given station, the number of arriving users does not equal the number of departing users, e.g., due to commuter traffic (Huang et al., 2018). To ensure that many users can rent a vehicle at any given location, vehicles must be rebalanced, i.e. driven from a low-demand station to a high-demand station. With driverless vehicles, rebalancing costs can be reduced significantly. Thus, vehicle sharing firms are most likely among the first to use driverless vehicles (Boston Consulting Group, 2020). However, driverless vehicles are costly, since the upfront investment is substantially larger than for human-driven vehicles.

Against this background, we study the technology choice and mix problem for partially driverless vehicle sharing systems consisting of finding the optimal number of human-driven and driverless vehicles, as well as the rebalancing strategy. We focus on operator-based rebalancing rather than user-based rebalancing via incentive mechanisms (e.g., Ströhle et al. (2019)) or dynamic pricing (e.g., Bimpikis et al. (2019)), since user-based rebalancing does not dissolve demand imbalances entirely, and operator-based rebalancing is more profitable if the willingness of users to move is low (Guda and Subramanian, 2019).

5.1.1 Contribution

We present a novel technology choice and mix problem to study the benefits of introducing driverless vehicles and the benefit of mixed fleets. Specifically, our contribution is 3-fold: (i) For the second stage rebalancing problem considering multiple vehicles, we provide two model formulations for the dynamic carsharing rebalancing problem with driverless and human-driven vehicles, based on Semi-Markov Decision Processes (SMDPs) and closed queueing networks by means of a fluid-based approximation respectively. (ii) For the first stage decision, we introduce a solution algorithm based on bound-and-enumerate. (iii) We investigate the benefit of driverless vehicles in vehicle sharing systems, i.e., we compare the highest achievable profits under mixed and human-driven fleets. As such, we support operators of vehicle sharing systems in their choice whether to use driverless vehicles. We validate our results on two sets of artificial instances and two real-world case studies with data from DiDi (China) and New York City. We find that in most case studies and instances, driverless vehicles are beneficial, and mixed fleets often result in profits increasing further. Driverless vehicles may be necessary to operate any fleet profitably.

5.1.2 Outline

Section 5.2 introduces the models and solution algorithms for first and second stage decisions. In Section 5.3, we introduce artificial instances and real-world case studies. Section 5.4 presents results including insights into the value of driverless vehicles both on artificial instances and real-life case study data. Section 5.5 concludes the Chapter.

5.2 Methodology

We focus on a strategic fleet composition and technology choice decision from the perspective of a vehicle sharing operator. Accordingly, this decision comprises fleet procurement and fleet rebalancing. Formally, a vehicle sharing provider decides how many driverless (m^d) and human-driven (m^h) vehicles to procure. $\alpha \in \{d, h\}$ refers to the vehicle type. Formally, the operator maximizes

$$\max_{m^d, m^h} \Pi\left(m^d, m^h\right) = \pi\left(m^d, m^h\right) - m^d \cdot F^d - m^h \cdot F^h. \tag{5.1}$$

For each vehicle, she faces annuities of F^d and F^h per period of time, respectively. These annuities are higher for driverless vehicles than for human-driven vehicles ($F^d \geq F^h$). $\pi\left(m^d, m^h\right)$ refers to the expected operational profits, i.e., contribution margins for served customers minus rebalancing costs. Customers use the vehicle sharing system, and the operator rebalances her fleet to maximize expected operational profits $\pi\left(m^d, m^h\right)$. Vehicles can be rented from n stations ($I = 1 \ldots n$) which may also approximate districts in a free-floating sharing system. Customers arrive at stations $i \in I$ following a Poisson process with arrival rate λ_i. We assume that customers randomly select one of the available vehicles. In carsharing systems, this is analogous to users selecting the closest available vehicle. In ride-hailing systems, operators assign vehicles, and if this assignment is not strategic, it can be approximated by random choice among vehicles available nearby. Customers then travel to station $j \in I$ with probability p_{ij} and exponentially distributed travel times $\frac{1}{\mu_{ij}}$. The operator's profit increases by a contribution margin of r_{ij}^{α} for every user she serves from station i to station j, given by revenue minus direct costs. Thus, contribution margins may be higher for driverless vehicles than for human-driven vehicles ($r_{ij}^d \geq r_{ij}^h \ \forall i, j \in I$). Customers who arrive at a station with no vehicles are lost, since customers often have an outside option such as bike or public transportation. To ensure a high availability throughout the vehicle sharing system, the operator can rebalance her fleet by sending a vehicle from station i to station j. Relocating a vehicle is associated with cost c_{ij}^{α} which consist of direct vehicle and labor costs. c_{ij}^{α} is higher for human-driven vehicles than for driverless vehicles ($c_{ij}^d \leq c_{ij}^h \ \forall i, j \in I$).

The decisions can be split into two different stages. The first stage comprises the fleet sizing and composition decision, the second stage comprises the operational rebalancing decisions. While both decisions can be integrated, such an integrated approach is computationally intractable even for small instances. Thus, we address the first and second stage decisions consecutively, and adopt a bound-and-enumerate algorithm in the first stage, and present two algorithms for the second stage decision. The algorithm for the first stage decision utilizes three properties that apply to both second stage algorithms. These properties will be introduced as part of the first stage decision, and proven when discussing the second stage algorithms.

5.2.1 First Stage Decision and Solution Method

Algorithm 1 presents pseudo-code for the bound-and-enumerate procedure used to determine the optimal fleet composition. Lemma 5.1 gives three conditions that guarantee the optimality of Algorithm 1.

Algorithm 1 Pseudo-Code of the Bound-and-Enumerate Algorithm

$\underline{\Pi} \leftarrow 0$
compute $m^{\max}$
for $m^d \leq m^{\max}$ **do**
  compute $\pi\left(m^d, 0\right)$
  $\underline{\Pi} \leftarrow \max \pi\left(m^d, 0\right) - F^d m^d$
end for
for $m^d + m^h \leq \max m$ **do**
  if $\pi\left(m^d + m^h, 0\right) - F^d m^d - F^h m^h \geq \underline{\Pi}$ **then**
    compute $\pi\left(m^d, m^h\right)$
    $\Pi\left(m^d, m^h\right) = \pi\left(m^d, m^h\right) - F^d m^d - F^h m^h$
  end if
end for

Lemma 5.1. *Algorithm 1 finds an optimal solution if the following properties hold.*

1. *The expected operational profit of a fleet that consists of driverless vehicles is an upper bound for any fleet of the same size* ($\pi\left(m^d, m^h\right) \leq \pi\left(m^d + m^h, 0\right)\ \forall m^d, m^h$).
2. *The expected operational profit of any fleet is bounded from below by* 0, *and from above by some constant* $\bar{\pi}$.
3. *The maximum total fleet size is bounded, i.e., there exists some* $m^{\max}$ *such that for all* m^d, m^h *s.t.* $m^d + m^h > m^{\max}$ *is not optimal* ($\Pi^\star > \Pi\left(m^d, m^h\right)$).

We refer to Appendix 5.A.1 for the proof of Lemma 5.1. $\pi\left(m^d, m^h\right)$ is not generally concave in m^d and m^h, but Algorithm 1 can still find the optimal fleet composition. To prove that the maximum fleet size is bounded (Property 3), we require Lemma 5.2 and Corollary 5.1. Appendices 5.A.2 and 5.A.3 contain the proofs.

Lemma 5.2. *Let* $\pi\left(m^d, m^h\right)$ *be known for some* $m = \{m^d, m^h\}$. *If there exists any upper bound* $\bar{\pi}$ *on the operational profit, then any* m' *such that*

$$\bar{\pi} - \pi\left(m^d, m^h\right) \leq F^d\left(m^{d\prime} - m^d\right) + F^h\left(m^{h\prime} - m^h\right), \tag{5.2}$$

cannot be optimal. If the above criterion holds for some m', *it also holds for all* $\hat{m} \geq m'$.

The above criterion can only apply if $m \leq m'$, since the left hand side of (5.2) is strictly positive. Lemma 5.2 directly implies that the number of vehicles cannot exceed an upper bound $\bar{m}^d$, $\bar{m}^h$:

Corollary 5.1. *If the expected operational profit cannot exceed some upper bound $\bar{\pi}$, then the optimal number of driverless/human-driven vehicles is bounded from above by*

$$m^d \cdot F^d + m^h \cdot F^h \leq \bar{\pi},$$

or independently per vehicle type by

$$m^d \leq \left\lfloor \frac{\bar{\pi}}{F^d} \right\rfloor; \qquad m^h \leq \left\lfloor \frac{\bar{\pi}}{F^h} \right\rfloor.$$

5.2.2 Dual Linear Programming Formulation for the Semi-Markov Decision Process

We solve the second stage SMDP for an optimal, state-dependent policy by means of linear programming using the formulation of Tijms (2003). On preliminary numerical experiments, linear programming solved slightly larger instances than value iteration or policy iteration. The decision variable $\hat{x}_{sa}$ is the long-run probability of being in state $s \in S$ taking action $a \in \mathcal{A}(s)$ weighted by the expected sojourn time of this state/action pair, $\pi_s(a)$ is the expected contribution margin in state s taking action a (expected revenue minus rebalancing costs), $p_{st}(a)$ is the transition probability from state s to state t when taking action a, and τ_{sa} is the average time in state s taking action a.

$$\max \pi\left(m^d, m^h\right) = \sum_{s \in S} \sum_{a \in \mathcal{A}(s)} \pi_s(a)\, \hat{x}_{sa} \tag{5.3a}$$

$$\sum_{a \in \mathcal{A}(t)} \hat{x}_{ta} - \sum_{s \in S} \sum_{a \in \mathcal{A}(s)} p_{st}(a)\, \hat{x}_{sa} = 0 \qquad \forall t \in S \tag{5.3b}$$

$$\sum_{s \in S} \sum_{a \in \mathcal{A}(s)} \tau_{sa} \hat{x}_{sa} = 1 \tag{5.3c}$$

$$\hat{x}_{sa} \in \mathbb{R}_0^+ \qquad \forall s \in S, a \in \mathcal{A}(s) \tag{5.3d}$$

The objective function (5.3a) maximizes the expected operational profit per period of time, given by the probability of entering a state s, taking action a, multiplied by the contribution margin. Constraints (5.3b) are flow-balancing constraints. Constraints (5.3c) establish that the system must be in exactly one state at any point in time. Constraints (5.3d) define the variable domain.

States

We define a state by the number of driverless and human-driven vehicles at each station i, y_i^d and y_i^h, and the number of vehicles in transit between any two stations i and j, y_{ij}^d and y_{ij}^h:

$$s = \langle y_1^\alpha, \dots y_n^\alpha, y_{11}^\alpha, \dots y_{nn}^\alpha \mid \forall i, j \in I, \alpha \in \{d, h\} \rangle, s \in S.$$

A state is feasible if the number of vehicles at stations and in transit of any type equals the total number of vehicles of this type. $y_i^\alpha(s)$ and $y_{ij}^\alpha(s)$ refer to station-inventories and in-transit-inventories of state s, respectively.

Actions and Sojourn Times

The operator chooses an action after each event (customer arrival or vehicle return), i.e., in every decision epoch. Since the state does not change between two events, rebalancing at other points in time will not improve the decision. In every decision epoch, the operator rebalances at most one vehicle, and preferably a driverless vehicle. As Benjaafar et al. (2018) prove, operators always return to the border of some "no-rebalancing zone", i.e. a steady state. Whenever the system leaves the steady state, rebalancing one vehicle is sufficient, since the operator can react after each event. Action $a \in \mathcal{A}(s)$ defines the rebalancing operation: $a = \langle i, j \rangle$ for moving a vehicle from i to j, or $a = \emptyset$ if the operator chooses to not relocate any vehicles. We define an auxiliary state s' (post-decision state) which does not trigger an event, but keeps track of all planned changes due to action a. The time the system remains in state s' only depends on the inventories in this state, and is calculated as the merged Poisson process over all customer arrival rates and vehicle return rates

$$\tau_{sa} = \tau_{s'} = \left(\sum_{i \in I} \lambda_i + \sum_{ij \in I} \mu_{ij} \sum_{\alpha \in \{d,h\}} y_{ij}^\alpha(s') \right)^{-1}.$$

State Transitions

State transitions from a post-decision state s' to the next state t can be triggered either by the arrival of a customer or the return of a vehicle. If the event is a customer arrival at a non-empty station i, the inventory at this station i decreases by 1 either for driverless or human-driven vehicles with probability $\frac{y_i^\alpha}{y_i^d + y_i^h}$, and one of the inventories in transit y_{ij}^d

(y_{ij}^h) increases by 1 (with probability p_{ij}). The probability of this state transition is

$$p_{s't}(\emptyset) = \frac{y_i^\alpha(s')}{y_i^d(s') + y_i^h(s')} \cdot \lambda_i p_{ij} \cdot \tau_{s'}.$$

If the event is a driverless (human-driven) vehicle return at station j from station i, the inventory at this station j increases by 1 driverless (human-driven) vehicle and the inventory in transit y_{ij}^α decreases by 1. The probability of this state transition is

$$p_{s't}(\emptyset) = y_{ij}^\alpha(s')\,\mu_{ij} \cdot \tau_{s'}.$$

If the event is a customer arrival at an empty station i ($y_i^d + y_i^h = 0$), the state does not change and the customer leaves the system. The probability of remaining in the previous state is

$$p_{s't}(\emptyset) = \sum_{i \in I} \lambda_i \cdot \mathbb{1}_{y_i^d + y_i^h = 0} \cdot \tau_{s'}$$

where $\mathbb{1}$ is the indicator function.

