

Effect of behavioral and service attributes and distribution of demand on ride-sharing efficiency and level of service

Arjan de Ruijter

May 1, 2019

TU Delft

Effect of behavioral and service attributes and distribution of demand on ride-sharing efficiency and level of service

Final report

Arjan de Ruijter / 4633954

A thesis presented in partial fulfillment of the requirements for the degree of

Master of Science
in Transport, Infrastructure & Logistics (TIL)

and to be defended publicly on Thursday May 9th, 2019



Graduation committee:
Prof.dr.ir. S.P. Hoogendoorn - Civil Engineering and Geosciences (CITG)
Dr. O. Cats - Civil Engineering and Geosciences (CITG)
Dr. J. Alonso Mora - Mechanical, Maritime and Materials Engineering (3mE)
Delft University of Technology

May 1, 2019

Preface

Dear reader,

In recent years it has become increasingly clear what the effect is of our everyday behavior on the liveability of our planet. Transportation, although essential in keeping society running, is with its large total emission partially to blame. Improvements are necessary in terms of how efficiently we are using our transport resources. At the same time, it should not be forgotten that resistance to climate-oriented policies is growing given the fear for economic consequences. The dichotomy that currently exists in the climate debate is my main motivation for this research. I hope to sharpen the expectations about the environmental benefits of one of the most promising innovations in transportation, ride-sharing, by considering specifically user acceptability.

With the rare opportunity to dive into one topic for six months, I found the main complexity of this thesis project in finding an interesting, challenging and practically relevant research topic. The very open-minded exploratory discussions I have had with Oded and Serge in the initial stages of my thesis have helped me along a great deal, and has finally led to the research presented in this report. I am very grateful for the time that both of you spent to help me with this and I feel that I could not have accomplished my (possibly excessive list of) personal goals without you. Thank you also for your continued support throughout the process, with special gratitude to Oded as my daily supervisor and, I am happy to say, soon-to-be colleague. I would also like to thank Javier for challenging me in our meetings and bringing the research to a whole new level.

I would like to end by giving a special word of appreciation to my parents Jacques and Tonny for their continued support throughout the seven years of my study, to my sisters Laura and Linda, my girlfriend Nina and all of the friends with whom I could always exchange thoughts and ideas.

Enjoy the ride!

Arjan de Ruijter
May 2019

EFFECT OF BEHAVIORAL AND SERVICE ATTRIBUTES AND DISTRIBUTION OF DEMAND ON RIDE-SHARING EFFICIENCY AND LEVEL OF SERVICE

Arjan de Ruijter
Delft University of Technology
Faculty of Civil Engineering and Geosciences
a.j.f.deruijter@tudelft.nl

Oded Cats
Delft University of Technology
Faculty of Civil Engineering and Geosciences
o.cats@tudelft.nl

Javier Alonso-Mora
Delft University of Technology
Faculty of Mechanical, Maritime and Materials Engineering
j.alonsomora@tudelft.nl

Serge Hoogendoorn
Delft University of Technology
Faculty of Civil Engineering and Geosciences
s.p.hoogendoorn@tudelft.nl

Word Count: 5384 words + 3 table(s) $\times$ 250 = 6134 words

Submission Date: April 30, 2019

ABSTRACT

Previous studies on ride-sharing potential have neglected that ride-sharing choice considers a trade-off of disbenefits with a financial reward. This study considers delays and psychological discomfort related to sharing a vehicle as the main ride-sharing costs. Next to finding how these behavioral preferences and the chosen discount structure affect the efficiency and level of service of ride-sharing, the effect of the directionality in demand is also considered. A graph-based approach is applied to allow for efficient assignment of vehicles to requests. The model is tested on an experiment representing an urban context. It was found that ride-sharing potential is strongly dependent on the willingness to share and somewhat less strongly on the delay tolerance. Implementation of a ride-sharing service is most successful when directionality in demand is low, while ride-based discounts can be effective in maximizing its societal benefits.

Keywords: Ride-sharing, Willingness to share, Delay tolerance, Demand distribution, Pricing mechanism

INTRODUCTION

Developments in communication and information technologies in recent years have led to the rise of real-time and on-demand ride-sharing platforms like UberPool and BlaBlaCar. Users of such platforms simultaneously share a vehicle, which offers opportunities for more efficient utilization of road space and consequently reduced congestion levels, improved air quality and better traffic safety (*1, 2*). From a user perspective, ride-sharing services can form an excellent alternative to expensive private and non-flexible public transit, thereby potentially solving the first- and last-mile problem that generally exists in public transit systems (*3, 4*).

Whether ride-sharing in practice can live up to these expectations is uncertain. Ride-sharing might for example be more competitive with public transit than with individual ride-hailing, leading to more rather than less vehicle kilometers on the roads, or the operation of a ride-sharing service may require excessive subsidization to cover for the discounted ride fares, and therefore not be viable. There are several other issues that can prevent a wide-scale adoption of ride-sharing services. Next to delays following from detouring to drop off other passengers, social issues related to sharing are an important deterrent for potential users. This includes a lack of privacy (*5, 6*), a feeling of dependence and the fear of having a negative social interaction with other users (*7, 8*). Social discomfort can explain why in 2016 only 20% of Uber's rides were made with its ride-sharing service UberPool (*9*).

Previous quantitative studies on societal benefits of ride-sharing showed promising results. A study by Ma et al. (*10*) for example stated that if the current fleet of taxis in New York allows for shared rides and users accept a maximum extra travel time of 5 minutes, 25% more users can be served and 13% of the total vehicle distance can be cut. Another study analyzed ride-sharing based on graph structures ('shareability graphs') and found for the same setting that 32% of the total vehicle distance of taxis can become abundant (*11*). When ride-sharing is executed with high-capacity vehicles of up to 10 seats, with less than one sixth of the size of the current taxi fleet 98% of the original requests can be served with a maximum delay of 3.5 minutes per passenger (*12*). Although the first three studies considered ride-sharing in New York, Tachet et al. (*13*) found that ride-sharing potential also exists for cities with a lower density than New York.

The overarching issue with previous ride-sharing studies is that they account for only one of three elements in the complex trade-off that users make between a delay, a discomfort and a financial reward when considering a shared ride. Each of these studies simplify ride-sharing choice by considering only a maximum allowed delay, meaning that they assume that users are principally willing to ride-share even if it gives them no benefit and a (small) delay. This study acknowledges that ride-sharing considers a trade-off of costs and benefits by accounting for a psychological sharing discomfort and discounted ride fares, with the aim of improving the validity of research on potential societal benefits of ride-sharing services. The question to be answered is how the operational efficiency of a ride-sharing service and the level of service that it offers depends on the attitudes of its potential users towards delays and co-riders. At the same time, the incorporation of the cost-benefit trade-off allows to give some first implications for the design of an effective discount structure to boost ride-sharing adoption and reduce the total vehicle distance on the road.

Another drawback of previous work in this field is that it consists mainly of case-specific analyses. It is largely unknown how the success of a ride-sharing service depends on the environment in which it is implemented, or in other words, in which markets ride-sharing has most potential. This study will focus specifically on the distribution of demand to find if efficiency gains

and level of service are higher or lower when there is more direction in demand.

This paper is structured into four parts. First, a detailed description is given of the methodology that was developed. The methodology is followed by a motivation for the design of the numerical experiment. The results for the different scenarios in the experiment are then presented, before stating the main conclusions that can be drawn in relation to the effect of users' behavioral preferences, the spatial distribution of demand and the pricing mechanism on the performance of a ride-sharing service.

METHODOLOGY

Several modelling approaches have been taken in previous research to assign requests to vehicles in order to simulate a ride-sharing service and determine the impact on the total vehicle movement and service quality. For example, one of the earlier studies considered ride-sharing assignment as a Vehicle Routing Problem (VRP) with time windows (*10*). Each incoming request was thereby individually allocated to a vehicle using a greedy algorithm. A study by Santi et al. (*11*) introduced the concept of shareability graphs (SG) to capture the shareability of two requests in a graph structure so that assignment can be performed with traditional graph-solving optimization methods. A follow-up study elaborated on the graph-based approach by the introduction of two more graph structures to allow for grouping of requests and consequently high-capacity ride-sharing (*12*). A request-group-vehicle (RGV) graph stores which requests are part of which request group and with which vehicles each request group can be served. In this way, the assignment problem is represented as an Integer Linear Problem (ILP). Finally, agent-based models (ABM) have been used to study ride-sharing before, whereby users and vehicles are modelled as agents that can interact (*14*).

Of the different approaches, the RGV-approach has most properties desired for the context of this study. First and foremost, it allows to model ride-sharing with more than two passengers per vehicle. Moreover, given that there is enough time available, the assignment yields an optimal solution. By representing the assignment in graph form, the problem's computational complexity is also minimized, which is useful as different scenarios are tested in this study. A final upside of the RGV-approach is that its explanatory power is strong with different graphs visualizing some of the main steps in the assignment procedure. As shown by Figure 1, in the RGV-approach requests are assigned to vehicles at fixed intervals, whereby each iteration consists of nine main steps.

Group-vehicle feasibility

Key in the approach is the way to determine whether a specific vehicle v in the fleet V, with passengers P_v on-board, can serve a group of pending requests Z or not. First, the complete set of routes K_v is identified with which v can potentially satisfy Z. Each route $S_v \in K_v$, defined as a sequence of stops, is then checked for feasibility based on a vehicle and user constraint. The vehicle constraint simply ensures that v cannot serve Z with route S_v if the vehicle capacity υ is exceeded between any of the stops in S_v. The user constraint is more complex and considers a trade-off between ride-sharing benefits and costs for each request r in Z or P_v. The ride-sharing benefit for r is dependent on the discount rate π_r that is applied to the ride fare c_r when a ride is shared. π_r can thereby be fixed or inverse with the level of service of the experienced ride. Disbenefits on the other hand follow from extra travel time and social discomfort when sharing a vehicle. The total disbenefit of a ride depends on how users perceive both attributes, in this study expressed by delay aversion β_r and reluctance to share γ_r. They indicate respectively what financial compensation

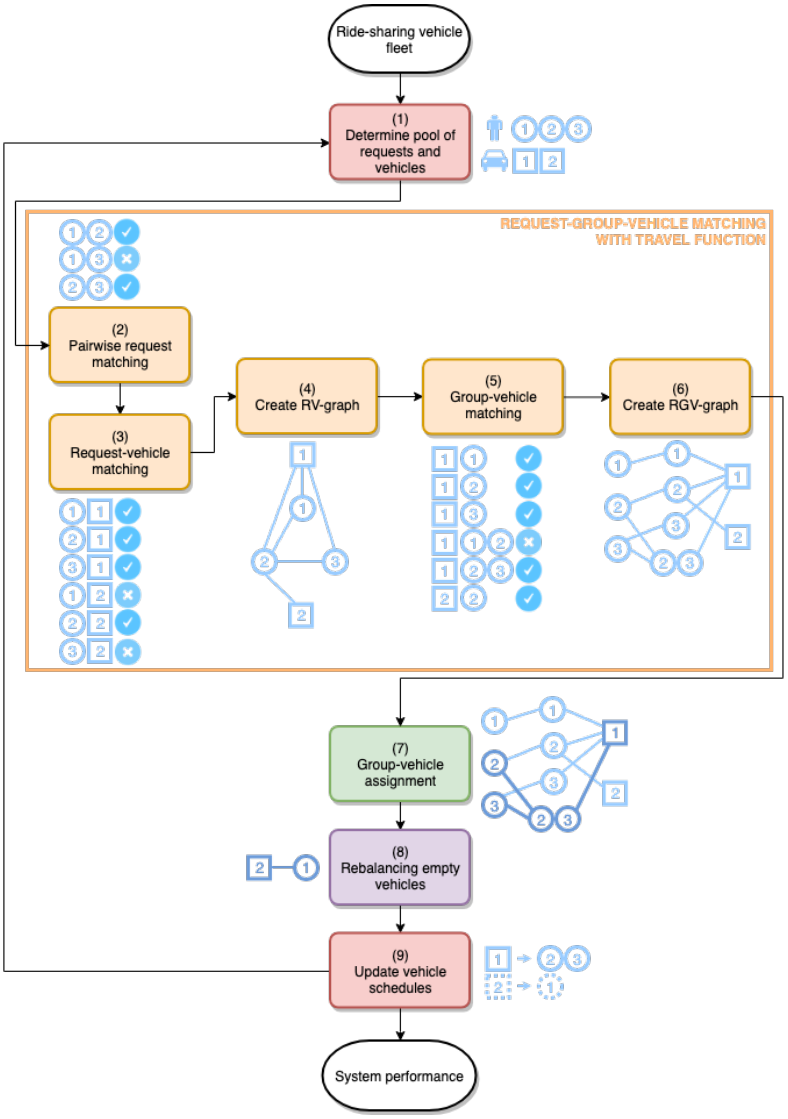


FIGURE 1 Overview of the methodology, including an example (in blue) with three requests and two (empty) vehicles

users require for an hour of delay and to share a vehicle with co-riders. A variable α_r is added to indicate how much more negatively users experience waiting time wt_r before pick-up compared to the in-vehicle delay. The waiting time of a request wt_r is thereby defined as the difference between the time of pick-up t_r^{pu} and time of request t_r^r. The total delay del_r is the difference between the actual time of drop-off t_r^d and the earliest possible time of drop-off $t_r^* = t_r^r + tt_{o_r,d_r}$ given immediate pick-up and a direct route with travel time tt_{o_r,d_r} between origin o_r and destination d_r.