Profit of State-Action Pairs

The profit of being in state s taking action a depends on two main components: relocation costs c_{ij}^α and expected contribution margins r_i^α. The contribution margin depends on the vehicle choice which is random to represent customers choosing the closest available vehicle. The former only apply if the action is to relocate some vehicles ($a \neq \emptyset$).

$$\pi_s(a) = \tau_{sa} \sum_{i \in I | y_i^d + y_i^h \neq 0} \sum_{\alpha \in \{d,h\}} \frac{y_i^\alpha}{y_i^d + y_i^h} r_i^\alpha \cdot \lambda_i - \sum_{i,j \in I} \mathbb{1}_{a = \langle i,j \rangle} \left(c_{ij}^d + \mathbb{1}_{y_i^d = 0} \left(c_{ij}^h - c_{ij}^d \right) \right) \quad (5.4)$$

Properties of the SMDP

We show that the expected operational profit derived using SMDP fulfills the three properties introduced in Lemma 5.1. Theorem 5.1 implies Property 1, as well as the lower bound for Property 2. If there exists an upper bound, Theorem 5.1 also implies Property 3.

Theorem 5.1. *The optimal expected operational profit function $\pi\left(m^d, m^h\right)$ of the second stage problem is (i) quasi-concave, (ii) increasing, and (iii) non-negative in the number*

of driverless and human-driven vehicles. (iv) an additional driverless vehicle increases expected operational profits more than an additional human-driven vehicle:

$$\pi\left(m^d - x, m^h + x\right) \leq \pi\left(m^d, m^h\right) \qquad \forall x \in \mathbb{N}\ \forall i, j.$$

The proof can be found in Appendix 5.A.4.

While the Dual Linear Program (DLP) provably results in optimal solutions, it is computationally infeasible even for small instances. The state space of the SMDP grows exponentially (due to the second Stars and Bars Theorem, Feller (1950, pg. 37)), e.g., for 2 locations and 20 vehicles per type, the number of state-action pairs already exceeds $5 \cdot 10^7$.

5.2.3 Fluid-Based Approximation

To solve the rebalancing problem on larger instances, we adapt a Fluid-Based Approximation Linear Program (FLP) from Braverman et al. (2019). Similar to Zhang et al. (2018), the system consists of two coupled closed queueing networks. At every station i, newly arriving customers are assigned to either the driverless or human-driven queueing network with a variable probability of p_i^α, and continue to another station j with probability p_{ij}. Each system consists of n single-server nodes representing stations, n^2 infinite-server nodes representing vehicles in which a customer is chauffeured from station i to station j, and $n \cdot (n-1)$ infinite-server nodes representing vehicles that currently rebalance from i to j ($i \neq j$). The average number of vehicles at each node (i.e., the queue length) is given by e_{ii}^α for the single-server nodes ("empty vehicles"), and f_{ij}^α ("full vehicles") and e_{ij}^α for the infinite-server nodes. If the number of vehicles is too low, the operator cannot satisfy all customers, even if those customers were willing to wait. Then, the linear programming formulation is as follows:

$$\begin{aligned}
\max \bar{\pi}\left(m^d, m^h\right) = &\sum_{i \in I, \alpha \in \{d,h\}} r_i^\alpha \lambda_i p_i^\alpha \\
&- \sum_{i,j \in I, \alpha \in \{d,h\}} c_{ij}^\alpha \mu_{ij} e_{ij}^\alpha && \text{(5.5a)} \\
\lambda_i p_{ij} p_i^\alpha = &\mu_{ij} f_{ij}^\alpha && \forall i, j \in I, \alpha \in \{d,h\} \quad \text{(5.5b)} \\
\mu_{ij} e_{ij}^\alpha \leq &\sum_{k \in I} \mu_{ki} f_{ki}^\alpha && \forall i, j \in I, \alpha \in \{d,h\} \quad \text{(5.5c)} \\
\sum_{k \in I, k \neq i} \mu_{ki} e_{ki}^\alpha \leq &\lambda_i p_i^\alpha && \forall i \in I, \alpha \in \{d,h\} \quad \text{(5.5d)}
\end{aligned}$$

$$\lambda_i p_i^\alpha + \sum_{j\in I, j\neq i} \mu_{ij} e_{ij}^\alpha = \sum_{k\in\{d,h\}, k\neq i} \mu_{ki} e_{ki}^\alpha + \sum_{k\in I} \mu_{ki} f_{ki}^\alpha \qquad \forall i \in I, \alpha \in \{d,h\} \qquad (5.5e)$$

$$p_i^d + p_i^h \leq 1 \qquad \forall i \in I \qquad (5.5f)$$

$$\sum_{i,j\in I} e_{ij}^\alpha + f_{ij}^\alpha = m^\alpha \qquad \forall \alpha \in \{d,h\} \qquad (5.5g)$$

$$0 \leq e_{ij}^\alpha, f_{ij}^\alpha \leq m^\alpha \qquad \forall i,j \in I, \alpha \in \{d,h\} \qquad (5.5h)$$

$$0 \leq p_i^\alpha \leq 1 \qquad \forall i \in I, \alpha \in \{d,h\} \qquad (5.5i)$$

The objective is to maximize expected operational profits, given by contribution margins for served customers and rebalancing costs for every vehicle that enters a rebalancing queue. Constraints (5.5b)–(5.5e) follow from Braverman et al. (2019) for each type of vehicle independently. Constraints (5.5b) makes sure that (i) a fraction p_{ij} of all customers who are assigned to vehicle type α at location i (in total $\lambda_i p_i^\alpha$) travel to location j, and thus increase the expected number of vehicles in the infinite-server node from i to j by $\mu_{ij} f_{ij}^\alpha$ (since they remain in this queue for μ_{ij} periods on average). Constraints (5.5c) and (5.5d) guarantee that the number of empty vehicles that leave station i in any period does not exceed the number of full vehicles arriving there; and that the number of empty vehicles that arrive at some station i does not exceed the total demand at this station. (5.5e) are flow balancing equalities for station i and vehicle type α. Constraints (5.5f) link both closed queueing networks. Constraints (5.5g) ensure that the number of vehicles per type equals the fleet size. It shall be noted that in the second stage problem, m^d and m^h cannot be endogenized, since $p_i^d + p_i^h$ is only an upper bound on the availability, and the actual availability is non-linear.

Corollary 5.2. *The solution to the linear program in* (5.5) *yields an upper bound on the operational profit* $\bar{\pi}$*, if* $m^h = 0$ *and* m^d *is sufficiently large such that there exists a feasible solution in which all demand can be satisfied and vehicles can be rebalanced on the cheapest (rather than fastest) rebalancing arcs.*

We refer to Appendix 5.A.5 for the proof.

The fluid-based approximation control is then to rebalance one vehicle with probability q_{ij}^α to some location j if the event is a vehicle return of type α at location i with

$$q_{ij}^\alpha = \frac{\mu_{ij} e_{ij}^\alpha}{\sum_{k\in I} \mu_{ki} f_{ki}^\alpha + \sum_{k\in I, k\neq i} \mu_{ki} e_{ki}^\alpha} \qquad \forall i,j \in I, j \neq i,$$

and no rebalancing otherwise. Rebalancing after rentals does not improve a state-

independent control.

Actual Availability

Since a_i is only a theoretical upper bound on the availability, we must calculate the actual availability $A_i^\alpha(m^\alpha)$. Analogous to Zhang and Pavone (2016), we use Mean Value Analysis (MVA) (Reiser and Lavenberg, 1980) to compute average waiting times $W_i^\alpha(m')$ and queue lengths $L_i^\alpha(m')$ for $1 \leq m' \leq \sum_{i,j} e_{ij}^\alpha + f_{ij}^\alpha$. The underlying principle behind MVA is to add one vehicle at a time and study how the expected number of vehicles per station – and consequentially the average waiting time of a vehicle at a station – increases. The expected time a vehicle spends in an infinite-server node (between two stations) is independent from the number of vehicles in the system, and given by $\frac{1}{\mu_{ij}}$ (the travel time). The expected time a vehicle remains at a station before being rented by a customer is given by the current queue length (assuming that the additional vehicle is placed at station i) multiplied by the inter-arrival time at this station ($\frac{1}{\tilde{\lambda}_i^\alpha}$). The mean length of a queue at a station is given by waiting time weighted by throughput, scaled to the size of the system and weighted waiting times in other queues.

$$
\begin{aligned}
W_{ij}(m') &= \frac{1}{\mu_{ij}(1)} && \forall i,j \in I \\
W_i(m') &= \frac{1}{\tilde{\lambda}_i^\alpha}\left(1 + L_i(m'-1)\right) && \forall i \in I \\
L_i(m') &= \frac{m'\pi_i W_i(m')}{\sum_j \pi_j W_j(m') + \sum_{jk} \pi_{jk} W_{jk}(m')} && \forall i \in I
\end{aligned}
$$

with $L_i(0) = 0$, and $\tilde{\lambda}_i^\alpha = \lambda_i \cdot p_i^\alpha + \sum_j \frac{e_{ij}^\alpha}{\mu_{ij}}$ being the adapted arrival rate accounting for the customers delegated to system α as well as the rate at which vehicles are rebalanced (Zhang and Pavone, 2016). The relative throughput of station i, π_i, solves the set of equations

$$
\pi_i = \sum_{k \in I} \pi_k \tilde{p}_{ki}^\alpha \qquad \forall i \in I
$$

and defines the relative throughput of the infinite-server nodes

$$
\pi_{ij} = \pi_j \tilde{p}_{ij}^\alpha \qquad \forall i,j \in I
$$

where $\tilde{p}_{ij}^\alpha = p_{ij} q_{ii}^\alpha + q_{ij}^\alpha$ are adapted transition probabilities.

The availability is

$$A_i^\alpha\left(m^\alpha\right) = \frac{L_i\left(m^\alpha\right)}{W_i\left(m^\alpha\right)\tilde{\lambda}_i^\alpha} \qquad \forall i \in I. \tag{5.6}$$

Operational Profit Function

The expected operational profit is

$$\pi\left(m^d, m^h\right) = \sum_{\alpha\in\{d,h\}, i\in I}\left(r_i^\alpha \lambda_i p_i^\alpha - \sum_{j\in I} c_{ij}^\alpha \mu_{ij} e_{ij}^\alpha\right)\cdot A_i^\alpha\left(m^\alpha\right). \tag{5.7}$$

Both payoffs and rebalancing costs only occur if a vehicle is available that the customer uses, or if a customer returns a vehicle and this vehicle is rebalanced.

The optimal solution to the linear program (5.5) can be degenerate, as the variables p_i^α, f_{ij}^α do not directly affect the objective function. The two symmetry-breaking constraints

$$\sum_{i\in I} p_i^\alpha \frac{\lambda_i}{\sum_j \lambda_j} = \sum_{i\in I}\frac{m^\alpha}{m^d + m^h}\cdot a_i \qquad \forall \alpha \in \{d, h\}, \tag{5.8}$$

ensure that a control with a "more even" usage of vehicle types is chosen. The idea behind this constraint is that in the balanced case, the operator assigns customers to the two different systems proportional to the fleet size per system.

Properties

To show that the expected operational profit of an entirely driverless fleet is an upper bound on the expected operational profit of any fleet of the same size, we show that the objective (5.5a) and the availability A^α are at least as high with only driverless vehicles as with any mixed fleet.

Lemma 5.3. *The availability of mixed vehicle fleets m^d, m^h is never higher than the availability of an driverless fleet of the same total size $(m^d + m^h, 0)$ if costs and contribution margins are proportional to the expected travel time.*

The proof of Lemma 5.3 can be found in Appendix 5.A.6.

Theorem 5.2 directly implies Property 1, as well as the lower bound for Property 2; the upper bound follows from Corollary 5.2.

Theorem 5.2. *The expected operational profit function $\pi\left(m^d, m^h\right)$ of the second stage problem derived via FLP is (i) concave if $m^d = 0$ or $m^h = 0$ (for $m^d + m^h \geq \underline{m}$), (ii) increasing in m^h (m^d) if $m^d = 0$ ($m^h = 0$) and (iii) non-negative in the number of driverless and human-driven vehicles. Also, (iv) operators can gain the highest expected operational profit if a fleet of constant size consists of only driverless vehicles*

$$\pi\left(m^d + m^h, 0\right) \geq \pi\left(m^d, m^h\right) \qquad \forall m^d, m^h \in \mathbb{N},\ \text{if } c_{ij}^d \leq c_{ij}^h \forall i, j \in I$$

if costs and contribution margins are proportional to the expected travel time.

We note that if both $m^d, m^h > 0$, the expected operational profit function for the FLP is neither concave nor provably increasing.

5.3 Experimental Design

We study three different sets of instances and case studies: (i) we consider small instances which allow to investigate effects using SMDP as a second stage algorithm; (ii) we generate artificial instances which resemble realistic environments for vehicle sharing systems; and (iii) we present results for two different real-world case studies settled in China and New York City. The artificial instances and case studies complement each other, as the case studies allow to quantify the benefit of driverless and mixed fleets, while the artificial instances can explain under which circumstances driverless vehicles are necessary to operate fleets profitably. Since the case studies resemble existing carsharing and ride-hailing services with currently only human-driven vehicles, we do not expect to see this effect there. Furthermore, the impact of demand imbalances becomes clearer on artificial instances. We define the demand imbalance as the difference between arrivals and departures per station:

$$\text{imbalance} = \frac{\sum_{i \in I} \left| \sum_{j \in I} \left(\lambda_i p_{ij} - \lambda_j p_{ji} \right) \right|}{2 \cdot \sum_{i,j \in I} \lambda_i p_{ij}}$$

5.3.1 Small Instances

The number of locations is $n = 2$, the arrival rate at station 2 is 1. All vehicles drive to the other location, and the expected travel times between the stations are 1. Rebalancing costs with driverless vehicles are $c_{12}^d = c_{21}^d = 1$, and contribution margins are 10. The design differs in the rebalancing costs for human-driven vehicles $(2, 3, 5)$, the

annuities for human-driven vehicles $(0.5, 0.625, 0.75)$, the annuities for driverless vehicles (relative to F^h, $\frac{F^d}{F^h} \in \{1.1, 1.2, 1.3\}$), and the arrival rate at the first station ($\lambda_1 = 1, 2$). Consequentially, the imbalance is either 0 or 1/3. c_{ij}^{α} and F^{α} are realistic \$ values if customers arrive in expectancy every 6 minutes. The contribution margins reflect a very high value for customer retention, and permit a positive fleet size. In total, we study 54 different instances.