A route S_v satisfies the user level of service constraint if the net benefit b_r of each request in Z and P_v, defined by $\pi_r \cdot c_r - (t_r^d - t_r^*) \cdot \beta_r - (t_r^{pu} - t_r^r) \cdot \alpha_r - \gamma_r$, is positive. If there exists at least

one feasible route $S_v \in K_v$ to serve Z, the combination of Z and v is considered feasible.

RGV-matching

The set of available requests for assignment R compiles rejected requests from the previous assignment (as long as they can still be served with a direct ride) and incoming requests since the last assignment. As mentioned earlier, the actual assignment involves the creation of a RGV-graph to find which group-vehicle combinations are feasible and how the benefit of different combinations compare. To prevent that all group-vehicle combinations have to be tested for feasibility, the process is divided into three main matching steps.

The purpose of the first of these three steps (corresponding to step 2 in Figure 1) is to find whether two requests in R can share a ride in the most optimal scenario in which there is an empty vehicle at the location of one of those requests, based on the procedure described in the previous subsection. By checking the match of all request pairs in R, the set of potentially feasible request groups G can be significantly reduced. The next step (step 3 in Figure 1) checks whether a vehicle $v \in V$ can serve a single request $r \in R$ given its location and available seats. The result of both steps can be combined into a RV-graph (step 4) with edges indicating that two requests, or a request and a vehicle, match. Each clique in the RV-graph represents a potentially feasible group-vehicle combination. In step 5 is actually considered whether a feasible route S_v exists to satisfy a group-vehicle combination within the user and capacity constraints. The RGV-graph consists of nodes representing the set of available requests R, the set of feasible request groups G and the set of vehicles V. Each edge between a request $r \in R$ and a request group $g \in G$ has a label a_{rg} indicating whether r is part of g ($a_{rg} = 1$) or not ($a_{rg} = 0$), and each edge between a request group $g \in G$ and vehicle $v \in V$ has a label b_{gv} indicating the sum of benefits of all requests in g and passengers in P_v given the optimal route S_v^*. If a group-vehicle combination gv is not feasible, it is assigned a very large penalty M^2, to ensure that this combination will not be chosen in the assignment.

Assignment

In this part of the procedure (step 7 in Figure 1), requests are assigned to vehicles based on the RGV-graph. The group-vehicle assignment is treated as an Integer Linear Problem (ILP) with binary decision variables x_{gv} indicating whether a group-vehicle combination with total benefit b_{gv} is chosen or not. The ILP is defined as follows:

$$\max \quad \sum_{g \in G} \sum_{v \in V} (b_{gv} + M \cdot \sum_{r \in R} a_{rg}) \cdot x_{gv} \tag{1}$$

$$\text{s.t.} \quad \sum_{g \in G} x_{gv} \leq 1, \ \forall v \in V \tag{2}$$

$$\sum_{g \in G} \sum_{v \in V} a_{rg} \cdot x_{gv} \leq 1, \ \forall r \in R \tag{3}$$

$$x_{gv} = [0,1], \ \forall g \in G, v \in V \tag{4}$$

The objective function (Equation 1) aims at a maximum total benefit for accepted requests and passengers, but prioritizes the acceptance of a maximum number of requests by adding a very large reward M for each request r in an assigned request group g. The total benefit of a group-vehicle combination gv therefore consists of the summed net benefit for all requests and passengers in this group plus a large reward M for each request that is part of the group. Note that the disbenefit of

infeasible group-vehicle combination was set to $-M^2$, so large that it will always compensate for the big M reward that is given to every assigned request in the objective function.

The Integer Linear Problem contains three types of constraints guaranteeing a maximum assignment of one request group g to every vehicle v (Equation 2), that each request r is not part of multiple assigned request groups in G (Equation 3), and that each decision variable is binary (Equation 4).

Rebalancing

If after the assignment both requests R^u and empty vehicles V^{np} are left unassigned, there is a spatial imbalance in supply and demand. Some unassigned vehicles in V^{np} are therefore assigned to move in the direction of unassigned requests in R^u to anticipate on new requests appearing areas that currently have undersupply (step 8 in Figure 1).

To guarantee efficient rebalancing and prevent unnecessary vehicle mileage of empty vehicles, a similar procedure as in the study of Alonso-Mora et al. (*12*) is applied. Like in the assignment step, rebalancing is solved as an ILP, with x_{rv} as decision variable to indicate whether vehicle $v \in V^{np}$ is or is not assigned to move in the direction of unassigned request $r \in R^u$. d_{rv} indicates the minimum distance needed for v to reach r. The ILP is defined as follows:

$$\min \quad \sum_{v \in V^{np}} \sum_{r \in R^u} x_{rv} \cdot d_{rv} \tag{5}$$

$$\text{s.t.} \quad \sum_{r \in R^u} x_{rv} \leq 1, \ \forall v \in V^{np} \tag{6}$$

$$\sum_{v \in V^{np}} x_{rv} \leq 1, \ \forall r \in R^u \tag{7}$$

$$\sum_{v \in V^{np}} \sum_{r \in R^u} x_{rv} = \min(|R^u|, |V^{np}|) \tag{8}$$

$$x_{rv} = [0,1], \ \forall r \in R^u, v \in V^{np} \tag{9}$$

The goal of the ILP is to minimize the total distance resulting from the rebalancing procedure (Equation 5), while ensuring a maximum number of vehicles to be assigned to rebalance (Equation 8). The latter means that either all requests in R^u or all vehicles in V^{np} are assigned, depending on whether more vehicles or requests exist in the assignment problem. The three other constraints in this ILP make sure that each vehicle can only move in the direction of one unassigned request (Equation 6), it is not allowed for multiple vehicles to rebalance to the same unassigned request (Equation 7), and that all decision variables are binary (Equation 9).

KPIs

After vehicles are assigned to pick-up requests, rebalance or remain idle, vehicle schedules are updated and the next assignment phase is prepared (step 9). To be able to comment on the performance of a ride-sharing service over a certain time period, many iterations have to be simulated.

The performance is measured with a few Key Performance Indicators (KPIs) capturing both its level of service (LoS) towards users and its operational efficiency for authorities and providers. If ride-sharing users value the same aspects as public transit users (*15–19*), the most important indicators for service quality are reliability, comfort, travel time and fare level. Ride fare in this case is not considered as a KPI as it is dependent directly on π_r and therefore considered as input. The main LoS KPIs for this study include the delay as percentage of the direct travel time (an

indicator for travel time), the acceptance rate (reliability) and the distribution of passenger time with a specific number of co-riders on-board (comfort).

An authority on the other hand will be most interested in the proportion of the vehicle distance that can be reduced with ride-sharing. A suitable KPI to express distance efficiency is the number of effective passenger kilometers served per vehicle kilometer, also known as the effective vehicle transportation distance ratio (*20*). For operators, the average vehicle occupancy in revenue service is an interesting KPI to find how much more efficient they can utilize their fleet.

Implementation

The simulation model is implemented in Python, using the open-source library Numpy (`http://www.numpy.org/index.html`) to enable efficient operations of large data structures in the model, such as creating and storing the edges of RGV-graphs when many requests and vehicles are considered. The Networkx package (`https://networkx.github.io/`) is used to compute shortest paths between two locations in the road network. The optimization problems that are part of the group-vehicle assignment and rebalancing procedure are solved with the MOSEK Optimizer API (`https://docs.mosek.com/8.1/pythonapi/index.html`)

EXPERIMENTAL DESIGN

Set-up

The experiment to which the model is applied is designed to mimic an urban context. The grid network assumed consists of 121 nodes with a distance of 500 meters between each of the nodes, leading to a maximum trip distance of 10 kilometers and a surface area of 25 km^2, comparable to the area inside the Ring Road of Amsterdam and Inner Ring of Berlin. The intermediate stop distance is relatively large, whereby implicitly is assumed that vehicles cannot stop at all road intersections and users are willing to walk to a stop. The assumed speed on the roads is slightly higher than in an average European city (*21*): 36 km/h.

The total demand for trips is set to 1,210 requests per hour, an average of 10 requests per hour per node. The way trips are distributed over the network is scenario-specific, but in any case trips with a ride distance of 2 kilometers or shorter are excluded, as such rides are uncommon (*22*) as well as undesirable in the context of a ride-sharing service. The gravity model (*23*) is applied to create the list of origin-destination pairs. Each such request r gets assigned a request time t_r^r by sampling from an exponential distribution based on the expected interval λ between two requests with a specific OD-combination, which follows, again, from the (scenario-specific) demand distribution.

The fleet of the investigated ride-sharing service consists of 150 vehicles with capacity $\upsilon = 3$, which are initially distributed over the network evenly. Assumed ride fares are based the regulated maximum taxi fares for the city of Amsterdam in 2019: a base fee of €3 and a kilometer fee of €2 (*24*).

For efficiency purposes, the total duration of the simulation is limited to two hours, with request groups being assigned to vehicles every minute (120 times in total). An additional warm-up period of 15 minutes applies to minimize the impact of the starting conditions.

Scenarios

The effect of the delay and waiting time aversion is tested by trying five different values for β_r (with α_r in this experiment set to half β_r). The delay aversion might be best compared to the

value of travel time reliability (VOR) in public transit. Since a study on travel time reliability for commuters in Barcelona (*25*) found a VOR of €34.4/h and a similar study for Australia (*26*) concluded a mean value of approximately €33/h, the base value for β_r assumed in this study is €30/h. The effect of β_r is tested by trying both values higher and lower than €30/h (as specified in scenarios 1-5 in Table 1). As little is known in literature about the reluctance to ride-share a vehicle with co-riders (γ_r), a relatively large range of values is tested in the numerical experiment: from €1 to €5 (scenarios 1 and 6-9 in Table 1). In all of the other scenarios, the median value of γ_r = €3 is assumed.

TABLE 1 Specification of all scenarios (with $\alpha_r = 0.5 \cdot \beta_r$)

#	Demand	β_r (€/h)	γ_r (€)	π_r (%)	Acronym
1	Uniform	30	3	50	U_30_3_50
2	Uniform	18	3	50	U_18_3_50
3	Uniform	24	3	50	U_24_3_50
4	Uniform	36	3	50	U_36_3_50
5	Uniform	42	3	50	U_42_3_50
6	Uniform	30	1	50	U_30_1_50
7	Uniform	30	2	50	U_30_2_50
8	Uniform	30	4	50	U_30_4_50
9	Uniform	30	5	50	U_30_5_50
10	Uniform	30	5	$50 + 7.5 \cdot n^{pax}$	U_30_5_D
11	Moderately directed	30	5	50	MD_30_5_50
12	Strongly directed	30	5	50	SD_30_5_50

Two scenarios (1 and 10 in Table 1) have been designed to test the effect of the pricing mechanism. The first of the two scenarios (or in fact all scenarios except 10) assumes a fixed 50% discount for all rides made by a ride-sharing service, independent of whether sharing actually occurred throughout the ride. In the alternative scenario, an similar discount is given to a user if he or she ends up being served privately (whereby the discount is basically a compensation for the risk of having to share), and an additional 7.5% discount is given for each co-rider n^{pax} in the vehicle during the busiest part of the ride.

The effect of directionality of demand is tested by three different scenarios. In the base scenario (1 in Table 1), demand is perfectly uniform, with equal production and attraction in each of the nodes. The other two scenarios (11 and 12 in Table 1) represent an increasingly concentrated demand pattern, with more production in the outer nodes in the network and more attraction in the central nodes, thereby mimicking a morning peak.

RESULTS

The comprehensive list of KPI values is presented in Tables 2 (level of service) and 3 (efficiency).