5.3.2 Artificial Instances.

We randomly sample $n \in \{5, 10, 20\}$ stations in a circle, with uniform distributions for distances from the center ($\in \{0, 10\}$ kilometers) and angle. Thus, the density of stations decreases towards the outer boroughs. The expected travel time between any two stations i and j ($1/\mu_{ij}$) is derived by dividing the Euclidean distance between i and j by an average velocity of 25km/h. Arrival rates (average number of customers per hour) are either similar across the entire operating area, or higher closer to the center (radius ≤ 5) and lower outside. In the case of similar arrival rates, they are sampled from either $\mathcal{U}[8, 12]$, case HS, or $\mathcal{U}[4, 6]$, case LS. In the latter case, we sample arrival rates from either $\mathcal{U}[6, 9]/\mathcal{U}[11, 14]$, case HD, or $\mathcal{U}[3, 4.5]/\mathcal{U}[5.5, 7]$, case LD. Decreasing demand towards the outskirts of the city is realistic in many major cities. To represent the significant differences in adoption between different shared mobility services, we additionally multiply arrival rates by a factor of 10 in all settings to generate instances (resulting in instance sets "base" and "x10"). Customers either travel to any other station with equal probability (case E), or the probability of traveling to another station is directly proportional to the distance (case P), reflecting the higher likelihood of motorized travel on longer distances. Relocation costs are 0.3\$/km (7.5\$/h) and 0.7\$/km (17.5\$/h) for driverless and human-driven vehicles, respectively, given by 0.3\$/km for fuel and usage, 25km/h / 10\$/h = 0.4\$/km for the driver. The driver costs approx. reflect minimum wages in Europe and North America, and the costs for fuel and usage are approx. in line with Lanzetti et al. (2020). The contribution margin per served customer kilometer is either 0.6\$ (cases LS, LD) or 1.5\$ (cases HS, HD), depending on whether customer retention is considered. These contribution margins realistically mimic the contribution margins of ShareNow with roughly 0.35-0.4\$ per minute and 0.3\$/km direct costs. The contribution margin per served customer kilometer in a human-driven vehicle decreases by 0.4\$ if customers do not drive themselves (cases HD, LD). The annuity per vehicle per hour in \$ is $\{F^d, F^h\} \in \{\{1.15, 1\}, \{1.3, 1\}, \{4.6, 4\}, \{5.2, 4\}\}$ (cases LL, LH, HL, HH). The rationale behind this choice is that the total cost of

ownership for human-driven vehicles is approx. 35,000$ (e.g., a BMW i3), and vehicles are deprecated over a period of 1-4 years; driverless vehicles are more expensive, but for bulk buyers or original equipment manufacturers, a relative price difference of 1.15-1.3 is achievable.

This results in 768 instances distinguished by number of vehicles n, arrival rates, transit probabilities, contribution margins per served customer kilometer, and annuity. For example, the base instance $5\_HS\_E\_LD\_LL$ is characterized by a 5 station system at which customer arrival rates are sampled from the distribution $\mathcal{U}[8, 12]$ at all stations, customers travel to all other stations with equal probability, only direct revenues are considered and human-driven vehicles are driven by employees of the the operator, and the annuity is low for both human-driven and driverless vehicles.

Table 5.1 lists the imbalance for all instances. The imbalance of the instances varies

	5 stations		10 stations		20 stations	
Arr. Type	E	P	E	P	E	P
HS	8.83	16.15	6.56	6.68	6.32	10.81
LS	7.25	10.66	3.59	9.12	5.67	12.57
HD	13.23	23.81	13.50	31.14	10.57	28.31
LD	12.98	18.47	13.50	25.64	12.01	20.09

Table 5.1: Imbalance (in %) of the Artificial Case Study Instances for Different Sizes (Header), Arrival Types, and Transition Types

between 3.6% and 31.1%, where the P instances are more imbalanced than the E instances since the expected number of incoming vehicles increases with the distance to the center in the P instances. The HD and LD instances are more imbalanced than HS and LS instances. Size of the instances and average arrival rate have a smaller impact.

5.3.3 Case Studies based on Real-Life Vehicle Sharing Systems

The case studies represent two different vehicle sharing systems from different continents.

DiDi Challenge, China

We use the dataset of Braverman et al. (2019) which reflects customer demand of the Chinese ride-hailing provider DiDi Chuxing. The dataset consists of average trip counts during the evening rush hour (5-6PM), divided into 10 minute intervals. During this

interval, taxi movements are almost balanced (imbalance 5.9%). The average number of trips per interval is 176, and those trips take on average 13.2 minutes. The trip fare is 1.41$ per 10 minute time interval, the costs for the driver amount to 0.82$ per time interval, and the costs for fuel and wear amount to 0.24$ per time interval, assuming an average velocity of 20km/h. We consider two settings for contribution margins of human-driven vehicles $r^h_{ij} = 0.35\frac{1}{\mu_{ij}}$, and $r^h_{ij} = 1.17\frac{1}{\mu_{ij}}$ which refer to customers being chauffeured ("ride-hailing") or driving themselves ("carsharing"), respectively. The annuities amount to 0.205$ for human-driven and 0.236$ or 0.256$ for driverless vehicles.

Lower Manhattan, NY, USA

The New York City Taxi and Limousine Commission publishes the booking data for all taxis (yellow cabs) in New York City. We use data for the end of the morning rush hour (9-10AM) on weekdays in January 2015 to generate average movement patterns within Lower Manhattan. Lower Manhattan consists of into 19 quadratic tiles with 500m edge length, and the total number of trips is counted for every pair of locations. On average, 1467 trips occurred within one hour. The average trip is 2km long, takes 8.34 minutes, and results in a fare of 7.52$. Taxi demand in Lower Manhattan faces substantial imbalances of 14.8%. The costs of a trip or rebalancing operation comprise costs for the vehicle, 0.36$ per km, and potentially driver, 42$ per hour. This cost structure is due to Lanzetti et al. (2020). Payoffs are calculated by subtracting the costs from the fare. Annuities are $F^d = 1.15$ or $F^d = 1.3$ and $F^h = 1$.

5.4 Computational Study

We study the value of driverless vehicles and the benefit of mixed fleets in vehicle sharing systems on the artificial instances and real-life case studies introduced in Section 5.3. Since service levels are an relevant criterion for users of the system, and since service levels allow to explain varying profits, we also report these.

We conduct numerical experiments on a personal computer with 16GB of memory on a single core. The bound-and-enumerate algorithm is implemented in Java, utilizing IBM ILOG CPLEX v12.10 for solving the linear programs of the second stage. For most experiments, we restrict ourselves to using FLP as a second stage algorithm, since the SMDPs can only solve small instances.

5.4.1 Fleet Sizing and Composition on Small Instances

In the following, we study service levels and the value of driverless vehicles on instances with two locations.

	SMDP						FLP					
	Service Levels			Profits			Service Levels			Profits		
Setting	SL	Δ SL h	Δ SL d	profit	benefit d	benefit h	SL	Δ SL h	Δ SL d	profit	benefit d	benefit h
Rebalancing Costs												
2	97.27	-0.11	0.18	17.25	0.79	2.05	83.49	-0.22	1.65	15.1	0.36	3.5
3	97.21	0.68	0.14	17.17	3	1.7	83.07	-0.22	1.47	14.92	1.88	3.22
5	97.14	4.34	0.1	17.16	9.43	1.68	82.9	-0.4	1.41	14.91	7.28	3.49
Annuities h												
0.5	97.72	1.7	0.01	18.32	4.89	1.32	84.54	0	0	16.06	3.63	2.76
0.625	97.07	1.52	-0.06	17.17	4.53	1.8	83.65	-0.18	2	14.94	3.65	3.25
0.75	96.84	1.69	0.47	16.09	4.35	2.31	81.27	-0.72	1.98	13.93	3.51	4.17
Annuities d												
1.1	97.52	1.86	0.11	17.44	5.37	0.69	83.24	-0.15	1.02	15.19	4.37	1.69
1.2	97.1	1.56	0.09	17.17	4.27	1.7	83.03	-0.48	1.24	14.96	3.53	3.24
1.3	97.01	1.49	0.21	16.97	4.05	3.03	83.2	-0.28	2.04	14.79	2.6	4.9
Imbalance												
1	96.74	1.88	0.28	13.49	3.98	3.53	82.27	-	4.59	12.33	-	6.97
2	97.68	3.03	0.13	20.9	9.92	1.9	84.05	3.03	0.13	17.63	8.71	2.86

Table 5.2: Average Service Level, Increase of Service Level (% increase over only human-driven/driverless if both exist), Average Profits (abs. values), Value of Driverless Vehicles (% increase over only human-driven vehicles), and Value of Keeping Human-Driven Vehicles (% increase over only driverless vehicles) in the Small Instances

Impact of Driverless Vehicles on the Service Level

The left-most and third block of Table 5.2 list service levels and the increase of service level relative to fleets consisting of only human-driven or only driverless vehicles. Due to the high contribution margins, average service levels are very high at 97.2%. This is 1.6% higher compared to only human-driven vehicles and 0.14% higher than for fleets only consisting of driverless vehicles. Using FLP as a second stage algorithm, the service levels are lower with 83.2% on average, but the trends and influencing factors are identical for both second stage algorithms. The main drivers for high services levels are low annuities and high imbalances. The former is unsurprising since lower annuities permit procuring larger fleets. The latter seems surprising, and can be explained by increasing fleet sizes to counteract the rebalancing demand as well. The difference in service levels of fleets only consisting of human-driven vehicles (driverless vehicles) and mixed fleets increases (decreases) (i) if the rebalancing cost of human-driven vehicles increases, (ii) if annuities of driverless vehicles decreases, and (iii) if the imbalance increases.

Impact of Driverless Vehicles on Profits

The second-left and right-most block of Table 5.2 report profits, and percentage increases of the profit relative to fleets consisting of only human-driven and only driverless vehicles. On average over all instances, the benefit of driverless vehicles w.r.t. profits is 4.6%, and the benefit of keeping human-driven vehicles, i.e., the profit difference between mixed and driverless fleets, is 1.8%. Using FLP as a second stage algorithm, the benefit of driverless vehicles is slightly lower (3.6%), and the benefit of human-driven vehicles is substantially higher (3.4%). The former occurs since with FLP, the operator does not rebalance her fleet if the imbalance is 0. As the imbalance is 0 in half of all instances, these instances do not benefit from driverless vehicles. On these very small instances, the gap between the profits of the two algorithms is substantial, but the profits increase faster for FLP than for SMDP as annuities F^h increase. Thus, the approximation error of FLP decreases as the fleet size increases which is the case in more realistic instances. Profits increase if the imbalance increases. The latter is an artefact of increasing customer demand (50% more demand), and does not persist when computing the profit per customer. Surprisingly, the profit only decreases very little if rebalancing costs of human-driven vehicles increase. This is because a the operator can smoothly increase the fraction of driverless vehicles, and only rebalance with those. Unsurprisingly, the benefit of driverless vehicles (human-driven vehicles) is driven by increasing (decreasing) contribution margins of human-

driven vehicles, decreasing (increasing) differences in annuities between vehicle types, and an increasing (decreasing) imbalance.

5.4.2 Fleet Sizing and Composition on Artificial Instances

In the following, we study service levels, fleet size and composition, and value of driverless vehicles on artificial instances.

Impact on Service Availability and Fleet Structure

For instances with equal contribution margins, Table 5.3 lists the average service levels (over those instances in which $m^d + m^h > 0$), the percentage increase in the service level over a fleet only consisting of human-driven/driverless vehicles, and the number of instances in which the service can be offered (with a mixed fleet, a human-driven, and an driverless fleet), separately for a base system and a system with 10-fold arrival rate (x10 system). Omitting different contribution margins from the analysis allows to interpret the results without the superimposing effect that different contribution margins always result in the optimal choice of only driverless vehicles. We further report the influence of different contribution margins in the last block of Table 5.3. When introducing driverless vehicles, it is possible that the vehicle sharing system can be operated even if it was unprofitable with only human-driven vehicles, resulting in an increase in the number of instances in which the service is offered with human-driven vehicles (column "# h") to mixed fleets (column "# total"). Further, the operators may reach a higher average service level, suggesting a better usage of the invested capacity. We find that vehicle sharing is not offered ($m^d + m^h = 0$) in many base instances (271 out of 384) which is mainly since current customer payoffs are insufficient for profitable systems. A system with a larger customer base (x10 instances) is more profitable, but there is still a large number of instances in which service is not offered (223 out of 384). Table 5.3 reports the number of instances which have a non-zero fleet size if the fleet is mixed, human-driven, or driverless. Unprofitability is mainly driven by the instance size (only 5 instances with $n = 20$ stations can be operated profitably in the base setting), high annuities for human-driven vehicles, and low contribution margins for driverless vehicles. We observe that a higher average arrival rate per location increases the service level. Since low contribution margins in the numerical design are direct revenues minus direct operating costs, we conclude that without subsidies (e.g., by the city or parent company) or a high valuation for customer retention (e.g., anticipation of a growing

	Base							x10						
Setting	SL	Δ SL h	Δ SL d	# total	# mixed	# h	# d	SL	Δ SL h	Δ SL d	# total	# mixed	# h	# d
Size														
5	48.96	9.52	-22.21	42	6	42	34	67.36	16.07	-17.03	48	20	48	47
10	40.49	2.75	-13.34	16	1	16	16	55.12	7.58	-9.09	21	6	20	19
20	22.04	1.07	-0.88	4	0	4	1	43.92	2.2	-6.43	16	0	16	10
Arrival Type														
HS	49.63	3.74	-13.31	18	2	18	15	68.11	8.05	-9.13	20	4	20	19
LS	53.91	0.85	-11.39	14	4	14	13	68.55	2.35	-9.18	22	8	22	19
HD	37.54	10.22	-10.65	16	0	16	11	53.07	16.75	-11.54	21	7	20	19
LD	38.82	2.97	-13.22	14	1	14	12	50.4	7.32	-13.55	22	7	22	19
Transition Type														
E	47.62	0.92	-10.94	26	4	26	23	61.16	4.17	-8.84	40	14	40	33
P	43.17	7.98	-13.35	36	3	36	28	58.83	13.06	-12.85	45	12	44	43
Annuities														
LL	49.62	12.12	-22.7	26	3	26	25	61.51	17.28	-18.87	35	10	34	33
LH	46.27	3.88	-21.99	26	4	26	23	61.13	14.59	-14.57	34	16	34	28
HL	29.93	0.89	-2.59	5	0	5	2	55.41	1.73	-5.42	8	0	8	8
HH	29.93	0.89	-1.29	5	0	5	1	52.42	0.87	-4.54	8	0	8	7
Contribution Margin														
LS	34.53	4.67	-8.75	16	2	16	15	51.34	9.13	-6.32	21	13	20	19
HS	48.69	4.23	-15.54	46	5	46	36	62.74	8.1	-15.37	64	13	64	57
LD	38.21	0	-8.22	15	0	0	15	52.89	0	-6.51	19	0	0	19
HD	56.94	16.45	-14.52	36	0	30	36	66.03	74.55	-14.46	57	0	52	57

Table 5.3: Average Service Level (if available), Increase of Service Level (% increase over only human-driven/driverless if both exist), and Numbers of Offered Fleets (total, mixed, human-driven and driverless) in the Artificial Instances

market), carsharing and ride-hailing services can hardly ever operate. Surprisingly, a 10-fold arrival rate does not increase the number of instances in which the service is offered

substantially if contribution margins are low. This indicates that even larger customer bases cannot alleviate the effect of low contribution margins, and can be explained by means of the upper bound $\bar{\pi}$ that reflects the maximum attainable operational profit: Rebalancing costs, contribution margins and fleet sizes due to the $\bar{\pi}$ scale linearily in the multiplicative factor on λ_i, and if $\bar{\pi}$ is already very low, profitability with non-zero fleet sizes is hardly ever reached when considering stochasticity. For high contribution margins, the increased availability by a fleet at scale allows to offer the fleet profitably in additional instances.

In the 21 out of 384 base instances and 25 x10 instances, the service can only be offered if driverless vehicles are available. In these instances, the fleet consists of driverless vehicles only. Driverless vehicles are necessary for offering the service if contribution margins differ between vehicle types, or the system has a high imbalance. Operators may benefit from mixed fleets if contribution margins are equal for both vehicle types, and if annuities are low. If contribution margins differ, using only driverless vehicle is usually a dominant strategy. If annuities are high, operators resort to a smaller fleet which more frequently consists of only a single vehicle type. Driverless vehicles can ensure operations in particular if travel flows are imbalanced (given by different arrival rates and transit probabilities proportional to the distance), if contribution margins for human-driven vehicles are lower than for driverless vehicles, and if annuities are high. Figure 5.1 lists the fleet composition of the x10 instances for different sets of instances, aggregated by number of instances, same or different arrival rates, and transition type. Though affected by some random differences, the trend is clear: With an increasing imbalance, driverless and mixed fleets become more frequent.

If driverless vehicles are permitted, the service level increases significantly (6.3% for the base case and 22.9% in the x10 instances), and decreases in 3 out of 384 base instances. This is due to inefficiencies in allocating vehicles (lost pooling benefits). Vice versa, when comparing the service level between mixed and driverless fleets, the service level decreases significantly (11.8% decrease for the base case and 10.7% decrease for the x10 instances). This is since rebalancing activities are drastically reduced in human-driven or mixed fleets, and customers are instead being rejected at stations with more rental requests than returns. This can be alleviated by forcing equal service levels at all stations, such as suggested by Zhang et al. (2018). Service levels are low in almost all base instances (47.9% on average if service is offered). With a 10-fold arrival rate, average service levels are higher (61.3%), but still lower than target service levels suggested in literature (e.g., 80% in He et al. (2017), or 85-95% in Zhang et al. (2018)). This gap

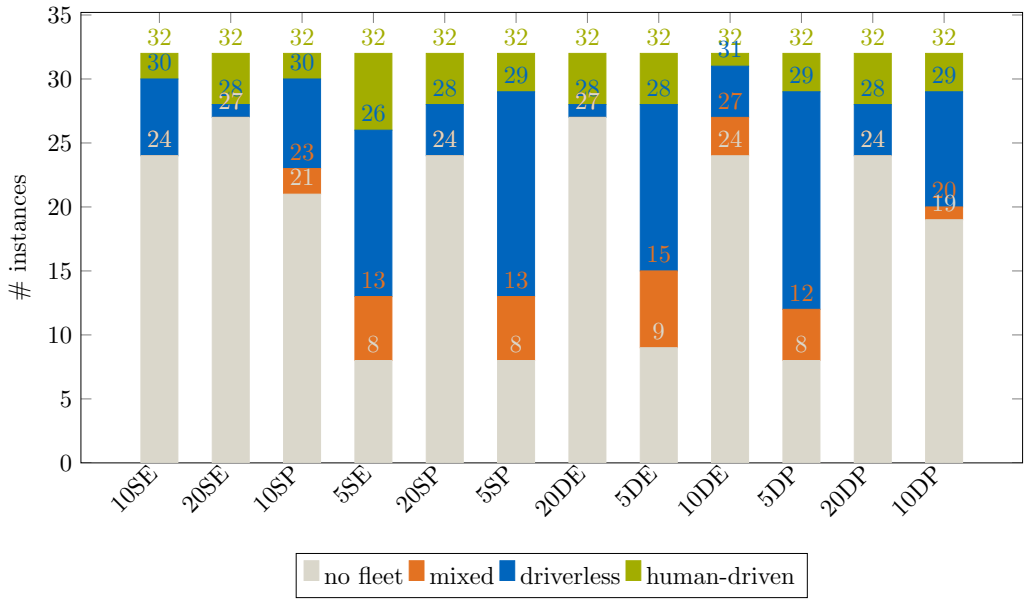

Figure 5.1: Fleet Composition (No fleet, mixed fleet, driverless fleet, human-driven fleet) for x10 Instances Grouped by No. of Stations, Same/Different Arrival Rate and Transition Type, Ordered by Imbalance

indicates that (i) optimizing profits for a given service level as suggested in literature is suboptimal for operators, but (ii) introducing service levels in the optimization may result in an unreliable service. The increasing average service level is due to improved pooling: The number of vehicles necessary to serve a customer trip decreases, allowing operators to serve more customers with a comparable fleet size. Service levels are mainly driven by an increasing arrival rate or a decreasing number of locations.

Impact on Profits and the Benefit of Driverless Vehicles

Table 5.4 lists the average profits of the vehicle sharing service, as well as the benefit of driverless vehicles and the benefit of keeping human-driven vehicles in the fleet, for the base case and the x10 instances. Obviously, the benefit of driverless (human-driven) vehicles is only available if the service is also offered if only human-driven (driverless) vehicles were available. For comparability to the service levels, Table 5.4 only lists those instances where $r^d = r^h$. The benefits of driverless and human-driven vehicles are represented in percentage values. For the base instances, the average profits over all instances are low, the hourly return is 7.59$. This very low value is partially due to the low number of instances in which the service is offered. If the service is operated, the

	Base			x10		
Setting	profit	benefit d	benefit h	profit	benefit d	benefit h
Size						
5	27.28	5.6	25.76	434.26	7.3	23.99
10	22.38	0.36	22.25	278.6	5.35	8.33
20	7.3	0	3.6	171.97	0	29.02
Arrival Type						
HS	38.95	1.01	18.38	548.38	7.99	12.18
LS	21.26	0.4	11.33	297.86	1.19	14.98
HD	21.82	4.15	1.32	347.97	5.55	40.28
LD	13.23	1.03	33.88	209.95	1.89	16.44
Transition Type						
E	19.93	0.74	20.87	264.77	3.53	27.18
P	28.19	2.67	11.88	419.02	5	13.48
Annuities						
LL	28.78	5.39	11.3	364.78	10.55	55.95
LH	27.47	1.97	44.88	362.77	7.53	12.66
HL	7.07	0	3.76	274.46	0.34	6.14
HH	7.07	0	7.23	268.68	0	16.29
Contribution Margin						
LS	6.95	0.78	7.94	135.01	1.52	5.93
HS	30.91	2.79	25.97	415.8	8.28	38.25
LD	6.02	0	0	133.5	0	0
HD	35.81	51.15	0	419.44	189.67	0

Table 5.4: Average Profits (abs. values), Value of Driverless Vehicles (% inc. over only human-driven vehicles), Value of Keeping Human-Driven Vehicles (% inc. over only driverless vehicles) in Artificial Instances

average hourly return increases to 25.8$. The influencing factors on the profit generated by the service are straight-forward: A smaller number of stations results in low random

imbalances, and thus high profits. A high arrival rate and low imbalances increase revenues and decrease the necessity for rebalancing. High contribution margins and low annuities directly influence the profit/cost structure. The influence of a high arrival rate also explains why the average profits increase to 145.5$ over all x10 instances (347.1$ for instances in which service is offered).

On average the benefit of driverless vehicles is 15.0% over all base instances in which a fleet of human-driven vehicles is operated. This is significantly higher for the x10 instances (54.8%). While the extent of this increase seems surprising, it can be explained by those instances which are only offered if the arrival rate increases. In these instances, the profit often increases by a factor of more than 2 when permitting driverless vehicles. The main drivers for the value of driverless vehicles in vehicle sharing systems are low annuities for human-driven vehicles, and high contribution margins which are higher for driverless than human-driven vehicles (case HD). The influence of contribution margins is self-explanatory, and the influence of annuities can be explained by an increasing number of instances in which the service is offered, regardless of the vehicle type. The benefit of human-driven vehicles is lower than the benefit of driverless vehicles, with 9.2% in the carsharing instances and 12.4% in the ride-hailing instances. Human-driven vehicles are particularily beneficial if the imbalance in the system is low.

5.4.3 Fleet Sizing and Composition in Real-Life Case Studies

Table 5.5 and 5.6 list the profits and fleet compositions for mixed, driverless and human-driven fleets. Bold font highlights instances with mixed fleets. The per-period profits

Instance		Carsharing			Ride-Hailing		
Data Set	F^d	Mix	h	d	Mix	h	d
DiDi	0.236	**1263.7**	1228.4	1259.3	1259.3	119.7	1259.3
DiDi	0.256	**1228.5**	1228.4	1165.5	1165.5	119.7	1165.5
NYC	1.15	6467.3	6225.0	6467.3	6467.3	295.0	6467.3
NYC	1.3	6413.6	6225.0	6413.6	6413.6	295.0	6413.6

Table 5.5: Average Profits (abs. values) of the Case Study Instances

are very high for the DiDi and NYC case studies. In the NYC carsharing instances, a mixed fleet does not increase profits compared to a fleet of only driverless vehicles. The operator benefits from using driverless vehicles in all instances. The value of driverless

Instance		Carsharing			Ride-Hailing		
Data Set	F^d	Mix	h	d	Mix	h	d
DiDi	0.236	**763, 1207**	1875	1854	1854, 0	1155	1854
DiDi	0.256	**1, 1865**	1875	1827	1827, 0	1155	1827
NYC	1.15	215, 0	308	215	215, 0	139	215
NYC	1.3	215, 0	308	215	215, 0	139	215

Table 5.6: Fleet Composition (m^d, m^h) in the Case Study Instances

vehicles is up to 2.9% in the DiDi carsharing instances, and up to 3.9% in the NYC carsharing instances. In the ride-hailing instances, the profit increases by a factor of up to 10.5 (DiDi) and 21.9 (NYC) due to introducing driverless vehicles.

We, thus, confirm the result that driverless vehicles can be beneficial from the experiment with artificial data. However, mixed fleets are less commonly beneficial, since the NYC case study is an extreme case with a very high imbalance and large differences in rebalancing costs. In both case studies, the optimal fleet size is consistently smaller than anticipated. This is since operators currently maximize revenues, service level, or market coverage rather than profits. In all instances with equal contribution margins, the fleet size is smaller using driverless vehicles than human-driven vehicles. This suggests (i) that the additional flexibility to relocate vehicles (quickly) is used, and (ii) that operators do not increase their upfront investment too much. In NYC, they even decrease their investment in the fleet.

For several instances, the optimal fleet mix is using only driverless vehicles, resulting in equal profits for mixed and driverless fleets. Unsurprisingly, the highly imbalanced system in NYC benefits from driverless vehicles to the extent that only these vehicles are being used, even if there are no differences in the contribution margins ("carsharing" setting). Surprisingly, even the low imbalance in the DiDi instances is sufficient to make mixed fleets preferred to only human-driven vehicles even if customers drive themselves. This is most likely due to the very high costs: Rebalancing a human-driven vehicle costs 1.06$ per time interval (vs. 0.24$ for driverless vehicles). Even counteracting the low imbalance is then too expensive with only human-driven vehicles. If only little excess capacity is available, fleets with a single vehicle type are preferential.

Sensitivity Analysis

We are interested in what drives the choice between modes. As such, we vary four different parameters: (i) The annuities of driverless vehicles F^d, (ii) the contribution margin of human-driven vehicles r^h (relative to r^d), and (iii) the rebalancing costs of human-driven vehicles c^h (relative to c^d). We use the three case study instances, and always study the instance with equal revenue and low annuities, unless stated otherwise.