Behavioral preferences

The acceptance rate (Figure 2a) is found to increase when the reluctance to share γ_r of users decreases, from 25.4% when γ_r = €5 to nearly 100% when γ_r = €1. The increase is approximately linear until the great majority of requests is accepted. The average vehicle occupancy (Figure 2b)

TABLE 2 KPIs representing ride-sharing level of service for all scenarios

Scenario	Acceptance rate	Average total delay per passenger (s)	Average waiting time per passenger (s)	Average in-vehicle delay per passenger (s)	Average passenger ratio delay / direct travel time	Average number of stops per passenger	Ratio of passenger time with 0 co-riders	Ratio of passenger time with 1 co-rider	Ratio of passenger time with 2 co-riders
U_30_3_50	76.0%	126.8	92.1	34.7	29.5%	0.95	50.5%	31.8%	17.7%
U_18_3_50	88.2%	200.5	122.4	78.1	48.7%	1.56	30.9%	33.0%	36.1%
U_24_3_50	81.8%	158.5	103.6	55.0	37.3%	1.24	40.2%	34.8%	25.0%
U_36_3_50	70.1%	106.2	84.8	21.4	24.6%	0.79	58.0%	30.6%	11.4%
U_42_3_50	64.2%	93.2	77.1	16.0	21.1%	0.67	63.3%	28.9%	7.8%
U_30_1_50	98.8%	219.5	138.2	81.3	61.3%	1.83	24.9%	37.8%	37.3%
U_30_2_50	97.6%	169.6	114.0	55.6	45.3%	1.40	35.3%	37.2%	27.5%
U_30_4_50	46.4%	102.7	82.9	19.8	20.5%	0.66	63.8%	27.4%	8.8%
U_30_5_50	25.4%	88.0	74.7	13.2	15.2%	0.39	76.3%	20.8%	2.9%
U_30_3_D	82.1%	221.7	118.7	103.0	54.0%	1.97	23.1%	31.8%	45.1%
MD_30_3_50	68.0%	117.7	92.54	25.1	29.1%	0.79	55.9%	30.4%	13.7%
SD_30_3_50	62.9%	130.7	106.53	24.2	31.9%	0.80	53.4%	32.6%	14.0%

TABLE 3 KPIs representing ride-sharing efficiency for all scenarios

Scenario	Total vehicle movement distance (km)	Total vehicle transportation distance (km)	Total deadheading distance (km)	Total rebalancing distance (km)	Effective vehicle transportation distance ratio	Corrected effective vehicle transportation distance ratio	Empty vehicle rebalancing distance ratio	Average vehicle occupancy	Ratio of non-empty vehicle time with occupancy 1	Ratio of non-empty vehicle time with occupancy 2	Ratio of non-empty vehicle time with occupancy 3
U_30_3_50	6,488	5,588	900	156	1.15	1.34	0.173	1.38	69.8%	22.0%	8.2%
U_18_3_50	6,357	5,690	667	78	1.27	1.42	0.117	1.68	52.0%	27.8%	20.3%
U_24_3_50	6,456	5,657	799	133	1.21	1.38	0.166	1.52	61.0%	26.4%	12.7%
U_36_3_50	6,350	5,453	898	156	1.12	1.30	0.174	1.30	75.2%	19.8%	4.9%
U_42_3_50	6,142	5,218	925	187	1.09	1.28	0.202	1.24	78.8%	18.0%	3.2%
U_30_1_50	6,389	5,867	522	9	1.36	1.48	0.017	1.78	44.3%	33.6%	22.1%
U_30_2_50	6,853	6,133	720	56	1.26	1.41	0.077	1.59	56.0%	29.5%	14.5%
U_30_4_50	4,966	4,330	636	106	1.10	1.26	0.167	1.24	79.3%	17.0%	3.7%
U_30_5_50	3,321	2,958	364	37	1.05	1.18	0.102	1.14	87.0%	11.9%	1.1%
U_30_3_D	5,687	5,186	501	125	1.35	1.48	0.249	1.85	42.7%	29.4%	27.8%
MD_30_3_50	6,000	4,829	1,171	393	1.06	1.31	0.335	1.32	73.9%	20.1%	6.0%
SD_30_3_50	6,288	4,552	1,736	910	0.95	1.31	0.524	1.35	71.8%	21.9%	6.3%

increases more than linearly when γ_r decreases, just like the waiting time and in-vehicle delay (Figure 3a). It is found that if users are very sensitive to sharing with other passengers, rides are hardly shared (average vehicle occupancy of 1.14), meaning that in-vehicle delays are close to zero. In such a scenario, the operational efficiency in terms of the number of effective passenger

kilometers per vehicle kilometer is as low as 1.05. This ratio is found to increase approximately linearly with an increase in the sharing tolerance, as the total effective vehicle distance (due to more requests served) increases more than the total vehicle movement distance, as shown in Figure 3b. This is the result of a more efficient assignment of vehicles to requests when users are more flexible. Also, deadheading is found to be relatively uncommon when users' sharing tolerance is high, as new requests can be picked-up by vehicles on their way to dropping off other passengers. If γ_r = €1 for example, the total effective passenger kilometers that are served with one vehicle kilometer rises to 1.36.

When considering the effect of delay aversion β_r instead of reluctance to share γ_r, similar results are found, although less pronounced. For example, the acceptance rate does not exceed 90% in any of the scenarios. Apparently, the level of service and operational efficiency are more sensitive to the tested values of γ_r than to those of β_r.

FIGURE 2 Effect of reluctance to share γ_r on (a) acceptance rate, and (b) vehicle occupancy and number of intermediate stops

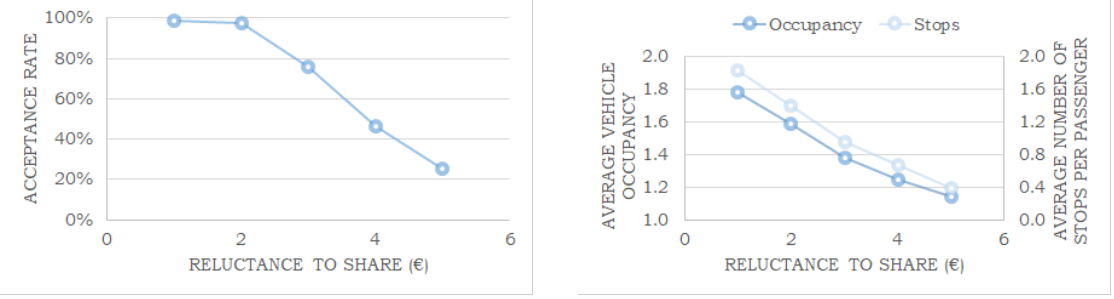
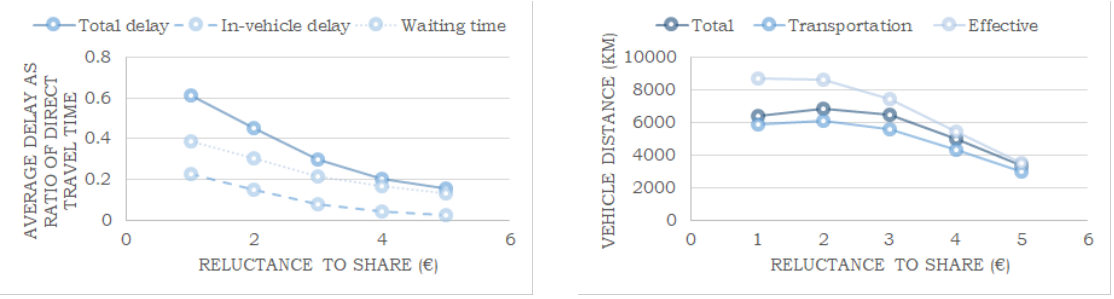


FIGURE 3 Effect of reluctance to share γ_r on (a) average total delay, in-vehicle delay and waiting time per passenger, and (b) total vehicle movement, total transportation distance and effective transportation distance



Discount mechanism

As expected, when users receive an additional 7.5% discount for every co-rider they share the busiest part of their ride with (U_30_3_D), the average vehicle occupancy increases quite dramatically (from 1.38 to 1.85, as shown by Figure 4a), and a similar increase is found in the relative time passengers spent in a vehicle at full capacity (from 17,7% to 45.1%). By utilizing the available

vehicle capacity more efficiently, the acceptance rate (also Figure 4a) increases from 76.0% to 82.1%, although at the cost of a higher average delay (Figure 4b). A higher vehicle occupancy will burden passengers with larger detours and consequently an in-vehicle delay more than three times as high as when no additional discount is offered (25.1% versus 8.1% of the direct travel time). Also, the waiting time of requests is marginally larger in the scenario with an occupancy-dependent discount, with pick-ups being complicated by the fact that many vehicles are driving around fully occupied.

The effective passenger kilometers per ride-sharing kilometer in the experiment increases from 1.15 to 1.35. In combination with a higher acceptance rate, a relatively large distance savings (Figure 5) can be achieved with an additional occupancy-dependent discount. The distance that a ride-sharing service can save in this way originates not only from more efficient transportation of requests (the transportation distance drops from 5,588 to 4,829 kilometers) but also from a reduction of the deadheading distance to access new requests (from 900 to 501 kilometers), as requests are being picked-up by non-empty vehicles on their way to dropping off other passengers.

FIGURE 4 Effect of discount structure on (a) acceptance rate and average vehicle occupancy, and (b) average delay per passenger

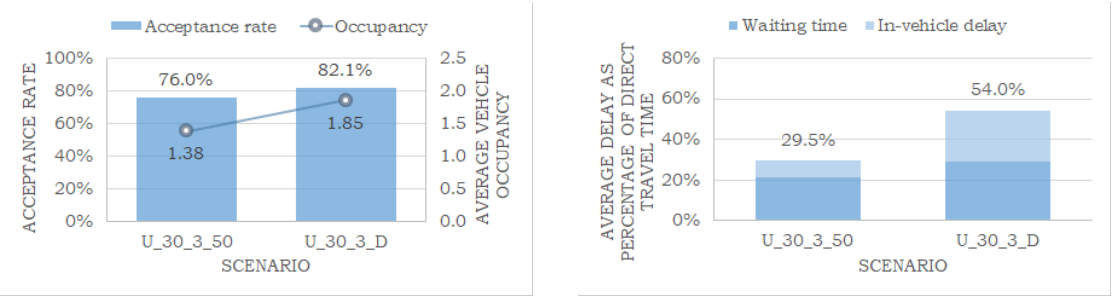
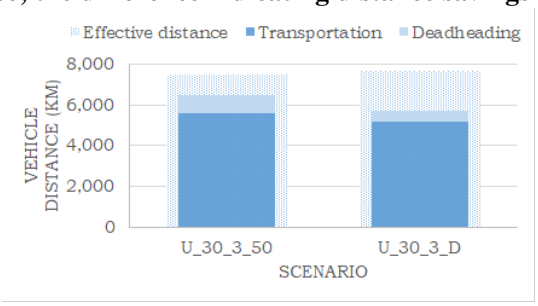


FIGURE 5 Effect of discount structure on total vehicle movement distance and effective transportation distance; the difference indicating distance savings



Demand distribution

More direction in demand led to more requests being rejected by the assumed ride-sharing service (37.1% when demand is strongly directed versus 24.0% when demand is perfectly uniform, Figure 6a). If demand is perfectly uniform, the average vehicle occupancy of vehicles in revenue mode (Figure 6a) is 1.38 and the average passenger delay (Figure 6b) 29.5% of the direct travel time.

The drop in the number of accepted requests when there is a moderate level of direction in demand leads to a drop of the vehicle occupancy (1.32) and average delay (29.1% of direct travel time). If the level of direction increases further however, the average delay starts to increase again, to 31.9% of the direct travel time in a scenario where demand is strongly directed. With a larger spatial inequality in pick-ups and drop-offs, average waiting times are low in areas with more attraction than production (Figure 7a). Since only the minority of requests originates here, the average waiting time is mainly determined by requests in areas with overproduction, with a waiting time nearly twice as high (27.1% versus 14.1%).

Deadheading to solve inequality in supply and demand is responsible for 27.6% of the vehicle kilometers (Figure 7b) in a scenario in which demand is strongly directed, while responsible for only 13.9% of the mileage when demand is uniform. The effective passenger kilometers per ride-sharing vehicle kilometer in respective scenarios are 0.95 and 1.15, meaning that when directionality in demand is large the total vehicle distance can be larger than the effective transportation distance.

FIGURE 6 Effect of directionality in demand on (a) acceptance rate and average vehicle occupancy, and (b) average delay per passenger

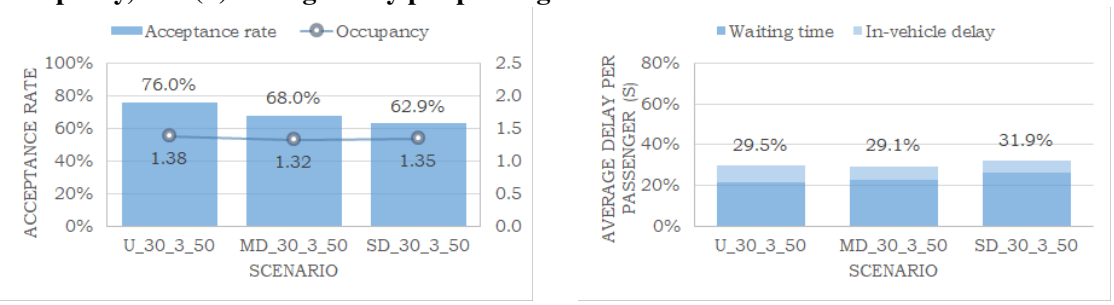
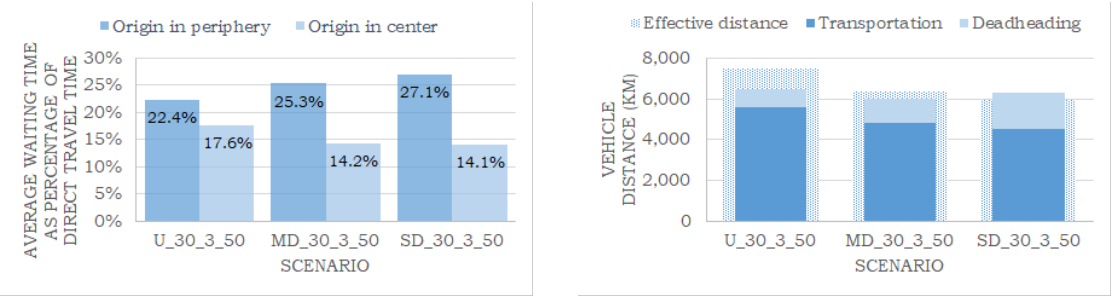


FIGURE 7 Effect of directionality in demand on (a) relative waiting time per origin location, and (b) total vehicle movement distance and effective transportation distance

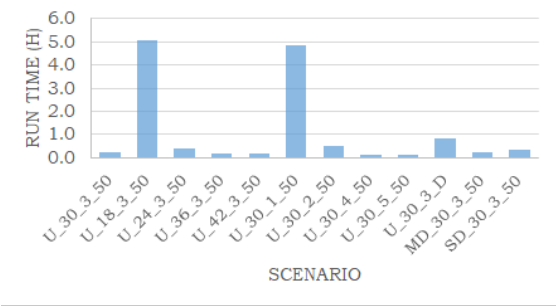


Computational complexity

Figure 8 shows the model run time with a single-core 2.30GHz processor for each scenario. The scenarios with the lowest delay aversion β_r and lowest reluctance to share γ_r are, by far, most computationally complex. This concerns scenarios where requests can be satisfied with large detours, meaning that large request groups are potentially feasible. It requires significant computational

time to test those as the set of possible routes to satisfy such groups is significantly (i.e. more than exponentially) larger than for small request groups. An occupancy-dependent additional discount (U_30_3_D) is also favorable for the feasibility of large request groups and therefore the computational complexity of this scenario is also relatively large.

FIGURE 8 Model run time of all scenarios with a single-core 2.30GHz processor



CONCLUSIONS

Based on the numerical experiment, both efficiency and level of service of a ride-sharing service are found to be strongly dependent on users' delay tolerance and especially the willingness to share a vehicle with others. The acceptance rate of a ride-sharing service varies between 25.4% and 98.8%, the average delay between 15.2% and 61.3% of the direct travel time, and the number of effective passenger kilometers per vehicle kilometer between 0.95 and 1.36, depending on the sharing and delay attitude of users. When an additional discount of 7.5% is offered on top of the standard 50% discount for every co-rider that a user shares the busiest part of its ride with, more requests can be accepted (82.1%) and the effective transportation distance ratio rises from 1.15 to 1.35. In return for a discount, users are on average burdened with an extra travel time of 24.5% of the direct travel time. Finally, it was found that the demand distribution can significantly influence ride-sharing efficiency and level of service. The larger the directionality of demand in the experiment, the fewer requests are accepted (from 76.0% when directionality is zero to 62.9% when the directionality is relatively large) and the lower the effective vehicle transportation distance ratio, even dropping below 1 when concentration in demand is strong. In such a situation, the level of service of ride-sharing users is also found to be low with long waiting times before pick-up and consequently a relatively large average delay of 31.9% of the direct travel time. A moderate level of directionality in demand on the other hand leads to a shorter average delay as the reduction in the in-vehicle delay resulting from a lower vehicle occupancy is then found to outweigh the additional waiting time.

This work is the first study to consider ride-sharing potential based on the trade-off that users are faced with when presented with the option of ride-sharing. Previous studies, such as that of Santi et al. (*11*) and Alonso-Mora et al. (*12*), assumed that all users are potentially willing to ride-share as long as their waiting time and total delay do not exceed a certain threshold. This is not very realistic as taxi users have no reason to share their ride (and accept a delay) when they do not get a benefit in return. Therefore, in this study the choice to ride-share considers the trade-off of ride-sharing disbenefits with a discount on the ride fare. The assumption is that users will only switch to a ride-sharing service if such a choice gives them a net positive utility over a conventional

taxi or ride-hailing ride. Also, this work accounts for the fact that negative utility can follow from social discomfort when sharing a vehicle with strangers, which in literature has been found to be one of the main barriers for a successful implementation of ride-sharing services. This research shows, in fact, that users' willingness to share might have a larger impact on ride-sharing potential than users' attitudes towards delays.

Compared to the findings in previous research, there is a clear indication that the expectation for potential distance savings should be negatively adjusted given that a trade-off of costs and benefits exists. In the most positive scenario in this research, which assumes that users have a high willingness to share, 27% of the vehicle kilometers can be removed, as opposed to for example 40% found in the study for New York by Santi et al. (*11*). In some of the scenarios in this work it was even found that a ride-sharing system can produce more vehicle kilometers than a ride-hailing system would need.

This study also shows that the potential of a ride-sharing system can be greatly dependent on external variables such as the distribution of demand. Demand in reality is likely to be at least somewhat concentrated due to spatial clustering of activities like work, residency and shopping. This study shows for example that, when most requests are going into the direction of the center of the grid, which is common in a morning peak, the performance of a ride-sharing system is relatively poor, both in terms of the level of service and efficiency. Most likely, this is also the case for (evening) peaks in which demand is directed outward.

Furthermore, this study has shown that the design of a ride-sharing service, such as the pricing structure, can potentially significantly affect the expected societal benefits and service quality. A small additional discount of 7.5% per co-rider with whom a user shares their ride at maximum occupancy, on top of the standard 50% discount assumed throughout this research, can more than double the total reduction in vehicle kilometers.

A serious concern for ride-sharing services can be its profitability. Given that there is a 50% discount for all users in this experiment, and not all requests can be accepted, a drastic reduction of the revenue is expected. The improvement in operational efficiency and the reduction of costs related to it, is limited to approximately 25% for the most positive scenario. Subsidies from authorities might therefore be necessary for a ride-sharing service to be financially sustainable.

Future research in this field can focus on a few aspects. Firstly, it would be interesting to apply the model to a real-world case. Also, behavioral, demand and pricing attributes of which the effect was investigated individually in this study can be combined to see how they interact and whether the whole is greater, equal or smaller than the sum of its parts. Moreover, improvements can be made to the assumption that the occupancy-dependent discount is dependent only on the number of co-riders during the busiest part of the ride. It would be recommended in future research to use a weighted measure for the sharing experience over time as experienced level of service is likely to be determined by the complete ride rather than by one segment. There are other potentially relevant attributes than the ones considered in this work of which the effect can be investigated. Future research can focus for example on the effect of fleet properties (capacity and fleet size), test a more complex pricing mechanism than the one assumed here, or investigate external variables like the number of stop locations in the network. If vehicle capacity is considered, it is relevant also to find out how the willingness to share depends on the number of co-riders in the vehicle. The developed model can also be used to investigate ride-sharing with a fleet of autonomous vehicles, for which it would be especially relevant to find how willingness to share depends on the absence of a driver.

Another main research opportunity lies in the assumption of most studies on ride-sharing potential, including this one, that traffic speeds are static and not effected by congestion. It would therefore be interesting to consider dynamic traffic speeds, but also the availability of parking spots. In terms of the assignment procedure, improvements can be made to prevent unwanted behaviour, such as route reversals and users staying in a vehicle after reaching their drop-off location in order to collect a larger discount (infrequently happening in the scenario with an occupancy-dependent discount). The current objective of the assignment can be rethought to exclude discounts, as they are costs from an operator's perspective. Alternative objectives could for example be a minimization of the sum of disbenefits or a minimum total vehicle distance. It might be interesting also to find whether efficiency of a ride-sharing service can be improved by rejecting requests that negatively affect ride-sharing performance on a system level, such as requests that are destined for a location far away from where new demand is expected.

The validity of ride-sharing studies can further be improved by considering ride-sharing in the context of a mode choice that includes private, shared and public transport. It would be interesting also to gain more knowledge on how ride-sharing can complement public transit by offering cheap and efficient first and last mile transport to and from transit stops. In the broader scope it is also essential to account for 'rebound' effects that can take place, whereby a reduction in vehicle kilometers on the road as a result of more shared rides leads to lower travel times and ultimately an increase in the relative attractiveness of the car compared to public transit (*27*). Finally, future research can address the equity of users in a ride-sharing setting, as users in remote areas are potentially more likely to be rejected by a ride-sharing service, which negatively affects their accessibility compared to other users.

REFERENCES

1. Wang, X., M. Dessouky, and F. Ordonez, A pickup and delivery problem for ridesharing considering congestion. *Transportation letters*, Vol. 8, No. 5, 2016, pp. 259–269.
2. Cici, B., A. Markopoulou, E. Frías-Martínez, and N. Laoutaris, Quantifying the potential of ride-sharing using call description records. In *Proceedings of the 14th Workshop on Mobile Computing Systems and Applications*, ACM, 2013, p. 17.
3. Bian, Z. and X. Liu, A Detour-Based Pricing Mechanism for First-Mile Ridesharing in Connection With Rail Public Transit. In *2018 Joint Rail Conference*, American Society of Mechanical Engineers, 2018, pp. V001T04A003–V001T04A003.
4. Cheng, S.-F., D. T. Nguyen, and H. C. Lau, A mechanism for organizing last-mile service using non-dedicated fleet. In *Proceedings of the The 2012 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology*, IEEE Computer Society, 2012, pp. 85–89.
5. Dueker, K. J., B. O. Bair, and I. P. Levin, Ride sharing: Psychological factors. *ASCE Transp Eng J*, Vol. 103, No. 6, 1977, pp. 685–692.
6. Teal, R., Carpooling: Who, how and why. *Transportation Research Part A: General*, Vol. 21, No. 3, 1987, pp. 203–214.
7. Morales Sarriera, J., G. Escovar ÃĄlvarez, K. Blynn, A. Alesbury, T. Scully, and J. Zhao, To Share or Not To Share: Investigating the Social Aspects of Dynamic Ridesharing. *Transportation Research Record: Journal of the Transportation Research Board*, Vol. 2605, 2017, pp. 109–117.
8. Correia, G. and J. Viegas, Carpooling and carpool clubs: Clarifying concepts and assessing value enhancement possibilities through a Stated Preference web survey in Lisbon, Portugal. *Transportation Research Part A: Policy and Practice*, Vol. 45, No. 2, 2011, pp. 81–90.
9. Lunden, I., Uber says that 20% of its rides globally are now on UberPool. *Tech Crunch*, 2016.
10. Ma, S., Y. Zheng, and O. Wolfson, T-share: A large-scale dynamic taxi ridesharing service. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*, IEEE, 2013, pp. 410–421.
11. Santi, P., G. Resta, M. Szell, S. Sobolevsky, S. Strogatz, and C. Ratti, Quantifying the benefits of vehicle pooling with shareability networks. *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 111, No. 37, 2014, pp. 13290–13294.
12. Alonso-Mora, J., S. Samaranayake, A. Wallar, E. Frazzoli, and D. Rus, On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 114, No. 3, 2017, pp. 462–467.
13. Tachet, R., O. Sagarra, P. Santi, G. Resta, M. Szell, S. Strogatz, and C. Ratti, Scaling law of urban ride sharing. *Scientific Reports*, Vol. 7, 2017.
14. Fiedler, D., M. Čertickỳ, J. Alonso-Mora, and M. Čáp, The impact of ridesharing in mobility-on-demand systems: Simulation case study in Prague. In *IEEE International Conference on Intelligent Transportation Systems*, 2018.
15. Bates, J., J. Polak, P. Jones, and A. Cook, The valuation of reliability for personal travel. *Transportation Research Part E: Logistics and Transportation Review*, Vol. 37, No. 2-3, 2001, pp. 191–229.

16. Edvardsson, B., Causes of customer dissatisfaction-studies of public transport by the critical-incident method. *Managing Service Quality: An International Journal*, Vol. 8, No. 3, 1998, pp. 189–197.
17. Hensher, D. A., P. Stopher, and P. Bullock, Service quality—-developing a service quality index in the provision of commercial bus contracts. *Transportation Research Part A: Policy and Practice*, Vol. 37, No. 6, 2003, pp. 499–517.
18. König, A. and K. W. Axhausen, The reliability of the transportation system and its influence on the choice behaviour. In *Swiss transport research conference (strc), monte verità/ascona*, 2002.
19. Friman, M. and T. Gärling, Frequency of negative critical incidents and satisfaction with public transport services. II. *Journal of Retailing and Consumer Services*, Vol. 8, No. 2, 2001, pp. 105–114.
20. Ehsani, S., R. Guglielmetti, M. Laumanns, I. Markov, and R. de Souza, *Simulation-based design and analysis of on-demand mobility services*. Bestmile, 2018.
21. Kfzteile24, *Best and Worst Cities to Drive 2017*, 2019.
22. Liu, T., P. Krishnakumari, and O. Cats, Exploring Demand Patterns of a Ride-Sourcing Service using Spatial and Temporal Clustering, 2019, unpublished.
23. Erlander, S. and N. F. Stewart, *The gravity model in transportation analysis: theory and extensions*, Vol. 3. Vsp, 1990.
24. Gemeente Amsterdam, *Wat kost een rit in een taxi?*, 2019.
25. Asensio, J. and A. Matas, CommutersâĂŹ valuation of travel time variability. *Transportation Research Part E: Logistics and Transportation Review*, Vol. 44, No. 6, 2008, pp. 1074–1085.
26. Li, Z., D. A. Hensher, and J. M. Rose, Willingness to pay for travel time reliability in passenger transport: A review and some new empirical evidence. *Transportation research part E: logistics and transportation review*, Vol. 46, No. 3, 2010, pp. 384–403.
27. Yin, B., L. Liu, N. Coulombel, and V. Viguié, Appraising the environmental benefits of ride-sharing: The Paris region case study. *Journal of Cleaner Production*, Vol. 177, 2018, pp. 888–898.