To study the influence of the first two parameters on the fleet composition, and the interdependency of the first two parameters, we use the DiDi case study instance. The influence of the parameter "rebalancing cost" on the fleet composition is studied using the NYC instance. The resulting graphs for fleet composition, profit and availability

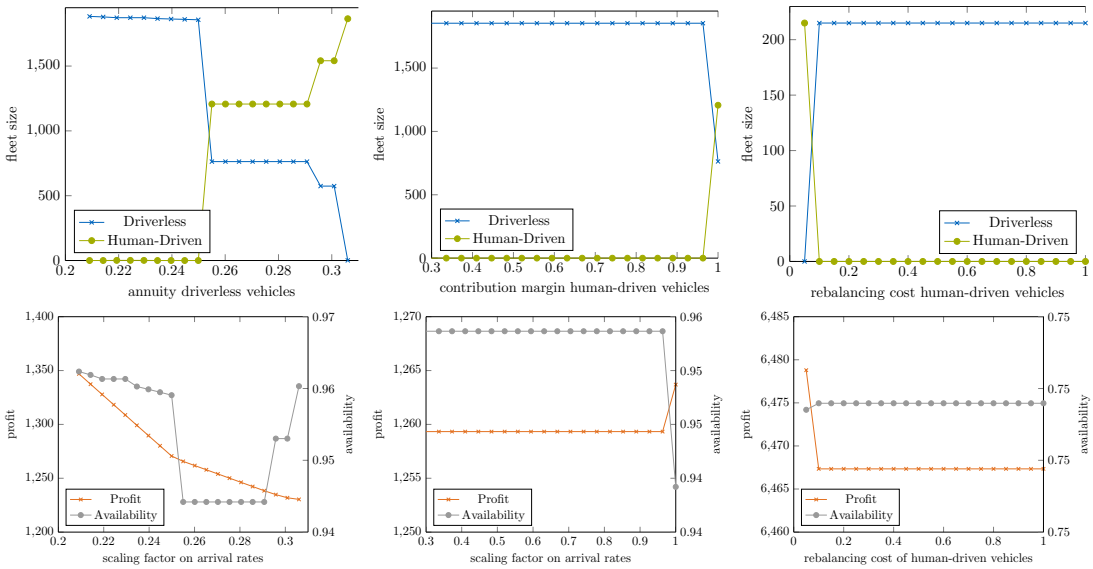

Figure 5.2: Fleet Size and Composition (1st row), as well as Profit and Availability (2nd row) for DiDi (varying annuity, contribution margin) and NYC (varying rebalancing costs) Instances

can be found in Figure 5.2. As one can see, an increasing annuity for driverless vehicles decreases the share of driverless vehicles. We observe that this does not happen continuously, but the share of human-driven vehicles directly increases from 0 to approx. 2/3. The availability is lower if the fleet is mixed, but at a similar level regardless of whether the fleet consists of only human-driven or driverless vehicles. The profits monotonely decrease in an increasing annuity. When varying contribution margins, human-driven vehicles are only beneficial if contribution margins for driverless and human-driven vehicles are identical. Otherwise, the operator is better off only procuring driverless vehicles. The availability decreases if the fleet is mixed, but the profits increase. For increasing

c^h, the fleet composition also switches directly from only driverless to only human-driven vehicles above some threshold very close to $c^h = c^d$.

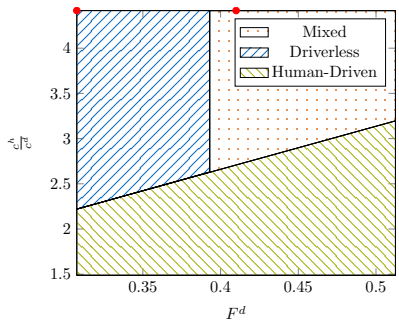

Figure 5.3: Barrier Graph for DiDi Instance

Figure 5.3 depicts the resulting barrier graphs for the interdependency between the parameters annuity and rebalancing costs. We do not show the interdependency between contribution margin and the other parameters, since even slightly unequal contribution margins result in always driverless fleets. When increasing the annuities between 0.3075 and 0.5125, the fleet composition shifts from mixed fleets from a fleet only consisting of one vehicle type which is human-driven vehicles below some threshold on $\frac{c^d}{c^h}$, and driverless vehicles above, to a mixed fleet above some threshold on $\frac{c^d}{c^h}$, and human-driven vehicles below. This threshold increases in F^d.

5.4.4 Discussion

The results of artificial instances and real-life case studies suggest that operators frequently benefit from introducing driverless vehicles into their fleet mix. If the contribution margin of driverless and human-driven vehicles differs (e.g., in ride-hailing systems), operators almost always resort to fleets only comprised of driverless vehicles. In these instances, driverless vehicles are often necessary to offer the service profitably. Otherwise, operators may not replace their entire fleet by driverless vehicles, but rather resort to mixed fleets.

In the case studies, mixed fleets are less frequently optimal than in the artificial instances. This is because the DiDi fleet and the NYC taxi fleet are extreme cases: The DiDi instance has highly diverging contribution margins. Once these are equal, mixed fleets become beneficial. In New York City, demand imbalances are very high, thus, in a mixed fleet even human-driven vehicles would be rebalanced very frequently.

5.5 Conclusion

This Chapter studies the technology choice and mix problem vehicle sharing operators face once driverless vehicles become available. While driverless vehicles are cheaper to rebalance and might result in higher contribution margins per served customer, their procurement cost is higher than that of human-driven vehicles. We conduct a large numerical experiment with both artificial instances and real-world case studies. Our main insights can be summarized as follows:

Driverless vehicles substantially improve the profitability of vehicle sharing systems. On average over a large set of artificial instances, introducing driverless vehicles increases the profit by 16.8% if the arrival rate is low and 10.6% if the arrival rate is high. Two case studies suggest that the profit increase may even be larger.

Driverless vehicles may be necessary to operate the service. Driverless vehicles are most helpful in cities in which demand patterns are highly imbalanced, i.e., some locations are have more incoming trips, while others have more outgoing trips, and if contribution margins of driverless vehicles exceed that of human-driven vehicles. Thus, driverless vehicles can help towards a more wide-spread adoption of vehicle sharing. Even if the service can be offered without driverless vehicles, introducing them frequently increases the optimal fleet size, improving the service quality.

Mixed fleets can be optimal if contribution margins for driverless and human-driven vehicles are equal. In such a mixed fleet, human-driven vehicles are being rebalanced significantly less, reducing operational costs, while not incurring as high annuities as if the entire fleet were comprised of driverless vehicles. If contribution margins of driverless vehicles exceed that of human-driven vehicles, the optimal fleet composition is almost always to procure only driverless vehicles.

Mixed fleets become more frequently optimal when the total demand increases. The number of artificial instances with mixed fleets as optimal procurement strategy increases from 7 to 26 if the arrival rates are multiplied by 10. This is since the larger demand results in larger fleets, which in turn increases pooling benefits and the availability.

Low contribution margins almost always result in unprofitable operations. We find that the profitability if very low in many instances, and most vehicle sharing systems are unprofitable if only direct revenues are considered, i.e., without subsidies (from the city government or a parent company), or a high value of customer retention (e.g. if the operator assumes that the number of customers or revenues will increase).

The technology choice and mix problem in vehicle sharing may be extended along

different avenues. So far, we have ignored that customers may have preferences for driverless or human-driven vehicles. The fleet mix will trivially be driven towards the preferred vehicle type, or towards a mix of both (if different customers have different preferences). Customer preferences become more relevant in competitive settings. Competition is a common trend in technology choice literature (e.g., Goyal and Netessine (2007)), and also commences in the literature on shared mobility (e.g., Chapter 3). In a competitive setting, an early adopter may increase the availability to prevent competitors from entering the market. In practice, demand patterns are non-stationary. So far, only few papers study non-stationary demand (e.g., Hao et al. (2020), He et al. (2020), and Tang et al. (2020)), and our analyses on fleet size and mix can be adapted to them. However, most results presented herein remain valid.

Appendix 5.A Proofs

5.A.1 Proof for Lemma 5.1

We prove by contradiction. Since the algorithm iterates all fleet compositions, unless (i) $m^d + m^h > m^{\max}$, or (ii) $\pi\left(m^d + m^h, 0\right) - F^d m^d - F^h m^h < \underline{\Pi}$, we must prove that no solution that does not fulfill both can be an optimal solution, or vice versa, that any optimal solution fulfills both. The upper bound on the fleet size follows from Corollary 5.1 (proven in Appendix 5.A.3). Case (ii) follows directly from Property 1, and will be proven for both second-stage algorithms independently.

5.A.2 Proof for Lemma 5.2

We prove by contradiction. Let's assume that inequality (5.2) holds, but $m^{d\prime}, m^{h\prime}$ is optimal. Then,

$$\Pi\left(m^d, m^h\right) < \Pi\left(m^{d\prime}, m^{h\prime}\right) \tag{5.9}$$

must hold. By adding (5.2) and (5.9), the contradiction becomes obvious:

$$\begin{aligned}
\Pi\left(m^d, m^h\right) + \bar{\pi} - \pi\left(m^d, m^h\right) &< \Pi\left(m^{d\prime}, m^{h\prime}\right) + F^d\left(m^{d\prime} - m^d\right) \\
&\quad + F^h\left(m^{h\prime} - m^h\right) \\
\pi\left(m^d, m^h\right) - F^d m^d - F^h m^h + \bar{\pi} - \pi\left(m^d, m^h\right) &< \pi\left(m^{d\prime}, m^{h\prime}\right) - F^d m^{d\prime} - F^h m^{h\prime} \\
&\quad + F^d\left(m^{d\prime} - m^d\right) + F^h\left(m^{h\prime} - m^h\right) \\
\bar{\pi} &< \pi\left(m^{d\prime}, m^{h\prime}\right)
\end{aligned}$$

This contradicts the assumption that $\pi\left(m^d, m^h\right)$ and $\pi\left(m^{d\prime}, m^{h\prime}\right)$ are bounded from above by $\bar{\pi}$.

The second part of this Lemma follows from the non-decreasing costs on the right hand side: Any $\hat{m} \geq m'$ necessarily incurs higher annuities, and can therefore not be optimal, either.

5.A.3 Proof for Corollary 5.1

We prove by contradiction. We apply Lemma 5.2 and compare to the fleet $\langle 0, 0 \rangle$. Then, any fleet such that

$$\bar{\pi} \geq F^d m^d + F^h m^h$$

cannot be optimal. Since F^d and F^h are positive, the largest number of vehicles of one type can be achieved if the other is 0. Thus, inequality (5.2) can be restated as

$$m^d \leq \frac{\bar{\pi}}{F^d}$$
$$m^h \leq \frac{\bar{\pi}}{F^h}.$$

The largest fleet size that fulfills these inequalities is then given by rounding down to the next integer.

5.A.4 Proof for Theorem 5.1

We prove (i) quasi-concavity, (ii) non-decreasingness, (iii) non-negativity, and (iv) a preference for driverless vehicles. Parts 1-3 are trivial properties of the SMDP, and part (iv) is proven by comparing policies.

Part (i): First, we prove quasi-concavity. Koole (2006) proves that quasi-concavity propagates through maximization and addition. When calculating the profit of the second stage using value iteration, these are the necessary operations. Thus, the profits of the second-stage problem are quasi-concave as well.

Part (ii): We observe that the profit function of the (exact) second-stage problem can only increase when adding an additional vehicle: Even if this vehicle remains at its original location, another vehicle will not decrease the profit.

Part (iii), the non-negativity, will follow from part (ii) and part (iv).

Part (iv): To prove that "making vehicles driverless" never decreases operational profits we compare policies. Assume that R^1 is the optimal policy with additional driverless vehicles and R^3 is the optimal policy with fewer driverless vehicles. To be able to compare policies, we assign x driverless vehicles in the solution 1 the type "human-driven". Thus, in the 3 solution they appear in the inventories for human-driven vehicles, but incur lower cost during rebalancing (reducing the expected rebalancing costs, but rebalancing costs of driverless vehicles remain strictly lower than rebalancing costs for

human-driven vehicles). The optimal policy in this case is R^2. Since policy R^3 is feasible for setting 2, and in all states either incurs equal rebalancing costs (if driverless vehicles or no vehicles are rebalanced) or higher rebalancing costs (if human-driven vehicles are rebalanced),

$$\pi_{m^d+x,m^h-x}\left(R^2\right) \geq \pi_{m^d,m^h}\left(R^3\right)$$

holds. Also, R^1 cannot return lower profits than R^2, since the "proper" rebalancing vehicles can be selected. In fact, human-driven vehicles will only be rebalanced, if no driverless vehicles are available at the location of origin. In setting 2, it is possible that this scenario there are driverless vehicles at the origin location, but they have been assigned the human-driven type. In these states, the marginal profit is lower. In all other states, the marginal profit is identical. Thus,

$$\pi_{m^d+x,m^h-x}\left(R^1\right) \geq \pi_{m^d+x,m^h-x}\left(R^2\right)$$

holds. Both previous inequalities together prove the assumption: Operators always prefer driverless vehicles over human-driven vehicles on an operational planning horizon.

5.A.5 Proof for Corollary 5.2

We first reduce the FLP to the model of Braverman et al. (2019) to prove that objective (5.5a) is an upper bound for any policy, and then show that this upper bound does not exceed some bound.

To prove that objective (5.5a) is an upper bound for any (state-dependent or state-independent) policy for a fleet of the same size, we first observe that the constraints in (5.5) with $p_i^d = 1\ \forall i \in I$ can be transformed into the model of Braverman et al. (2019) by dividing all λ_i, f_{ij}^d, and e_{ij}^d by m^d. Then, Theorem 2 and Remark 5 of Braverman et al. (2019) hold which state that if the objective function is (i) non-decreasing in a (true since $a = p_i^d$ is monotonely increasing), (ii) non-decreasing in f_{ij} (true since $\lambda_i p_i^d = \sum_{j\in I} \mu_{ij} f_{ij}^d$ is monotonely increasing), (iii) non-increasing in $e_{ij}, i \neq j$ (true sind monotonely decreasing), independent from e_{ii} (true since e_{ii} does not appear in the objective), and concave in e_{ij} and f_{ij} (true since linear or independent), the optimal objective function value is a (strict) upper bound on the profit of any routing policy.

Second, we show that $\bar{\pi}\left(m^d, m^h\right)$ is bounded. If the fleet size is infinitely large, all customers will be served, assuming that the revenues are sufficiently high compared to

the rebalancing cost. The rebalancing problem can then be reformulated as a transportation problem with the demand imbalance (difference between number of incoming and outgoing vehicles at a station) as pickup or delivery demand. This transportation problem has a cost-minimizing solution which translates to rebalancing rates (number of vehicles the operator should rebalance per period of time, λ'_{ij}) which in turn can be transformed into the expected number of empty vehicles, $e^d_{ij} = \frac{\lambda'_i}{\mu_{ij}}$. Let's denote the minimum number of vehicles necessary to rebalance in the most cost-efficient manner by $\underline{m}^d$. Any additional vehicle above $\lfloor \underline{m}^d \rfloor$ cannot increase revenues, and does not decrease cost, and thus does not influence the objective function, resulting in a (constant) upper bound on any state-dependent policy.

5.A.6 Proof for Lemma 5.3

We define the system with m^d driverless and m^h human-driven vehicles as two coupled closed queueing networks. Analogously, we can define the system with m^d+m^h driverless vehicles (without human-driven vehicles) as two coupled closed queueing networks with m^d and m^h vehicles, respectively. Then, the routing, and thus the availability, in the first system remains the same. The routing in the second system remains feasible, but not necessarily optimal in the most general case. However, if rebalancing costs c^α_{ij} and contribution margins r^α_{ij} are proportional to the traversal time $\frac{1}{\mu_{ij}}$, a routing solution that is optimal for human-driven vehicles is also optimal for driverless vehicles.

We show that adapting the routing decision cannot improve profits: The routing decision can only improve in the formerly human-driven queueing system. In this system, all demand assigned to the system is served, and since $r^d_{ij} \geq r^h_{ij}$ and $c^d_{ij} \leq c^h_{ij}$ as well as $r^d_{ij} \geq c^d_{ji}$, all demand will be served after merging. Thus, we can safely ignore contribution margins r^d_{ij}, and instead focus on minimizing the cost. Henceforth, the let γ^α be defined such that $c^\alpha_{ij} = \frac{\gamma^\alpha}{\mu_{ij}}$. The assumption that $c^d_{ij} \leq c^h_{ij}\ \forall i,j$ then translates to $\gamma^d \leq \gamma^h$. Let's assume that some rebalancing policy $\mathbf{e}$ is optimal for the human-driven system, but some other rebalancing policy $\mathbf{e}'$ is optimal if the same demand is served with driverless vehicles. This then results in the following two optimality conditions:

$$\sum_{i \neq j} e_{ij}\mu_{ij}c^h_{ij} < \sum_{i \neq j} e'_{ij}\mu_{ij}c^h_{ij}$$
$$\sum_{i \neq j} e'_{ij}\mu_{ij}c^d_{ij} < \sum_{i \neq j} e_{ij}\mu_{ij}c^d_{ij}$$

When replacing the costs c_{ij}^d and c_{ij}^h, this reads

$$\sum_{i \neq j} e_{ij} \gamma^h < \sum_{i \neq j} e'_{ij} \gamma^h$$
$$\sum_{i \neq j} e'_{ij} \gamma^d < \sum_{i \neq j} e_{ij} \gamma^d$$

which cannot be satisfied.

If the since the calculation of the availability only depends on the arrival rates $\tilde{\lambda}_i$ and expected travel times $\frac{1}{\mu_{ij}}$, it directly follows that if the routing does not change, the availability A^α does not change either.

If both systems are being merged, the availability is thus at least the same (and often higher).

5.A.7 Proof for Theorem 5.2

We prove (i) concavity (if either $m^d = 0$ or $m^h = 0$), (ii) non-decreasingness (if either $m^d = 0$ or $m^h = 0$), (iii) non-negativity, and (iv) a preference for driverless vehicles.

Part (i): We first prove that the expected operational profit function is concave and increasing for $m^d \geq \underline{m}$, and $m^h = 0$. Here, $\underline{m}$ is the number of vehicles necessary such that all rebalancing can be performed on the cheapest rather than fastest rebalancing arc, and still fulfilling all customer requests. It follows from Corollary 5.2 that the optimal solution to (5.5) is constant in m^d.

It is obvious that the objective function value of (5.5) is non-negative, since $p_i^d = p_i^h = 0 \forall i \in I$, $e_{ij}^d = e_{ij}^h = 0 \forall i, j \in I, i \neq j$, $e_{ii}^\alpha = \frac{m^\alpha}{n} \forall i \in I, \alpha \in \{d, h\}$, and $f_{ij}^d = f_{ij}^h = 0 \forall i \in I$ is a feasible solution that yields the objective function value 0.

Shanthikumar and Yao (1987) prove that MVA yields concavely increasing availabilities if m^d increases. Obviously, concavity and increasingness survive multiplication by a non-negative constant (the objective function value of (5.5)). By symmetry, the same holds for $m^d = 0$ and $m^h \geq \underline{m}$.

Part (ii): Next, we show that the expected operational profit is also increasing for $m^d < \underline{m}$ and $m^h = 0$ ($m^h < \underline{m}$ and $m^d = 0$). The expected operational profit is a product of the objective function value to (5.5) and the availability. Since both are clearly non-negative and increasing, the expected operational profit must also be non-negative and increasing.

Part (iii): By the same argument, non-negativity persists even if both m^d and m^h are strictly positive. Then, the objective function value to (5.5) and the availability are

both non-negative, and so is their product.

Part (iv): If costs and contribution margins are proportional to the expected travel time, Lemma 5.3 states that the availability of an purely driverless fleet is at least as high as the availability of any mixed fleet of the same total size. Also, as stated before, the availability is always non-negative. Also, the objective function value of (5.5) necessarily increases the higher the share of driverless vehicles in the system grows. It is obvious that the objective function value of (5.5) is non-negative. The expected operational profit of a purely driverless fleet is the product of two non-negative values that are both at least as high as the corresponding values for mixed fleets. Thus, a purely driverless fleet always incurs at least as high expected operational profits as any mixed fleet of the same size.

Chapter 6

Conclusions

6.1 Summary

This thesis studies different aspects of rebalancing shared mobility systems, focusing on one-way station-based and one-way free-floating carsharing systems as well as ride-hailing. In such systems, rebalancing is paramount to increase availability of vehicles and avoid piling of vehicles in some regions. Consequentially, rebalancing leads to higher customer satisfaction and higher profits for operators.

Chapter 3 investigates profit gains due to considering the presence of competition and losses due to the presence of competition, as well as features that drive profit gains and losses. We present a novel model formulation, the Competitive Pickup and Delivery Orienteering Problem (C-PDOP), that models competition between multiple one-way carsharing operators that rebalance their fleets simultaneously. The C-PDOP is solved for pure strategy Nash equilibria using Iterated Best Response (IBR) and Potential Function Optimizer (PFO).

RQ 1.1 *How much can operators gain from considering the presence of competition in their rebalancing operations with regards to gross profits? Put differently, what is the price of ignoring the presence of competition?*

Using artificial data, profit gains can be up to several orders of magnitude. Using data from a case study settled in Munich, Germany, with two operators, we quantify the profit gains to be 35% on average over both operators, if the operators otherwise ignore the presence of competition (optimistic). If the operators otherwise assume that their competitor has vehicles at all competitive locations (pessimistic), the operators can gain on average 12% of their profits.

RQ 1.2 *How much is lost by competing in comparison to jointly optimizing fleet re-*

balancing with regard to gross profits, and how do alternative business models under competition compare to each other?
We find that welfare-maximization in the routing (joint optimization of the sum over both objectives) does not improve profitability by a large margin, compared to equilibrium profits. For most operators, these minor profit improvements do not justify the additional coordination effort. Under two minor restrictions, merging the fleets, or outsourcing rebalancing to a third-party operator provably increases profitability. Merging the fleets results in at least as high profits as outsourcing, and equal profits if all customers have memberships with all operators. On artificial data, this improvement can be several orders of magnitude, whilst in a Munich, Germany, case study, operators only lose 10% by competing rather than merging.
RQ 1.3 *Which features drive the gains from considering competition, and the losses due to the presence of competition?*
The profit gain due to considering competition increases if competition is fierce, given by a high fraction of customers with multiple memberships. The profit loss due to presence of competition decreases with increasingly fierce competition. An increasing instance size decreases the percentage loss due to the presence of competition. The profit gain also increases in a high number of operators, but decreases if the size of stations (excess demand per station) increases. Further, inhomogeneous payoffs (under full competition), and customers preferring one operator over the other (only for the less preferred operator), increase the profit gain due to considering competition.

In Chapter 4 we study the mode selection problem of carsharing operators: Vehicles can either be loaded onto a truck, or be driven by workers. In the latter case, the workers have to continue to the next vehicle (or return to the depot) by biking, hitching rides with colleagues, or by using public transit. We support operators in their modal decision and their decision to hybridize among different modes by building classifiers and providing insights into the key features driving the modal choice.
RQ 2.1 *Can a good mode be selected a-priori based upon features of the fleet and city?*
Yes, the best mode can be determined with an accuracy of more than 90% using linear regression, logistic regression, and one of the three suggested decision trees. With these classifiers, the excess cost due to misclassification is low (less than 10% of the total cost on average over all misclassified instances). Most instances in which the algorithms fail to return the optimal mode are borderline cases in the sense that given the set of features, the optimal mode differs between days, and hybridization among various modes often reduces costs.

RQ 2.2 *Which features drive the choice of the optimal rebalancing mode?*
We find that in most instances, the optimal mode is either bike or truck, and bike is preferred, unless worker wages are very high, or costs for vehicles (fuel, wear, tear, deprecation) is not too low. Public transit and car are only the best mode, if the modes bike and truck are impeded, e.g., by very low accessibility by truck and very low velocities for the mode bike.

Chapter 5 focuses on the technology choice problem of shared mobility operators arising when driverless vehicles become available. While driverless vehicles are cheaper to rebalance and might result in higher contribution margins per customer, they incur substantially higher investment costs, in particular in an early stage after their introduction. Operators thus face the problem of determining the optimal fleet size and mix. The technology choice problem is solved using a bound-and-enumerate algorithm, and the rebalancing problems (for a fixed fleet mix) are solved using either Semi-Markov Decision Processes (SMDPs) or Fluid-Based Approximation Linear Program (FLP).
RQ 3.1 *Should shared mobility operators use fleets with only one vehicle type, or mix among driverless and human-driven vehicles?*
While in some instances driverless vehicles do not increase profits, there are many instances for which the possibility of introducing driverless vehicles in the fleet mix increases profits. If contribution margins differ between vehicle types, operators should almost always only use driverless vehicles. Mixed fleets become increasingly important if the fleet size is large due to a large customer base.
RQ 3.2 *Under which circumstances can shared mobility operators benefit from introducing driverless vehicles in their fleet?*
Operators benefit from introducing driverless vehicles in their fleet mix in most instances, unless rebalancing costs and contribution margins of driverless and human-driven vehicles are almost identical, and annuities differ largely between vehicle types.
RQ 3.3 *How much can operators gain with respect to total profits from using mixed fleets comprised of driverless vehicles?*
In some instances, driverless vehicles are necessary for the shared mobility operator to reach profitability, and thus offer their service. They can improve their profitability by several orders of magnitude if instances are highly imbalanced and contribution margins differ. The benefit of mixed fleets over fleets only consisting of driverless vehicles is substantially lower (on average 2%).

6.2 Limitations and Future Research

The methodological contributions and managerial insights developed in this thesis can be extended in different directions.

In Chapter 3, we assume that the profit functions of the competitors are known. In practice, operators must first observe their operators' reactions, to be able to determine the profit function. Future research could investigate how to learn a profit function from the limited information available (e.g., vehicle position, pricing structure). Once a profit function is established for the non-competitive case, additional research can extend this to learning own demand if competition is present by including the availability of the competitor in the data for the demand prediction algorithms. Further, timing of rebalancing operations may influence the insights: Our numerical experiments show that it is possible that one operator refrains from rebalancing altogether. Possibly, operators might benefit from rebalancing on different days (then, the problem should be studied as a repeated game). Further, operators may benefit from delaying their rebalancing activities until the other operator has finished (in particular if multiple equilibria exist). If both operators delay their actions, they might reach a prisoner's dilemma. Future research can extend the C-PDOP to a repeated game, and consider the timing of rebalancing (timing and frequency of rebalancing operations is also an open problem in the non-competitive case).

Chapter 4 currently only studies operator-based rebalancing in a homogeneous operating area. A similar feature-based classifier can help to determine if an operator should rebalance entirely herself, or additionally offer incentives for users to rebalance. If user-based rebalancing is used, this may also affect the optimal mode with which the remaining fleet is rebalanced. Extending the feature-based classifier to different features within the operating area (e.g., city center with high density of locations and low velocities vs. suburban areas with lower density and higher velocities) is not trivial, and future research can address this line of research either by integrating the feature-based mode selection with districting, or by enhancing the algorithms to explicitly consider the variability of features. Currently, the Carsharing Relocation Problem (CRP) is restricted to rebalancing during the night when demand is negligible. If the operator also rebalanced during the day to counteract short-term imbalances, the longer vehicle unavailability and inflexibility when loaded onto a truck might deter operators from using this mode. Future research can address this tradeoff between a mode with a higher capacity and less flexibility, both from a modelling and an analytics perspective. We model the CRP

as a cost-minimizing problem. In line with Chapter 3, the CRP could be include the profitability of locations. While this could result in some locations being frequently omitted (and thus not fulfill the goal of a balanced fleet), this could also improve the performance of modes with a high variance in rebalancing costs (e.g., bike).