Contents

1 Introduction **1**
- 1.1 Background . . . 1
- 1.2 Related work . . . 2
- 1.3 Problem context . . . 3
- 1.4 Research gap . . . 4
- 1.5 Relevance of research . . . 4
- 1.6 Research questions . . . 5
- 1.7 Outline . . . 7

2 Methodology **8**
- 2.1 Selection of modelling approach . . . 8
- 2.2 Assignment procedure . . . 10
- 2.3 Mathematical problem definition . . . 13
  - 2.3.1 Parameters . . . 13
  - 2.3.2 Sets . . . 13
  - 2.3.3 Demand and network characteristics . . . 14
  - 2.3.4 Sequences . . . 14
  - 2.3.5 Variables . . . 15
- 2.4 Main steps . . . 15
  - 2.4.1 RGV-matching with travel function . . . 16
  - 2.4.2 Group-vehicle assignment . . . 18
  - 2.4.3 Rebalancing empty vehicles . . . 20
  - 2.4.4 Update vehicle status and request pool . . . 21
- 2.5 Key performance indicators . . . 23
  - 2.5.1 Level of service . . . 23
  - 2.5.2 Efficiency . . . 23
- 2.6 Implementation . . . 24

3 Experimental design **26**
- 3.1 Set-up . . . 26
- 3.2 Scenario design . . . 28
  - 3.2.1 Base scenario . . . 28
  - 3.2.2 Homogeneous preferences . . . 30
  - 3.2.3 Heterogeneous preferences . . . 30
  - 3.2.4 Occupancy-dependent reluctance to share . . . 30
  - 3.2.5 Occupancy-dependent discount . . . 30
  - 3.2.6 Demand distribution . . . 31

4 Results **32**
4.1 Homogeneous preferences . . . 32
4.2 Heterogeneous preferences . . . 38
4.3 Occupancy-dependent reluctance to share . . . 41
4.4 Occupancy-dependent discount . . . 43
4.5 Demand distribution . . . 46
4.6 Correspondence of objectives . . . 49
4.7 Computational complexity . . . 51

5 Conclusions & recommendations **53**
5.1 Key findings . . . 53
5.2 Implications . . . 58
5.2.1 Vehicle mileage . . . 59
5.2.2 Profitability . . . 60
5.2.3 Equity . . . 60
5.3 Practical recommendations . . . 60
5.4 Future research . . . 61
5.4.1 Application . . . 61
5.4.2 Behavioral attributes . . . 61
5.4.3 Operations . . . 62
5.4.4 Network . . . 63
5.4.5 Assignment procedure . . . 63
5.4.6 Public transit interaction . . . 64
5.4.7 Spatial equity . . . 64

References **65**

A List of results **68**

B Demand distribution **71**

List of Figures

1.1 Efficiency (red), level of service (blue) and profitability (green) of ride-sharing platforms based on individual, collective and service attributes . . . 6

2.1 Overview of the methodology, including an example (in blue) of an assignment phase with three requests and two vehicles (without passengers) . . . 11

3.1 Reluctance to share γ_z as function of the number of co-riders 31
3.2 Discount percentage π_z as function of the number of co-riders 31

4.1 Acceptance rate as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z 33
4.2 Average delay as ratio of the direct travel time, as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z 34
4.3 Ratio of passenger time with specific number of co-riders on-board as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z 34
4.4 Average vehicle occupancy and number of intermediate stops per passenger as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z 34
4.5 Distribution of vehicle time in each state as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z 35
4.6 Total vehicle movement, transportation movement and effective transportation movement as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z 36
4.7 Effective vehicle transportation distance as ratio of the total vehicle distance, including and excluding deadheading, as a function of (a) delay aversion β_z and (b) reluctance to share γ_z 36
4.8 Total deadheading distance as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z 37
4.9 Effect of heterogeneity on the acceptance rate 38
4.10 Reluctance to share of accepted users versus reluctance to share of all requests in *U_30_H_50* 38
4.11 Delay aversion of accepted users versus the delay aversion of all requests in *U_H_3_50* 38
4.12 Effect of heterogeneity on average delay per passenger 39
4.13 Effect of heterogeneity on the distribution of delays 39
4.14 Effect of heterogeneity on average vehicle occupancy in revenue mode and average number of intermediate stops 39
4.15 Effect of heterogeneity on passenger time with different number of co-riders 39
4.16 Delay per passenger as function of its delay aversion β_z 40
4.17 Delay per passenger as function of its reluctance to share γ_z 40

4.18 Effect of heterogeneity on relative distance savings, including and excluding deadheading . . . 40
4.19 Effect of heterogeneity on the total deadheading distance . . . 40
4.20 Acceptance rate in scenarios *U_30_3_50* and *U_30_D_50* . . . 41
4.21 User distribution of maximum number of co-riders in ride in scenarios *U_30_3_50* and *U_30_D_50* . . . 41
4.22 Average vehicle occupancy and number of intermediate stops in scenarios *U_30_3_50* and *U_30_D_50* . . . 42
4.23 Average total delay per passenger in scenarios *U_30_3_50* and *U_30_D_50* . 42
4.24 Proportion of vehicle time in different states in scenarios *U_30_3_50* and *U_30_D_50* . . . 42
4.25 Rebalancing distance as ratio of the vehicle transportation distance in scenarios *U_30_3_50* and *U_30_D_50* . . . 43
4.26 Effective vehicle transportation distance ratio in scenarios *U_30_3_50* and *U_30_D_50* . . . 43
4.27 Acceptance rate in scenarios *U_30_3_50* and *U_30_3_D* . . . 44
4.28 Distribution of maximum number of co-riders during ride in scenarios *U_30_3_50* and *U_30_3_D* . . . 44
4.29 Average occupancy of vehicles in revenue mode and average number of intermediate stops in scenarios *U_30_3_50* and *U_30_3_D* . . . 44
4.30 Distribution of passenger time with specific number of co-riders in scenarios *U_30_3_50* and *U_30_3_D* . . . 45
4.31 Average total delay per passenger in scenarios *U_30_3_50* and *U_30_3_D* . . 45
4.32 Cumulative percentage of passengers with certain delay, based on their discount (*U_30_3_D*) . . . 45
4.33 Percentage of vehicle time in different states in scenarios *U_30_3_50* and *U_30_3_D* . . . 45
4.34 Total vehicle movement distance in scenarios *U_30_3_50* and *U_30_3_D* . . . 46
4.35 Effective vehicle transportation ratio in scenarios *U_30_3_50* and *U_30_3_D*, including and excluding deadheading . . . 46
4.36 Effect of demand distribution on acceptance rate . . . 47
4.37 The share of requests originating in outer rings for accepted versus total demand, given different distributions of demand . . . 47
4.38 Effect of demand distribution on average total delay per passenger . . . 47
4.39 Effect of demand distribution on the average vehicle occupancy in revenue mode and average shortest path trip length between a request's origin and destination . . . 47
4.40 Effect of demand distribution on distribution of passenger time with specific number of co-riders . . . 48
4.41 Percentage of requests in *SD_30_3_50* with no in-vehicle delay given the trip direction, "O" meaning location is in one of the three outer rings, "I" in any of the other rings . . . 48
4.42 Total passenger delay in *SD_30_3_50* given trip direction, "O" meaning location is in one of the three outer rings, "I" in any of the other rings . . . 48
4.43 Relative user waiting time given request location in scenarios *U_30_3_50*, *MD_30_3_50* and *SD_30_3_50* . . . 48
4.44 Effect of demand distribution on effective vehicle transportation distance ratio . . . 49
4.45 Effect of the demand distribution on rebalancing distance ratio . . . 49
4.46 Relation between the acceptance rate and average user delay for all scenarios, including the Pareto frontier . . . 50

4.47 Relation between the acceptance rate and total distance savings for all scenarios 50
4.48 Relation between the average user delay and total distance savings for all scenarios, including the Pareto frontier 51
4.49 Model run time of all scenarios with a single-core 2.30GHz processor 51

5.1 Most relevant relations found between individual, collective and service attributes and ride-sharing level of service and efficiency indicators (green = positive relation, red = negative relation, numbers correspond to the research question with which a relation was found) 58

B.1 Production and attraction per hour per node in different scenarios 71

List of Tables

1.1 Overview of studies on ride-sharing potential and how this work fits in . . . 5

2.1 Suitability of each approach to fulfill a certain desired property (from impossible to implement (N/A) to already capable (+++) 9
2.2 Input for travel function in each of the matching steps 16
2.3 Summary of the configuration of Algorithm 1 for different sub research questions 18
2.4 KPIs representing Level of Service of a dynamic ride-sharing system 24
2.5 KPIs representing the efficiency of a dynamic ride-sharing system 25

3.1 Parameter values in numerical experiment 26
3.2 Specification of all fifteen scenarios (with α_z always half of β_z) and the sub research questions they help to answer 28

A.1 KPIs representing ride-sharing level of service for all scenarios 69
A.2 KPIs representing ride-sharing efficiency and computational complexity for all scenarios 70

Chapter 1

Introduction

In this chapter, the scope of the research is introduced. First, background information is provided on ride-sharing and specifically why it has gained new attention in recent years (Section 1.1). Section 1.2 summarizes the most relevant previous studies in this field, with Section 1.3 commenting on the main drawbacks or missing elements of those studies. Based on this, a research gap is identified that currently exists in the field of ride-sharing (Section 1.4). In Section 1.5, the added practical value of a contribution to closing this knowledge gap, and thereby of this specific study, is motivated. The research gap is translated into a main research question that is divided into six sub questions to structure the approach, as presented in Section 1.6. Finally, an outline on the structure of the report is given in Section 1.7.

1.1 Background

In an era in which sustainability is becoming a more and more predominant issue, the sharing economy offers new opportunities to limit the spoilage of resources in various sectors. In the field of mobility, the principle of the sharing economy is found in recently emerged car-sharing services like Zipcar, that improve the utilization of a car over time, and ride-sharing services, like BlaBlaCar and UberPool.

Ride-sharing, as opposed to single-occupant car-sharing or traditional ride-hailing (like Uber X), can lead to a more efficient utilization of road space as it allows road users to occupy vacant seats in a vehicle traveling between an origin and destination that matches theirs, while allowing for a small detour. In this way, a single vehicle can simultaneously serve multiple passengers, thus reducing the total vehicle kilometers needed to serve each of the requests. On a city scale, this can give societal benefits including reduced congestion levels, better air quality and improved traffic safety (Wang, Dessouky, & Ordonez, 2016; Cici, Markopoulou, Frías-Martínez, & Laoutaris, 2013). From a user perspective, ride-sharing services can bridge the gap between private and public transit, by offering users a more flexible alternative to conventional public transport that is cheaper than a private taxi. Moreover, ride-sharing can potentially solve the first- and last-mile problem that exists for public transportation users in low-density areas (Bian & Liu, 2018; Cheng, Nguyen, & Lau, 2012). Given that ride-sharing can also be an alternative to public transport, there is no certainty that congestion levels will actually drop by the introduction of ride-sharing services, as is often assumed. The opposite might even be true if most ride-sharing users originally traveled by public transport and not in single-occupant vehicles.

The recent growth of online ride-sharing services like BlaBlaCar can be considered a reinvention of the much older concept of carpooling. In the 1970s, the oil crisis led commuters to share their vehicle with colleagues in order to save fuel costs (Handke & Jonuschat, 2012). Although the end of the oil crisis did not mark the end of carpooling,

its usage dropped over time as psychological barriers related to sharing a vehicle with strangers and constraints related to time and flexibility started to outweigh the reduced economic incentives (Correia & Viegas, 2011).

Renewed online sharing platforms are found to be able to address some of these issues. In the last few years, developments in communication and information technologies in combination with the rise of the smartphone have allowed ride-sharing to take place in real-time and on-demand. This has greatly improved the flexibility of ride-sharing by giving users the possibility to alter their plans at any moment and any place. Sharing platforms also benefit from economies of scale, resulting in more efficient trip matching and consequently reduced time losses when sharing a ride. With the improved level of service, ride-sharing may have become a more attractive alternative to private and public transport than ever, opening up new opportunities for reducing road congestion in urban environments.

1.2 Related work

Along with the emergence of dynamic ride-sharing platforms, ride-sharing and its potential for more efficient urban mobility has gained attention in research. The first study to quantify ride-sharing benefits on a larger scale shows that in a city with 33,000 taxi's, 13% of the travel distance can be saved by real-time ride-sharing while serving 25% more users if those users allow for a pick-up and drop-off time window of 5 minutes (Ma, Zheng, & Wolfson, 2013). This study is likely to underestimate the potential as it uses a greedy algorithm to assign requests to vehicles. Although sequential assignment of requests based on incoming order grants direct feedback to the user and is computationally more efficient, on a system level it does not necessarily lead to a minimum number of vehicle kilometers.

This issue is addressed in a study that identifies the potential of taxi-sharing in Manhattan (Santi et al., 2014). It represents the ride-sharing assignment problem as a graph problem, where graphs contain information on whether individual trip requests can be shared. Passenger inconveniences are incorporated in the framework by allowing a maximum detour (in minutes) for any shared request. Requests are collected within a specific time window and collectively assigned so that the total travel time in the system is minimized. When users allow for a 5 minute delay, the study shows that 32% of the total vehicle kilometers can be saved by introduction of a ride-sharing scheme that gives feedback to users within a minute. In this approach, vehicles with spare capacity will not accept any new requests after starting their trip.