We assume that when driverless vehicles become available, operators will start offering shared mobility, or directly replace their entire fleet. In practice, operators replace vehicles over time. Since our numerical study showed that having single driverless vehicles is often not beneficial, technology diffusion may thwart the adoption of driverless vehicles. Further, the cost difference between driverless and human-driven vehicles will decrease over time. Thus, operators may put off procuring driverless vehicles. Future research may address the question how quickly the fleet is replaced to reach the target state according to Chapter 5.

An interesting line for future research is integrating competition and fleet mix, that are studied separately in Chapters 3 and 5. Competition might speed up the adoption of driverless vehicles if frequent rebalancing is necessary to compete on availability, but might as well slow down the adoption if operators increase their fleet size to catch more customers and thus rebalance less.

In all chapters, we assume that the demand processes at different stations (or in different districts in the case of free-floating systems) are independent, i.e., a customer who cannot be served at a station is lost, rather than walking to a nearby station if the other station has ample supply. There exist first approaches for measuring this effect (e.g., Kabra et al. (2020)). Future research could (i) develop modelling approaches for addressing this behavior in rebalancing problems which can for example draw upon results from attractivity in facility location problems, and (ii) measure the cost reduction and profit increase that becomes possible if the customer's willingness to walk is explicitly considered. In Chapter 3, this mainly affects the payoff function and results in non-linearities for which a novel algorithmic treatment needs to be developed. In Chapter 5, one can first study how much the availability increases, and subsequently also devise methods to address interdependent demand. Currently, we assume that the demand (or the demand distribution) is stationary and exogeneous. While demand stationarity is not very restrictive if vehicles are rebalanced periodically during the night, demand is subject to change during the day and, thus, affects dynamic rebalancing. Varying demand may impact the optimal fleet composition, since driverless vehicles can better react to spatio-temporal demand imbalances, but also incur higher costs during periods of low demand (e.g., the night). Future research could first investigate the optimal fleet

size under non-stationary demand, and then extend this to the optimal fleet composition. If demand is partially endogeneous, that is if we consider return trips which require that a vehicle was available for the first part of the trip, and that longer periods of low availability reduce the trust in vehicle sharing, we must integrate the demand process into the rebalancing models.

Bibliography

Ackermann, Heiner, Heiko Röglin, and Berthold Vöcking (2008). On the impact of combinatorial structure on congestion games. *Journal of the ACM* 55 (6), p. 25.

Afeche, Philipp, Zhe Liu, and Costis Maglaras (2018). Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. Columbia Business School Research Paper.

Al-Kanj, Lina, Juliana Nascimento, and Warren B Powell (2020). Approximate dynamic programming for planning a ride-hailing system using autonomous fleets of electric vehicles. *European Journal of Operational Research* 284 (3), pp. 1088–1106.

Albiński, Szymon and Stefan Minner (2020). Competitive rebalancing in one-way carsharing. Working Paper.

Almeida Correia, Gonçalo Homem de and António Pais Antunes (2012). Optimization approach to depot location and trip selection in one-way carsharing systems. *Transportation Research Part E: Logistics and Transportation Review* 48 (1), pp. 233–247.

Ampudia-Renuncio, María, Begoña Guirao, Rafael Molina-Sánchez, and Cristina Engel de Álvarez (2020). Understanding the spatial distribution of free-floating carsharing in cities: Analysis of the new Madrid experience through a web-based platform. *Cities* 98, p. 102593.

Anshelevich, Elliot, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden (2008). The price of stability for network design with fair cost allocation. *SIAM Journal on Computing* 38 (4), pp. 1602–1623.

BMW Group (2018). BMW Group und Daimler AG vereinbaren Bündelung ihrer Mobilitätsdienste. `https://www.press.bmwgroup.com/deutschland/article/detail/T0279654DE/bmw-group-und-daimler-ag-vereinbaren-buendelung-ihrer-mobilitaetsdienste`.

Balac, Milos, Henrik Becker, Francesco Ciari, and Kay W Axhausen (2019). Modeling competing free-floating carsharing operators – A case study for Zurich, Switzerland. *Transportation Research Part C: Emerging Technologies* 98, pp. 101–117.

Banerjee, Siddhartha, Daniel Freund, and Thodoris Lykouris (2017). Pricing and optimization in shared vehicle systems: An approximation framework. *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 517–517.

Banerjee, Siddhartha and Ramesh Johari (2019). Ride sharing. *Sharing Economy*. Springer, pp. 73–97.

Baptista, Patrícia, Sandra Melo, and Catarina Rolim (2014). Energy, environmental and mobility impacts of car-sharing systems. Empirical results from Lisbon, Portugal. *Procedia-Social and Behavioral Sciences* 111 (0), pp. 28–37.

Baron, Opher, Oded Berman, and Mehdi Nourinejad (2018). Introducing autonomous vehicles: formulation and analysis. Available at SSRN 3250557.

Basciftci, Beste, Shabbir Ahmed, and Siqian Shen (2020). Distributionally robust facility location problem under decision-dependent stochastic demand. *European Journal of Operational Research*, Forthcoming.

Bellos, Ioannis, Mark Ferguson, and L Beril Toktay (2017). The car sharing economy: Interaction of business model choice and product line design. *Manufacturing & Service Operations Management* 19 (2), pp. 185–201.

Benjaafar, Saif, Daniel Jiang, Xiang Li, and Xiaobo Li (2018). Dynamic inventory repositioning in on-demand rental networks. Available at SSRN 2942921.

Bernstein, Fernando, Gregory A DeCroix, and N Bora Keskin (2020). Competition between two-sided platforms under demand and supply congestion effects. *Manufacturing & Service Operations Management*, Forthcoming.

Bimpikis, Kostas, Ozan Candogan, and Daniela Saban (2019). Spatial pricing in ride-sharing networks. *Operations Research* 67 (3), pp. 744–769.

Boston Consulting Group (2020). Can self-driving cars stop the urban mobility meltdown? `https://www.bcg.com/publications/2020/how-autonomous-vehicles-can-benefit-urban-mobility`.

Boyacı, Burak, Konstantinos G Zografos, and Nikolas Geroliminis (2015). An optimization framework for the development of efficient one-way car-sharing systems. *European Journal of Operational Research* 240 (3), pp. 718–733.

Braverman, Anton, Jim G Dai, Xin Liu, and Lei Ying (2019). Empty-car routing in ridesharing systems. *Operations Research* 67 (5), pp. 1437–1452.

Brook, David (2004). Carsharing - Start up issues and new operational models. *Transportation Research Board 83rd Annual Meeting, Washington, DC*.

Bruck, Bruno P, Fábio Cruz, Manuel Iori, and Anand Subramanian (2019). The static bike sharing rebalancing problem with forbidden temporary operations. *Transportation Science* 53 (3), pp. 882–896.

Bruglieri, Maurizio, Alberto Colorni, and Alessandro Luè (2014a). The relocation problem for the one-way electric vehicle sharing. *Networks* 64 (4), pp. 292–305.

— (2014b). The vehicle relocation problem for the one-way electric vehicle sharing: an application to the Milan case. *Procedia-Social and Behavioral Sciences* 111, pp. 18–27.

Bruglieri, Maurizio, Ferdinando Pezzella, and Ornella Pisacane (2017). Heuristic algorithms for the operator-based relocation problem in one-way electric carsharing systems. *Discrete Optimization* 23, pp. 56–80.

— (2018). A two-phase optimization method for a multiobjective vehicle relocation problem in electric carsharing systems. *Journal of Combinatorial Optimization*, pp. 1–32.

Chen, Le, Alan Mislove, and Christo Wilson (2015). Peeking beneath the hood of Uber. *Proceedings of the 2015 Internet Measurement Conference*, pp. 495–508.

Dandl, Florian and Klaus Bogenberger (2018). Comparing future autonomous electric taxis with an existing free-floating carsharing system. *IEEE Transactions on Intelligent Transportation Systems* 20 (6), pp. 2037–2047.

Datner, Sharon, Tal Raviv, Michal Tzur, and Daniel Chemla (2019). Setting inventory levels in a bike sharing network. *Transportation Science* 53 (1), pp. 62–76.

Dror, Moshe, Dominique Fortin, and Catherine Roucairol (1998). *Redistribution of self-service electric cars: A case of pickup and delivery*. Tech. rep. INRIA.

Fabrikant, Alex, Christos Papadimitriou, and Kunal Talwar (2004). The complexity of pure Nash equilibria. *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*. ACM, pp. 604–612.

Feillet, Dominique, Pierre Dejax, and Michel Gendreau (2005). Traveling salesman problems with profits. *Transportation Science* 39 (2), pp. 188–205.

Feller, William (1950). *An introduction to probability theory and its applications*. Vol. 1. Wiley.

Feng, Guiyun, Guangwen Kong, and Zizhuo Wang (2020). We are on the way: Analysis of on-demand ride-hailing systems. *Manufacturing & Service Operations Management*, Forthcoming.

Fernández, Elena, Mireia Roca-Riu, and M Grazia Speranza (2018). The shared customer collaboration vehicle routing problem. *European Journal of Operational Research* 265 (3), pp. 1078–1093.

Fine, Charles H and Robert M Freund (1990). Optimal investment in product-flexible manufacturing capacity. *Management Science* 36 (4), pp. 449–466.

Fink, Andreas and Torsten Reiners (2006). Modeling and solving the short-term car rental logistics problem. *Transportation Research Part E: Logistics and Transportation Review* 42 (4), pp. 272–292.

Firnkorn, Jörg and Martin Müller (2011). What will be the environmental effects of new free-floating car-sharing systems? The case of car2go in Ulm. *Ecological Economics* 70 (8), pp. 1519–1528.

Freund, Daniel, Shane G Henderson, and David B Shmoys (2018). Minimizing multimodular functions and allocating capacity in bike-sharing systems. *Production and Operations Management* 27 (12), pp. 2346–2349.

— (2019). Bike sharing. *Sharing Economy*. Springer, pp. 435–459.

Freund, Daniel, Ashkan Norouzi-Fard, Alice Paul, Carter Wang, Shane G Henderson, and David B Shmoys (2020). Data-driven rebalancing methods for bike-share systems. *Analytics for the Sharing Economy: Mathematics, Engineering and Business Perspectives*. Springer, pp. 255–278.

Gambella, Claudio, Enrico Malaguti, Filippo Masini, and Daniele Vigo (2018). Optimizing relocation operations in electric car-sharing. *Omega* 81, pp. 234–245.

George, David K and Cathy H Xia (2011). Fleet-sizing and service availability for a vehicle rental system via closed queueing networks. *European Journal of Operational Research* 211 (1), pp. 198–207.

Ghosh, Supriyo, Pradeep Varakantham, Yossiri Adulyasak, and Patrick Jaillet (2017). Dynamic repositioning to reduce lost demand in bike sharing systems. *Journal of Artificial Intelligence Research* 58, pp. 387–430.

Goyal, Manu and Serguei Netessine (2007). Strategic technology choice and capacity investment under demand uncertainty. *Management Science* 53 (2), pp. 192–207.

Greenblatt, Jeffery B and Samveg Saxena (2015). Autonomous taxis could greatly reduce greenhouse-gas emissions of US light-duty vehicles. *Nature Climate Change* 5 (9), pp. 860–863.

Guda, Harish and Upender Subramanian (2019). Your Uber is arriving: Managing on-demand workers through surge pricing, forecast communication, and worker incentives. *Management Science* 65 (5), pp. 1995–2014.

Hao, Zhaowei, Long He, Zhenyu Hu, and Jun Jiang (2020). Robust vehicle pre-allocation with uncertain covariates. *Production and Operations Management* 29 (4), pp. 955–972.

He, Long, Zhenyu Hu, and Meilin Zhang (2020). Robust repositioning for vehicle sharing. *Manufacturing & Service Operations Management* 22 (2), 241–256.

He, Long, Ho-Yin Mak, and Ying Rong (2019). Operations management of vehicle sharing systems. *Sharing Economy*. Springer, pp. 461–484.

He, Long, Ho-Yin Mak, Ying Rong, and Zuo-Jun Max Shen (2017). Service region design for urban electric vehicle sharing systems. *Manufacturing & Service Operations Management* 19 (2), pp. 309–327.

Herbawi, Wesam, Martin Knoll, Marcus Kaiser, and Wolfgang Gruel (2016). An evolutionary algorithm for the vehicle relocation problem in free floating carsharing. *2016 IEEE Congress on Evolutionary Computation (CEC)*. IEEE, pp. 2873–2879.

Huang, Kai, Goncalo Homem de Almeida Correia, and Kun An (2018). Solving the station-based one-way carsharing network planning problem with relocations and non-linear demand. *Transportation Research Part C: Emerging Technologies* 90, pp. 1–17.

Huang, Kai, Kun An, Jeppe Rich, and Wanjing Ma (2020). Vehicle relocation in one-way station-based electric carsharing systems: A comparative study of operator-based and user-based methods. *Transportation Research Part E: Logistics and Transportation Review* 142, p. 102081.

Illgen, Stefan and Michael Höck (2019). Literature review of the vehicle relocation problem in one-way car sharing networks. *Transportation Research Part B: Methodological* 120, pp. 193–204.

Jiang, Zhong-Zhong, Guangwen Kong, and Yinghao Zhang (2020). Making the most of your regret: Workers' relocation decisions in on-demand platforms. *Manufacturing & Service Operations Management*, Forthcoming.