Another study on the same data set shows that potential ride-sharing benefits in New York can be further improved when a ride-sharing scheme includes high-capacity vehicles with 4 or even 10 available seats per vehicle, while at the same time enabling dynamic rerouting based on new trip requests (Alonso-Mora, Samaranayake, Wallar, Frazzoli, & Rus, 2017). With a minimal extra waiting time of 2.8 minutes and a trip delay of 3.5 minutes, 98% of the demand can be fulfilled with 2,000 vehicles of capacity 10 or 3,000 vehicles of capacity 4 as opposed to the 13.500 taxis in the base scenario without ride-sharing. Such a high-capacity ride-sharing scheme thus offers efficiency that resembles that of public transport, while only causing small travel delays compared to direct private alternatives.

The high potential for ride-sharing is not specific for New York City (Tachet et al., 2017), a city with an extraordinary dense traffic demand. Research that considers the shareability of taxi trips in New York, Singapore, San Francisco and Vienna, reveals that a scaling law exists that can be used to predict the potential of ride-sharing in a city based on accepted travel delay, the number of trip requests per hour and two city-specific parameters, namely the average daily traffic speed and surface area. As the shareability of

trips as a function of the trip density follows a saturation curve that saturates relatively quickly, also in cities with densities much lower than New York there is great potential for sharing (Santi et al., 2014).

The potential for sharing rides is not limited to taxi services. As in the traditional concept of carpooling, road users can also share rides privately with road users going in a similar direction. As this does not involve a driver occupying a seat without fulfilling any travel demand, it can potentially offer greater societal benefits than presented in previous taxi-sharing studies. However, the downside of traditional carpooling, or more specifically private car ownership, is a lack of efficiency over time, resulting in large parking space and fleet size requirements (Fagnant & Kockelman, 2014). In a study for the city of Prague, a Mobility-on-Demand (MoD) service is proposed that allows for simultaneous car- and ride-sharing (Fiedler, Čertickỳ, Alonso-Mora, & Čáp, 2018). It shows that such a hybrid system can potentially reduce the vehicle kilometers resulting from private car usage in the city by 40%, while at the same time reducing the pressure on parking areas throughout the city.

Table 1.1 summarizes previously mentioned studies based on the type of ride-sharing they study. This includes vehicle type, whether assignment takes place real-time, whether assigned requests can be reassigned before being picked-up and whether empty vehicles can rebalance to anticipate on new demand. It also indicates the capacity of vehicles in the ride-sharing system, the properties of user preferences and the magnitude of the ride-sharing system studied.

1.3 Problem context

In a few major cities, extensive ride-sharing systems are already operating, including Uber's UberPool. It appeared that in 2016 only 20% of all Uber rides were made with its ride-sharing service UberPool (Lunden, 2016). Assuming that UberPool has not generated new demand by its lower prices and that each vehicle is shared with just one other passenger, it means that 10% of the originally needed Uber vehicle trips have been replaced by sharing, and less than 10% of the vehicle kilometers due to the detouring that ride-sharing involves. When considering that not all ride-sharing users were actually using the conventional taxi service Uber X before, the figures might turn out worse. The question arises whether the previous studies cover the real potential of ride-sharing or whether oversimplifications are made in relation to price, delay etc. that lead to unrealistically high reductions in vehicle kilometers resulting from those services.

In all studies on ride-sharing potential, user discomfort is accounted for by means of a maximum accepted travel time delay. Naturally, if a shorter travel delay is accepted, the potential for sharing decreases (Santi et al., 2014; Alonso-Mora et al., 2017; Fiedler et al., 2018). Hence, possibly the generally assumed accepted delay of five minutes is too high, or there are other factors that play a role in a user's decision to share or not that result in its currently limited usage in the cities where ride-sharing is available.

Surveys reveal that next to time delays, privacy issues are the most important deterrent for potential users or ride-sharing services (Dueker, Bair, & Levin, 1977; Teal, 1987). Not only do road users negatively value the feeling of dependence, also the fear of having a negative social interaction with other users appears to be a relevant deterrent (Morales Sarriera et al., 2017; Correia & Viegas, 2011). The opportunity to have a positive social interaction on the other hand is hardly considered an incentive for ride-sharing.

The most relevant incentive for ride-sharing is a cut in the travel costs (Teal, 1987). As ride-sharing considers a trade-off between extra travel time caused by picking up and dropping off other passengers and a financial compensation (Morales Sarriera et al., 2017), it is remarkable that none of the studies on ride-sharing potential implicitly incorporates

discounts in travel costs. If the willingness to allow a delay of five minutes every trip for example requires excessive price cuts, the potential presented in earlier works might not be economically feasible.

What adds to the problem is that every user has its own preferences. Various studies show that there is significant heterogeneity in users' behavioral preferences towards ride-sharing. A latent class analysis, for example, distinguishes two classes when high-occupancy vehicle (HOV) lanes are present: one with a higher and one with lower tendency to ride-sharing. It appears that people that are more sensitive to both time and cost are more likely to share a ride than to drive individually (Xiong, Hetrakul, & Zhang, 2014). Classification based on users' preferred roles gives a different insight. For people who prefer to drive, social interaction can be an important incentive to share a vehicle, while for users who prefer the passenger role, safety and flexibility are the most important issues. For users open to both roles, cost reductions were an important incentive (Park, Chen, & Akar, 2018). Other personal characteristics that influence tendency to ride-share include gender and character. For example, social safety in the vehicle is valued higher by women than men (Morales Sarriera et al., 2017). Also, introverts were found to be less likely to share rides than extroverts, as the latter group valued social interaction more positively (DeLoach & Tiemann, 2012).

1.4 Research gap

In existing work on the potential societal and user impact of ride-sharing services, it has not been acknowledged that the choice for ride-sharing regards a complex trade-off between delays, a financial reward and psychological discomfort following from a lack of privacy and independence. Previous studies have simplified users' choice by considering merely a maximum time delay that users tolerate. The incorporation of psychological discomfort and discounted ride fares in ride-sharing choice is new and will provide more realistic insights into the potential of ride-sharing.

As behavioral preferences have been found to vary among users, heterogeneity in its parameters can influence the results. The fact that some people have a higher willingness to ride-share than others can potentially have a crucial impact on the attractiveness and ultimately the success of a ride-sharing service. Consideration of such heterogeneity among users when analyzing trip shareability would also be new in research, as illustrated by Table 1.1.

Finally, little is known about how ride-sharing potential for improving user mobility and reducing the total vehicle distance is dependent on the spatial distribution of demand. The results of the previous studies are mostly context-specific. One of the few studies accounting for the spatial lay-out of requests, by Bimpikis, Candogan, and Saban (2016), state that both user benefit (expressed by means of the consumer surplus) and operator profit are higher when the spatial distribution of demand is more balanced, i.e. the demand originating at each location is relatively close to the demand for those locations as destination. It is interesting to find out whether this also applies when the choice to share a ride consists of a trade-off between delays, a psychological discomfort and a financial compensation.

1.5 Relevance of research

With the introduction of a ride-sharing service, users can potentially benefit from lower prices compared to ride-hailing services if they are willing to accept a minimum delay and other passengers in their vehicle. Operators can benefit too, as ride-sharing provides them with the opportunity to execute travel requests more efficiently and consequently to cut the

Study	Vehicle type	Real-time	Reassignment	Rebalancing	Vehicle capacity	Heterogeneity in user preferences	User cost-benefit trade-off	Scale
Ma et al., 2013	Taxi	✓			3			Large
Santi et al., 2014	Taxi				2-3			Large
Alonso-Mora et al., 2017	Taxi	✓	✓	✓	1-10			Large
Tachet et al., 2017	Taxi				2-3			Medium to large
Fiedler et al., 2018	MoD	✓		✓	5			Medium
This study	Taxi	✓	✓	✓	3	✓	✓	Small

Table 1.1: Overview of studies on ride-sharing potential and how this work fits in

total costs of operation. A government will be most interested in a reduction in the number of vehicles on the road and the total distance they produce. To stimulate the adoption and success of ride-sharing services, authorities can use regulation. It can range from changing the tendering procedures of transport services, to adapting and implementing standards for ride-sharing services, to creating incentives for users and operators to participate in ride-sharing initiatives. An example of the latter is a subsidy on a part of the standard taxi fare if such a ride is shared with other users.

By analyzing the effect of the spatial distribution of demand as well as the attitudes of potential users towards delays and sharing, this research can help to identify in which markets the introduction of a ride-sharing service is most likely to be successful. Thus, it can help authorities and service providers to get a first idea of whether the application of a ride-sharing service has potential to generate efficiency, service quality and revenue gains in their specific context. At the same time, it will also give insights into the effect of specific design parameters on the success of a ride-sharing system. For example, as this study considers ride-sharing choice as the trade-off between sharing inconveniences and a financial reward, it is possible to study the effect of different discount schemes on the societal benefits of a ride-sharing system. It is interesting especially to find out whether ride-sharing efficiency can be boosted by awarding higher discounts to users if they share their ride with more co-riders.

To summarize, this research can help authorities and service providers to get a first estimate on the societal benefits that a ride-sharing service can generate given the situation at hand. Also, it includes implications for the design of a successful ride-sharing service.

1.6 Research questions

The research gap is translated into the following research question, in which efficiency is evaluated in terms of the total number of vehicle kilometers in the system and the level of service in terms of the acceptance rate and average delay:

> *How do efficiency and level of service of dynamic ride-sharing services depend on behavioral and service attributes and the distribution of demand?*

This research will not consider all behavioral and service attributes but rather focus on a few specific ones. Specifically, this research will consider users' willingness to accept extra

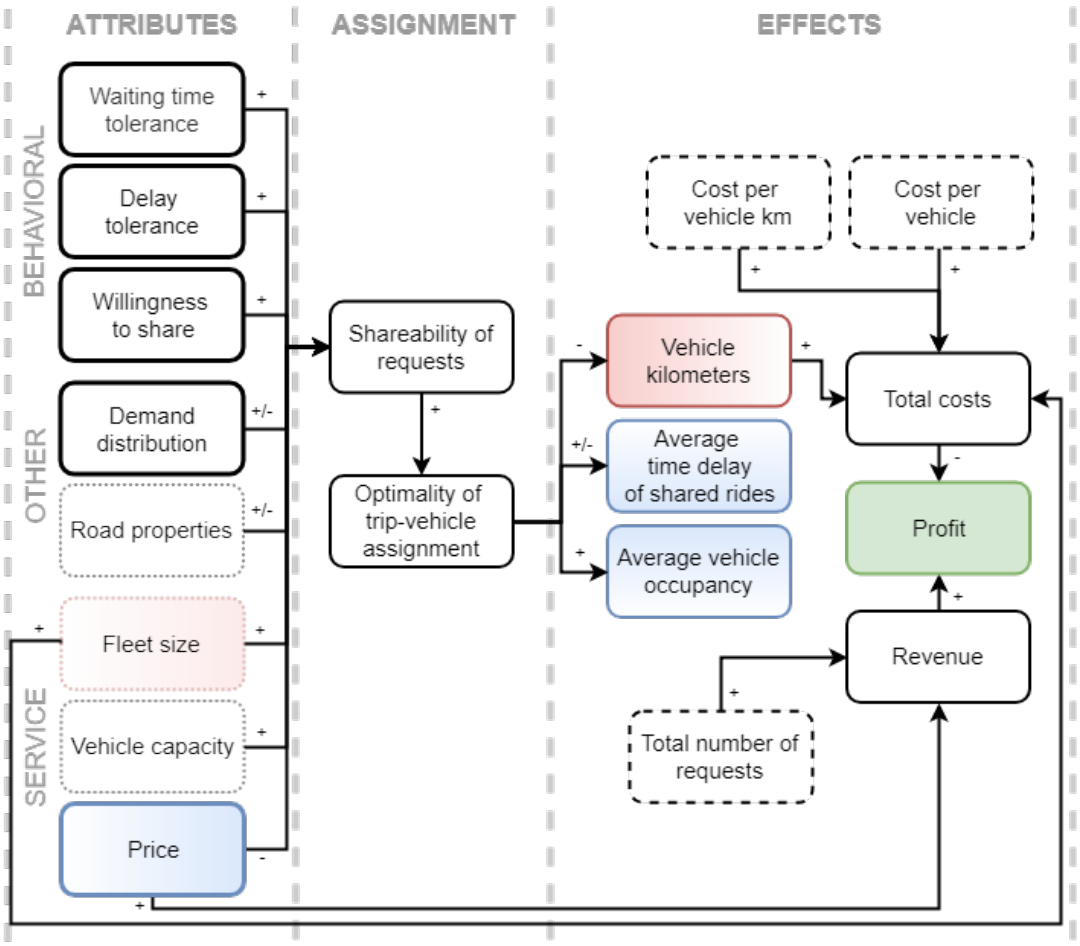


Figure 1.1: Efficiency (red), level of service (blue) and profitability (green) of ride-sharing platforms based on individual, collective and service attributes

waiting and total travel time, their willingness to share a vehicle with other passengers, the spatial distribution of demand and ride fares. Fleet properties are thus not in the focus of this research. Figure 1.1 highlights the five attributes to be analyzed with their potential effects.

The main research question is broken down into several sub questions. The first sub question considers the effect of the delay tolerance and willingness to share on the efficiency and level of service of a ride-sharing system. Sub question two provides an additional insight on this by including heterogeneity when analyzing the impact of behavioral preferences. It is also interesting to consider how the efficiency changes when is assumed that the willingness to share depends on the number of passengers that a ride is shared with. The effect of this assumption on ride-sharing efficiency and level of service will be investigated with research question three. Similarly, the fourth sub research question analyzes the efficiency and level of service when the number of co-riders is incorporated into the discount, as opposed to a fixed discount per ride that is executed with the ride-sharing service. Research question five looks at the spatial distribution of demand to find whether a ride-sharing system performs better when demand patterns are more balanced / requests are spread more evenly over the network. Finally, research question six addresses the relationship between the different ride-sharing objectives to find if they are mutually exclusive, meaning that a trade-off is always required, or whether multiple objectives can be achieved at the same time. Specifically, the relationship between the acceptance rate, average user delay and the vehicle distance savings are considered in this research question.

To summarize, the following list of sub questions will be answered with this research:

1. *How do efficiency and level of service of a dynamic ride-sharing service depend on the delay tolerance, waiting time tolerance and willingness to share, assuming uniformity across users?*

2. *What is the effect of heterogeneity in the delay tolerance, waiting time tolerance and willingness to share on the efficiency and level of service of a dynamic ride-sharing service?*

3. *How do efficiency and level of service of a dynamic ride-sharing service depend on how willingness to share relates to the number of co-riders in the vehicle?*

4. *How do efficiency and level of service of a dynamic ride-sharing service change when the fare discount for a shared ride depends on the number of co-riders?*

5. *How do efficiency and level of service of a dynamic ride-sharing service depend on the spatial distribution of demand over the network?*

6. *What is the relation between the acceptance rate, the average delay of shared rides and the total reduction in vehicle distance following from a ride-sharing service?*

1.7 Outline

The sub research questions are used as a guideline throughout the remainder of the report. Chapter 2 textually and mathematically describes the followed approach, thereby referring to the sub research questions. It also includes an overview of the main indicators that are considered for the level of service and efficiency of a ride-sharing service, as well as a description of the model implementation. In Chapter 3, the experiment to which the method is applied is introduced, distinguishing the general set-up of the experiment and the design of different scenarios to answer the individual sub research questions. Chapter 4 presents the main results for each sub research question, following from the model application. Finally, Chapter 5 concludes the report by summarizing and synthesizing the main findings and stating the main practical recommendations. It also includes recommendations for further research. For the sake of completeness, the full list of results is presented in Appendix A. Appendix B includes a detailed specification of the three scenarios related to sub research question 5.

Chapter 2

Methodology

This section describes the approach that is followed in this study. First, the approach directions of previous studies on ride-sharing are examined in Section 2.1. This section also substantiates which approach direction is most suitable given the objectives of this study. Section 2.2 then gives an outline on the chosen approach and how it incorporates the previously addressed trade-off between ride-sharing discomfort, delay and a lower ride fare. Section 2.3 introduces the symbology used in the more detailed description of the different steps of the method, as presented in Section 2.4. Section 2.5 contains the definitions of the main Key Performance Indicators (KPIs) that are used in this study to evaluate ride-sharing efficiency and level of service. Finally, the main choices regarding the implementation of the model are summarized in Section 2.6.

2.1 Selection of modelling approach

In previous research on ride-sharing potential, a wide variety of approaches has been applied. This section distinguishes and summarizes the main approach directions taken in such research and motivates the approach direction pursued in this study.

One of the first studies in the field considered ride-sharing in the light of a Dial-a-Ride problem (DARP) (Ma et al., 2013), which is also known as the Vehicle Routing Problem with time windows. The static version of this problem is applied widely in the field of transportation when planning vehicle schedules. To allow for dynamic ride-sharing in which requests emerge in real-time, the DARP has to be applied dynamically. In this case, each incoming request is individually allocated to a vehicle using a greedy algorithm. Therefore, it does not lead to an optimal solution at a system level.

To acquire an optimal solution, requests need to be assigned concurrently rather than sequentially. A study on ride-sharing in New York by Santi et al. (2014) addressed the optimality issue by introducing the concept of shareability graphs (SG). Nodes in such a graph structure, representing trip requests, are connected if the respective requests can be shared with a maximum acceptable delay for each individual request. This study takes a static approach, meaning that is assumed that all requests are known beforehand. Also, it is limited to two passengers per vehicle. The fact that information on the pairwise shareability of requests is stored in a graph structure means that optimization can be done with traditional graph-solving methods. This offers computational efficiency essential for large-scale problems.

In a follow-up study by Alonso-Mora et al. (2017) on the same data set graph structures were applied as well, however in a different way to allow for dynamic ride-sharing with more than two passengers per vehicle. While the information in a shareability graph only includes information about pairwise shareability of requests, the graphs proposed in their research included also vehicles and request groups (combinations of requests to be

Property	DARP	SG	RGV	ABM
Congestion	+	N/A	+	+
Heterogeneity	+	++	++	+++
High-capacity	++	+	+++	+
Optimality	N/A	+++	+++	N/A
Real-time	+++	N/A	+++	+++
Rebalancing	+++	N/A	+++	+++
Scalability	+	+	+++	+++

Table 2.1: Suitability of each approach to fulfill a certain desired property (from impossible to implement (N/A) to already capable (+++)

executed with one vehicle, called 'trips' in the respective study). Firstly, new requests are collected in a small time window (e.g. 30 seconds). A Request-Vehicle (RV-)graph is then composed with information on both the pairwise shareability of requests (like the previously mentioned approach) and the compatibility of a request with a vehicle, thereby considering the itineraries of passengers already on-board of the vehicle. The graph containing information about request-request and request-vehicle matches is then translated to one that incorporates request groups: the Request-Group-Vehicle-graph (from hereon described as the RGV-graph). This graph gives information about which request group can be executed with which vehicle and at what cost. Finally, groups are assigned to vehicles by solving the RGV-graph as a simple integer linear problem (ILP), allowing for assignment of large numbers of requests and vehicles within a limited computation time.

Another study considered ride-sharing in a Mobility-on-Demand (MoD) system, using an agent-based simulation called AgentPolis (Fiedler et al., 2018). In an agent-based model (ABM), individuals' travel demand is based on the sequence of activities that they perform during a day. This allows for more realistic interactions between users and vehicles. However, these interactions also make this type of simulation computationally heavy. For the study of Prague, this problem was mitigated by estimating travel time based on the Eucledian distance and by assigning requests to vehicles based on an inserting heuristics approach. Like DARP, trip requests are assigned individually to minimize travel time, thus not necessarily leading to a system-optimum.

Table 2.1 summarizes the four approaches based on whether they fulfill or can easily be adapted to fulfill each of the following desired properties:

- Effect of ride-sharing on *congestion* on the road network and vice versa
- *Heterogeneity* in attribute values among different users
- *High-capacity* ride-sharing (more than 3 passengers per vehicle)
- Guarantee of an *optimal* assignment of requests on a system level, given a batch of requests and a long enough run time
- (Near) *real-time* assignment of requests (possibly with a small time window to gather those requests)
- *Rebalancing* empty vehicles based on expected demand
- *Scalable* to large networks and high demand levels

The RGV-approach fulfills the most properties desired given the context of this study. It already allows studying dynamic ride-sharing with high-capacity vehicles that can be reassigned if a better match is found and that can rebalance to areas with high demand if unassigned. Assignment of request groups to vehicles happens optimally given there

is sufficient time to run the complete model, otherwise it will yield a suboptimal result. Finally, incorporation of heterogeneity of individual attributes relating to the willingness to share and the delay tolerance, one of the aims of this study, can be introduced fairly easily given the model properties. The main downside of this approach is the relative difficulty to incorporate the interaction between request assignment and road congestion, which however also applies to the other methods. All in all, the RGV-approach is found most suitable for the aims of this study.

2.2 Assignment procedure

In the RGV-approach as introduced in the previous section, requests are first collected during a small time window and then collectively assigned to vehicles, thereby considering the dynamic nature of a ride-sharing system as well as allowing for multiple requests to be performed with a single vehicle. As shown by Figure 2.1, each assignment phase consists of nine steps:

1. Each assignment phase starts with an exploration of the current ride requests. The pool of requests for assignment is created by aggregating all potentially feasible unassigned requests and all assigned but not yet picked-up requests from the previous assignment phase, as well as new ride requests that came in since the last assignment. Requests that were left unassigned in the previous assignment are only reconsidered for assignment if they can still be served given an optimum ride without detours, with an immediate pick-up and with a maximum discount (which means that in theory it is still possible for the user to receive a net benefit). All other requests can impossibly be assigned to a ride in the next, nor in any of the following assignments, and are therefore rejected by the service. The assumption in this step is that users are willing to wait to be assigned to a ride, even without the guarantee that a suitable ride can be found for them. In this step also the current status of vehicles is loaded, based on the schedule that was assigned to each vehicle in the previous assignment phase.
   In the example: there are three requests in the batch; the fleet consists of two (currently empty) vehicles.

2. Each request in the pool of requests is pairwise matched to each of the other requests in the batch to find if a pair of requests can be served with a single (hypothetical) vehicle, meaning that both requests experience a net benefit. This step only considers combinations of two requests, not more (irrespective of the capacity of vehicles in the fleet). Neither does it consider the status of vehicles in the fleet at the time of assignment.
   Requests 1 and 2, and 2 and 3 possibly form feasible request pairs.

3. Each request in the pool of requests is individually matched with each of the vehicles in the fleet (possibly with passengers on-board), to find for each request-vehicle match whether the request can be fulfilled with that particular vehicle, meaning that the capacity of the vehicle is not exceeded and the request experiences a net benefit. This step does not consider the combination of requests into request groups.
   All requests can be served by vehicle 1, only request 2 can be served by vehicle 2.

4. A request-vehicle-graph (RV) is created that stores the pairwise request matches (step 2) and the request-vehicle matches (step 3). The nodes in the RV-graph are vehicles and requests, the presence of an edge indicates that a request can be served with a vehicle.
   RV-graph as shown in Figure 2.1.