Jordan, William C and Stephen C Graves (1995). Principles on the benefits of manufacturing process flexibility. *Management Science* 41 (4), pp. 577–594.

Jorge, Diana and Gonçalo Correia (2013). Carsharing systems demand estimation and defined operations: A literature review. *European Journal of Transport and Infrastructure Research* 13 (3).

Kabra, Ashish, Elena Belavina, and Karan Girotra (2020). Bike-share systems: Accessibility and availability. *Management Science* 66 (9), pp. 3803–3824.

Karp, Richard M (1972). Reducibility among combinatorial problems. *Complexity of Computer Computations*. Springer, pp. 85–103.

Kek, Alvina GH, Ruey Long Cheu, Qiang Meng, and Chau Ha Fung (2009). A decision support system for vehicle relocation operations in carsharing systems. *Transportation Research Part E: Logistics and Transportation Review* 45 (1), pp. 149–158.

Kök, A Gürhan, Kevin Shang, and Şafak Yücel (2020). Investments in renewable and conventional energy: The role of operational flexibility. *Manufacturing & Service Operations Management* 22 (5), pp. 925–941.

Koole, Ger (2006). Monotonicity in Markov reward and decision chains: Theory and applications. *Foundations and Trends in Stochastic Systems* 1 (1), pp. 1–76.

Kortum, Katherine, Robert Schönduwe, Benjamin Stolte, and Benno Bock (2016). Free-floating carsharing: City-specific growth rates and success factors. *Transportation Research Procedia* 19, pp. 328–340.

Koutsoupias, Elias and Christos Papadimitriou (1999). Worst-case equilibria. *Annual Symposium on Theoretical Aspects of Computer Science*. Vol. 1563. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg, pp. 404–413.

Krishnamoorthy, Mukkai S (1975). An NP-hard problem in bipartite graphs. *ACM SIGACT News* 7 (1), pp. 26–26.

Krumke, Sven O, Alain Quilliot, Annegret K Wagler, and Jan-Thierry Wegener (2013). Models and algorithms for carsharing systems and related problems. *Electronic Notes in Discrete Mathematics* 44, pp. 201–206.

Kypriadis, Damianos, Grammati Pantziou, Charalampos Konstantopoulos, and Damianos Gavalas (2018). Minimum walking static repositioning in free-floating electric car-sharing systems. *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, pp. 1540–1545.

Lanzetti, Nicolas, Maximilian Schiffer, Michael Ostrovsky, and Marco Pavone (2020). On the interplay between self-driving cars and public transportation: A game-theoretic perspective. Les Cahiers du GERAD G–2020–24, GERAD, HEC Montreal, Canada.

Laporte, Gilbert, Frédéric Meunier, and Roberto Wolfler Calvo (2015). Shared mobility systems. *4OR* 13 (4), pp. 341–360.

— (2018). Shared mobility systems: an updated survey. *Annals of Operations Research* 271 (1), pp. 105–126.

Liang, Xiao, Gonçalo Homem de Almeida Correia, and Bart Van Arem (2016). Optimizing the service area and trip selection of an electric automated taxi system

used for the last mile of train trips. *Transportation Research Part E: Logistics and Transportation Review* 93, pp. 115–129.

Lu, Mengshi, Zhihao Chen, and Siqian Shen (2018). Optimizing the profitability and quality of service in carshare systems under demand uncertainty. *Manufacturing & Service Operations Management* 20 (2), pp. 162–180.

Lyft (2020). Testing Self-Driving Rides With Employees. `https://medium.com/lyftlevel5/testing-self-driving-rides-with-employees-eba94dc23d15`.

Martin, Elliot W and Susan A Shaheen (2011). Greenhouse gas emission impacts of carsharing in North America. *IEEE Transactions on Intelligent Transportation Systems* 12 (4), pp. 1074–1086.

Martin, Layla and Stefan Minner (2020). Feature-based selection of carsharing relocation modes. Working Paper.

Martin, Layla, Stefan Minner, Marco Pavone, and Maximilian Schiffer (2020a). It's all in the mix: Technology choice for vehicle sharing. Working Paper.

Martin, Layla, Stefan Minner, Diogo Poças, and Andreas S. Schulz (2020b). The competitive pickup and delivery orienteering problem for balancing carsharing systems. Working Paper.

Ming, Liu, Tunay I. Tunca, Yi Xu, and Weiming Zhu (2020). An Empirical Analysis of Market Formation, Pricing, and Revenue Sharing in Ride-Hailing Services. Available at SSRN 3338088.

Moccia, Luigi and Gilbert Laporte (2016). Improved models for technology choice in a transit corridor with fixed demand. *Transportation Research Part B: Methodological* 83, pp. 245–270.

Monderer, Dov and Lloyd S Shapley (1996). Potential games. *Games and Economic Behavior* 14 (1), pp. 124–143.

Nair, Rahul and Elise Miller-Hooks (2011). Fleet management for vehicle sharing operations. *Transportation Science* 45 (4), pp. 524–540.

— (2014). Equilibrium network design of shared-vehicle systems. *European Journal of Operational Research* 235 (1), pp. 47–61.

Nash, John (1951). Non-cooperative games. *Annals of Mathematics*, pp. 286–295.

Nourinejad, Mehdi, Sirui Zhu, Sina Bahrami, and Matthew J Roorda (2015). Vehicle relocation and staff rebalancing in one-way carsharing systems. *Transportation Research Part E: Logistics and Transportation Review* 81, pp. 98–113.

Oliveira, Beatriz Brito, Maria Antónia Carravilla, and José Fernando Oliveira (2017). Fleet and revenue management in car rental companies: A literature review and an integrated conceptual framework. *Omega* 71, pp. 11–26.

Ostrovsky, Michael and Michael Schwarz (2019). Carpooling and the economics of self-driving cars. *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 581–582.

Pandey, Venktesh, Julien Monteil, Claudio Gambella, and Andrea Simonetto (2019). On the needs for MaaS platforms to handle competition in ridesharing mobility. *Transportation Research Part C: Emerging Technologies* 108, pp. 269–288.

Parragh, Sophie N, Karl F Doerner, and Richard F Hartl (2008). A survey on pickup and delivery problems. *Journal für Betriebswirtschaft* 58 (1), pp. 21–51.

Perboli, Guido, Francesco Ferrero, Stefano Musso, and Andrea Vesco (2018). Business models and tariff simulation in car-sharing services. *Transportation Research Part A: Policy and Practice* 115, pp. 32–48.

Pfrommer, Julius, Joseph Warrington, Georg Schildbach, and Manfred Morari (2014). Dynamic vehicle redistribution and online price incentives in shared mobility systems. *IEEE Transactions on Intelligent Transportation Systems* 15 (4), pp. 1567–1578.

PriceWaterhouseCoopers (2017). Five trends transforming the Automotive Industry. `https://www.pwc.com/hu/hu/kiadvanyok/assets/pdf/five_trends_transforming_the_automotive_industry.pdf`.

Qi, Wei and Zuo-Jun Max Shen (2019). A smart-city scope of operations management. *Production and Operations Management* 28 (2), pp. 393–406.

Reiser, Martin and Stephen S Lavenberg (1980). Mean-value analysis of closed multichain queuing networks. *Journal of the ACM* 27 (2), pp. 313–322.

Repoux, Martin, Mor Kaspi, Burak Boyacı, and Nikolas Geroliminis (2019). Dynamic prediction-based relocation policies in one-way station-based carsharing systems with complete journey reservations. *Transportation Research Part B: Methodological* 130, pp. 82–104.

Rosenthal, Robert W (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2 (1), pp. 65–67.

Roughgarden, Tim and Éva Tardos (2002). How bad is selfish routing? *Journal of the ACM* 49 (2), pp. 236–259.

Santos, Gonçalo Gonçalves Duarte and Gonçalo Homem de Almeida Correia (2019). Finding the relevance of staff-based vehicle relocations in one-way carsharing sys-

tems through the use of a simulation-based optimization tool. *Journal of Intelligent Transportation Systems* 23 (6), pp. 583–604.

Schaefers, Tobias (2013). Exploring carsharing usage motives: A hierarchical means-end chain analysis. *Transportation Research Part A: Policy and Practice* 47, pp. 69–77.

Schmöller, Stefan, Simone Weikl, Johannes Müller, and Klaus Bogenberger (2015). Empirical analysis of free-floating carsharing usage: The Munich and Berlin case. *Transportation Research Part C: Emerging Technologies* 56, pp. 34–51.

Schuijbroek, J, Robert C Hampshire, and W-J Van Hoeve (2017). Inventory rebalancing and vehicle routing in bike sharing systems. *European Journal of Operational Research* 257 (3), pp. 992–1004.

Schulz, Andreas S and N Stier Moses (2003). On the performance of user equilibria in traffic networks. *Proceedings of Symposium on Discrete Algorithms (SODA)*, pp. 86–87.

Shaheen, Susan A and Adam P Cohen (2013). Carsharing and personal vehicle services: worldwide market developments and emerging trends. *International Journal of Sustainable Transportation* 7 (1), pp. 5–34.

Shaheen, Susan A, Adam P Cohen, and J Darius Roberts (2006). Carsharing in North America: Market growth, current developments, and future potential. *Transportation Research Record* 1986 (1), pp. 116–124.

Shanthikumar, J George and David D Yao (1987). Optimal server allocation in a system of multi-server stations. *Management Science* 33 (9), pp. 1173–1180.

Shu, Jia, Mabel C Chou, Qizhang Liu, Chung-Piaw Teo, and I-Lin Wang (2013). Models for effective deployment and redistribution of bicycles within public bicycle-sharing systems. *Operations Research* 61 (6), pp. 1346–1359.

Sixt SE (2019). Sixt Share: Maximale Flexibilität und Vielfalt – Sixt verbindet Autovermietung und Carsharing. `https://about.sixt.com/websites/sixt_cc/German/2999/news-details.html?newsID=1757303`.

Smith, Stephen L, Marco Pavone, Mac Schwager, Emilio Frazzoli, and Daniela Rus (2013). Rebalancing the rebalancers: Optimally routing vehicles and drivers in mobility-on-demand systems. *American Control Conference (ACC), 2013*. IEEE, pp. 2362–2367.

Ströhle, Philipp, Christoph M Flath, and Johannes Gärttner (2019). Leveraging customer flexibility for car-sharing fleet optimization. *Transportation Science* 53 (1), pp. 42–61.

Tang, Qinshen, Yu Zhang, and Minglong Zhou (2020). Vehicle repositioning under uncertainty. Available at SSRN 3612626.

Tijms, Henk C (2003). *A first course in stochastic models*. John Wiley and sons.

Trentini, Andrea and Federico Losacco (2017). Sampling car-sharing data to evaluate urban traffic behaviour. *International Conference Applied Computing*. IADIS, pp. 295–299.

Van Mieghem, Jan A (1998). Investment strategies for flexible resources. *Management Science* 44 (8), pp. 1071–1078.

Vasconcelos, Ana S, Luis M Martinez, Gonçalo HA Correia, Daniel C Guimarães, and Tiago L Farias (2017). Environmental and financial impacts of adopting alternative vehicle technologies and relocation strategies in station-based one-way carsharing: An application in the city of Lisbon, Portugal. *Transportation Research Part D: Transport and Environment* 57, pp. 350–362.

Volgenant, Ton and Roy Jonker (1987). On some generalizations of the travelling-salesman problem. *Journal of the Operational Research Society*, pp. 1073–1079.

Vosooghi, Reza, Jakob Puchinger, Marija Jankovic, and Göknur Sirin (2017). A critical analysis of travel demand estimation for new one-way carsharing systems. *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, pp. 199–205.

Wagner, Sebastian, Tobias Brandt, and Dirk Neumann (2016). In free float: Developing Business Analytics support for carsharing providers. *Omega* 59, pp. 4–14.

Wall Street Journal (2017). The End of Car Ownership. `https://www.wsj.com/articles/the-end-of-car-ownership-1498011001`.

Wei, Qinshuang, Ramtin Pedarsani, and Samuel Coogan (2020). Mixed Autonomy in Ride-Sharing Networks. *IEEE Transactions on Control of Network Systems*.

Weikl, Simone and Klaus Bogenberger (2013). Relocation strategies and algorithms for free-floating car sharing systems. *IEEE Intelligent Transportation Systems Magazine* 5 (4), pp. 100–111.

— (2015). A practice-ready relocation model for free-floating carsharing systems with electric vehicles–Mesoscopic approach and field trial results. *Transportation Research Part C: Emerging Technologies* 57, pp. 206–223.

Wittenbrink, Paul (2014). *Transportmanagement: Kostenoptimierung, Green Logistics und Herausforderungen an der Schnittstelle Rampe*. Springer-Verlag.

Yang, Bo, Shen Ren, Erika Fille Legara, Zengxiang Li, Edward YX Ong, Louis Lin, and Christopher Monterola (2020). Phase transition in taxi dynamics and impact of ridesharing. *Transportation Science* 54 (1), pp. 250–273.

You, Peng-Sheng and Yi-Chih Hsieh (2014). A study on the vehicle size and transfer policy for car rental problems. *Transportation Research Part E: Logistics and Transportation Review* 64, pp. 110–121.

Yu, Jiayi Joey, Christopher S Tang, Zuo-Jun Max Shen, and Xiqun Michael Chen (2020). A balancing act of regulating on-demand ride services. *Management Science* 66 (7), pp. 2975–2992.

Zhang, Rick and Marco Pavone (2016). Control of robotic mobility-on-demand systems: a queueing-theoretical perspective. *The International Journal of Robotics Research* 35 (1-3), pp. 186–203.

Zhang, Rick, Federico Rossi, and Marco Pavone (2018). Analysis, control, and evaluation of Mobility-on-Demand systems: a queueing-theoretical approach. *IEEE Transactions on Control of Network Systems* 6 (1), pp. 115–126.