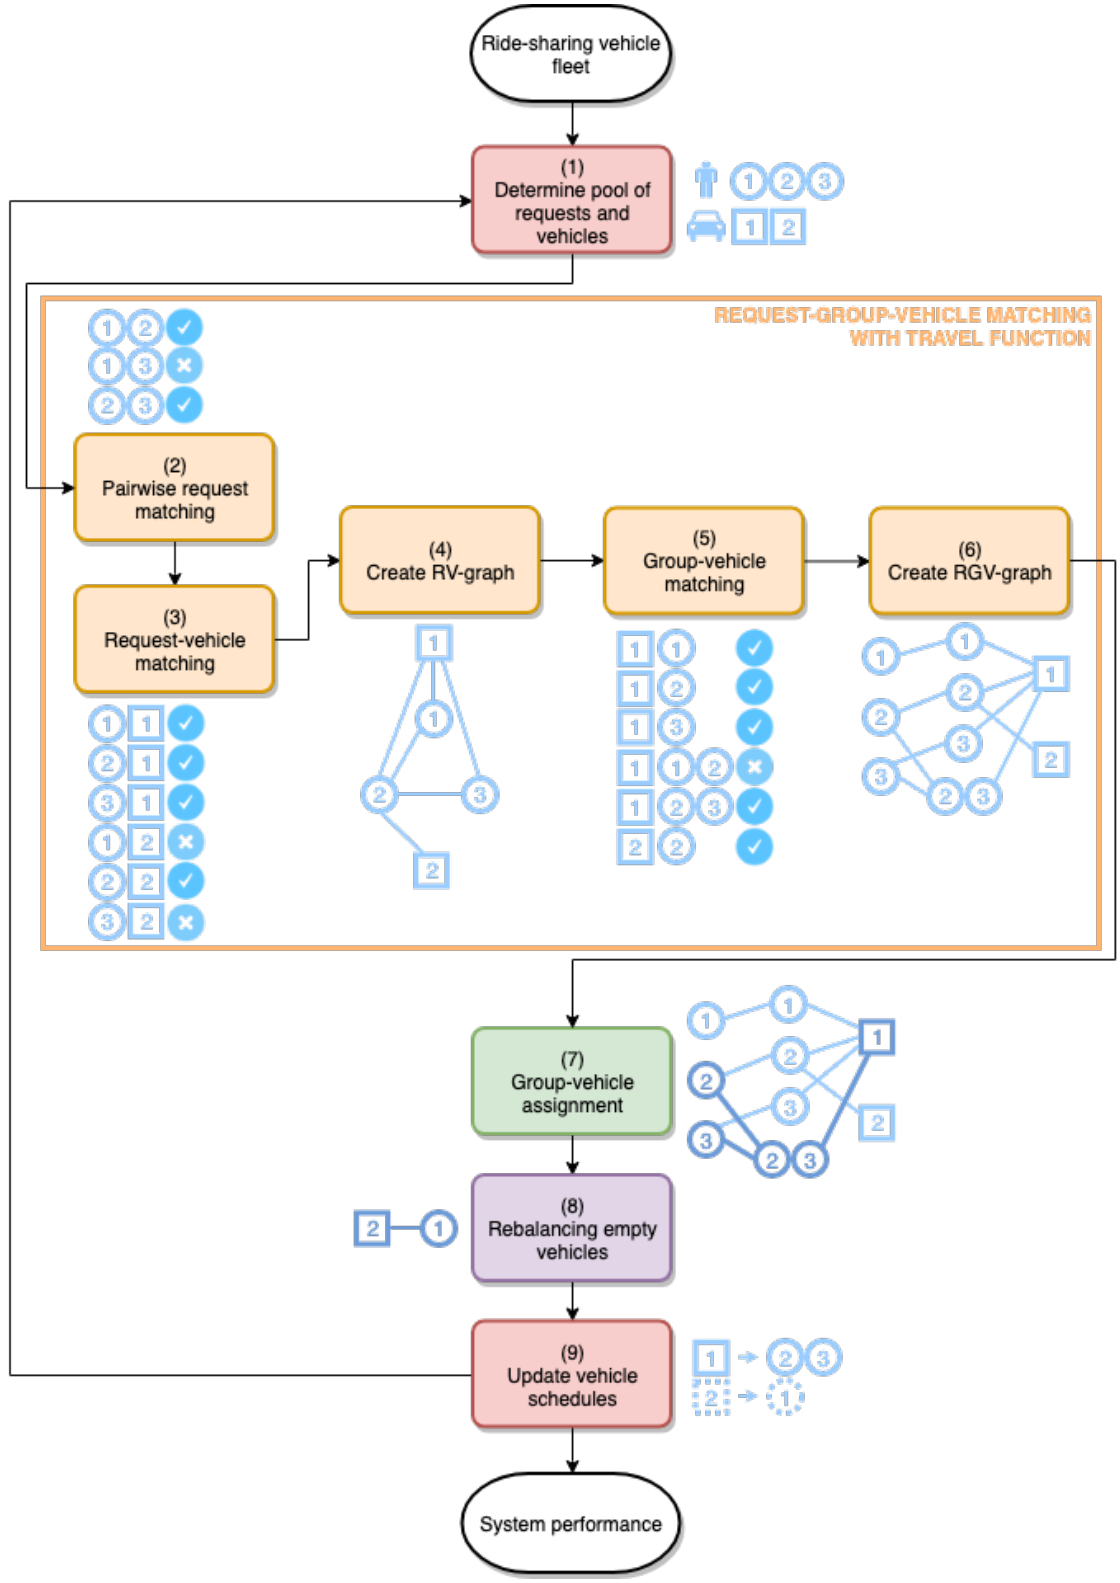


Figure 2.1: Overview of the methodology, including an example (in blue) of an assignment phase with three requests and two vehicles (without passengers)

5. Based on the cliques in the RV-graph, individual requests are combined into groups that can possibly be served using a specific vehicle, without yet considering a route to do so. The list of feasible groups is then narrowed down by checking whether a group-vehicle (GV) combination actually has a feasible route (i.e. there is a net benefit for all passengers in the group and vehicle capacity is never exceeded). A route not only visits the request origins and destinations but also destinations of passengers already on-board of the vehicle, if any.
   Based on the RV-graph, a group consisting of requests 1 and 2, and a group of requests 2 and 3 can be potentially be served using vehicle 1. Besides, all groups with only one request can potentially be served by vehicle 1, and a group containing only request 2 by vehicle 2, as found in step 4. Of all those group-vehicle combinations, only the group containing requests 1 and 2 executed with vehicle 1 does not appear to be feasible when considering routes to serve all requests in a group with a particular vehicle.

6. The result of the previous step is stored in the request-group-vehicle graph (RGV). This graph contains request, group and vehicle nodes. A group and vehicle node are connected by an edge if in the previous step was found that a feasible route exists to serve the request group with the particular vehicle. The maximum sum of net benefits of all requests in the group and all passengers on-board the vehicle is stored as label of the edge, considering all routes that can satisfy the group-vehicle combination. The graph also includes which requests are part of a group by the presence of edges between request nodes and group nodes.
   RGV-graph as shown in Figure 2.1, with request (left), group (middle) and vehicle (right) nodes.

7. The RGV-graph that followed from the group-vehicle matching step is used as input for the group-vehicle assignment, in which vehicles are actually assigned to serve a group of requests. The graph structure allows solving the group-vehicle assignment as an Integer Linear Problem with the costs of all GV-combinations stored in the edges of the RGV-graph. The group-vehicle assignment is optimized based on an objective function, aiming for a maximum total benefit for all passengers in the system or a maximum number of served requests.
   A group containing requests 2 and 3 is assigned to vehicle 1. Vehicle 2 cannot serve request 1 and is left unassigned.

8. If the group-vehicle assignment leaves both requests and vehicles unassigned, rebalancing takes place. Empty vehicles move in the direction of the origin of an unassigned request, anticipating on new requests that may arise in this unserved area. Vehicles are assigned to requests based on the travel time between vehicle locations and origin requests. The aim is to mimimize the total travel time for rebalancing vehicles.
   Request 1 and vehicle 2 are both unassigned. Vehicle 2 moves to the origin of request 1.

9. Vehicle schedules are updated based on the output of the assignment and the rebalancing step. This includes updating the location, passengers (considering pick-ups and drop-offs) and operating status (moving or halting) of every vehicle to the time of the next iteration.
   The status of vehicle 1 is updated based on its assignment to the group consisting of requests 2 and 3, the status of vehicle 2 is updated based on its assignment to rebalance in the direction of the origin of request 1.

Although the three matching steps (2,3,5) have a different purpose, they are in fact fairly similar. Each step checks the feasibility of one or a sequence of requests to be executed with a (hypothetical or actual) vehicle by varying the order of pick-ups and drop-offs and finding the route with the maximum total benefit, considering the difference between the discount received and the monetized delay and sharing costs. The feasibility of routes is thereby restricted by demand and supply constraints. A 'travel function' is designed to check the feasibility of one or multiple requests to be executed with a specific vehicle by checking all sequences of pick-ups and drop-offs for the relevant constraints and selecting the sequence that has maximum total benefit while satisfying all constraints. The travel function can be used for each of the three matching steps.

2.3 Mathematical problem definition

This section lists the parameters, sets, sequences and variables used in the more detailed description of the methodology. Note the difference in usage of 'stop locations' (i.e. all locations where a vehicle may stop / where requests can be picked up and dropped off) as compared to 'stops' (i.e. stop locations where a specific vehicle schedules to stops to drop-off passengers or pick-up requests). Also, waiting time tolerance, delay tolerance and willingness to share from hereon will be defined as the the inverse, i.e. respectively waiting time aversion, delay aversion and reluctance to share.

2.3.1 Parameters

τ^{sim}	Total duration of simulation
τ^{int}	Time interval between consecutive assignment phases
μ	Fleet size, i.e. number of vehicles in operation
κ	Vehicle capacity (a homogeneous fleet is assumed)
π_z	Percentage of discount offered to request z for a shared ride
α_z	Waiting time aversion associated with request z
β_z	Delay aversion associated with request z
γ_z	Reluctance to share a vehicle with other passengers associated with request z

2.3.2 Sets

V	Set of vehicles in fleet $V = \{v_1, \ldots, v_\mu\}$, in which each v is defined by $\{q_v, t_v, l_v, h_v\}$
Z	Set of all incoming requests in time period of simulation $[0, \tau^{sim}]$ $Z = \{z_1, \ldots, z_{\lvert Z\rvert}\}$, in which each z is defined by $\{o_z, d_z, t_z^r, c_z, \alpha_z, \beta_z, \gamma_z\}$
Q	Set of stop locations (intersections) within a road network $Q = \{q_1, \ldots, q_{\lvert Q\rvert}\}$
E	Set of road segments in a road network $E = \{e_1, \ldots, e_{\lvert E\rvert}\}$, in which an e is defined by $\{tt_e, d_e, q_{e1}, q_{e2}\}$
R	Set of all active requests at the start of an assignment phase $R \subseteq Z; R = \{r_1, \ldots, r_{\lvert R\rvert}\}$
R^{new}	Set of new requests from Z since the last assignment ($t_{-1} < t_z^r < t$) $R_{new} \subseteq Z$
R_v	Set of requests assigned to vehicle v $R_v \subseteq R$; $R = \{r_1, \ldots, r_{\lvert R_v\rvert}\}$

P_v Set of passengers on-board of vehicle v. A passenger is a former request r that has been picked-up in the past by v and has not yet been dropped off
$P_v = \{p_1, \ldots, p_{l_v}\}$, in which each p is defined by $\{r, t_p^{pu}\}$

V^a Set of all vehicles that are assigned to fulfill requests in R
$V^a \subseteq V$

V^u Set of all vehicles that are not assigned to fulfill any requests in R. V^a and V^u are mutually exclusive and collectively exhaustive
$V^u = (V^a)^c$

R^a Set of all requests assigned to a vehicle in the assignment phase
$R_a \subseteq R$

R^{ap} Set of requests in R^a picked-up by a vehicle before the next assignment
$R^{ap} \subseteq R^a$

R^{anp} Set of requests in R^a not picked-up by a vehicle before the next assignment. V^{ap} and V^{anp} are mutually exclusive and collectively exhaustive
$R^{anp} = (R^{ap})^c$

R^u Set of all requests not assigned to a vehicle in the assignment phase. R^a and R^u are mutually exclusive and collectively exhaustive
$R^u = (R^a)^c$

V^p Set of all vehicles with passengers P_v on-board that are not assigned to any requests in R
$V^p \subseteq V^u$

V^{np} Set of all empty vehicles that are not assigned to any requests in R. V^{np} and V^p are mutually exclusive and collectively exhaustive
$V^{np} = (V^p)^c$

V^{reb} Set of empty unassigned vehicles that are assigned to rebalance
$V^{reb} \subseteq V^{np}$

K_v Set of routes visiting all locations of R_v and P_v with vehicle v

G Set of request groups that are feasible given at least one vehicle
$G = \{g_1, \ldots, g_{|G|}\}$

A Set of all assignment phases in $0 < t < \tau^{sim}$

2.3.3 Demand and network characteristics

o_z Origin of request z

d_z Destination of request z

t_z^r Incoming time of request z

c_z Standard fare associated with request z

tt_e Travel time on edge e (static)

d_e Length of edge e

2.3.4 Sequences

S_v Route for vehicle v, consisting of a sequence of stops s (starting at its current location q_v)
$S_v \in K_v$

S_v^* Optimal route in set of routes K_v
$S_v^* \in K_v$

$S_v{}'$ Sequence of all stops in S_v reached before next t
$S_v{}' \subseteq S_v$

U_v	Sequence of all stop locations (note: not necessarily locations where vehicle v stops) in route S_v, given shortest path between consecutive stops s $U_v \supseteq S_v$
$U_v{}'$	Sequence of all stop locations in U_v that can be reached before the next t $U_v{}' \subseteq U_v$
W_v	Sequence of all stop locations along shortest path from q_v to o_r for a rebalancing vehicle in V^{reb} assigned to r
L_v	On-board load between all respective stops s in route S_v $L_v = (l_1, \ldots, l_{\lvert S_v \rvert - 1})$
$L_{v,p}$	On-board load between respective stops s in route S_v while passenger p is on-board $L_{v,p} \subseteq L_v$
S_r	Sequence of stops experienced by request r between pick-up and drop-off

2.3.5 Variables

t	Current time in simulation
q_v	First location q that vehicle v reaches after next assignment phase commences
t_v	The time by which vehicle v reaches location q_v
l_v	Passenger load on-board vehicle v at location q_v $l_v \in \{0, \ldots, \kappa\}$
h_v	Binary variable indicating if vehicle v halts at location q_v for a pick-up or drop-off $h_v \in \{0, 1\}$
n_p^{comax}	Maximum simultaneous number of co-riders with which passenger p shares a ride
t_s	Time a vehicle v reaches stop s in S_v
t_p^{pu}	Time passenger p has been picked-up by vehicle v
tt_{q_i,q_j}	Travel time between two locations q_i and q_j (given shortest path)
$dist_{q_i,q_j}$	Distance between two locations q_i and q_j (given shortest path)
del_r	Delay experienced by request r when following route S_v
wt_r	Waiting time for a request r when a route S_v is followed
dsc_r	Absolute fare discount for request r when sharing a ride
b^a	Total benefit of an assignment

2.4 Main steps

This section describes in more detail the assignment procedure as introduced in Section 2.2. The description of the methodology is divided into four main parts, each corresponding to one color in Figure 2.1. First, the algorithm that is used for each of the three RGV-matching steps (pairwise request matching, request-vehicle matching and group-vehicle matching, respectively steps 2, 3 and 5 in Figure 2.1) is explained in Subsection 2.4.1. Then, the procedure for assigning request groups to vehicles (step 7 in Figure 2.1) is clarified in Subsection 2.4.2. The algorithm used to determine which vehicles rebalance in the direction of which requests (step 8 in Figure 2.1) is described in Subsection 2.4.3. Finally, Subsection 2.4.4 gives information on how vehicle schedules are updated given the result of an assignment and how the status of vehicles and the new pool of requests are determined for the next assignment (steps 9 and 1 in Figure 2.1).

In the description of the methodology, special attention should be devoted to the incorporation of a cost-benefit trade-off for users in the system, as this is a novelty for

studies on ride-sharing potential. In previous approaches, level of service was considered only in terms of a maximum acceptable waiting time and delay, not allowing for their trade-off with a reduction in travel costs when a ride is shared. The application of the cost-benefit trade-off in this approach is threefold:

1. As a constraint in the analysis of feasible request (group) - vehicle combinations (described in Section 2.4.1): every request / passenger requires a net positive benefit
2. In the objective function when looking for the most beneficial feasible route to serve one or multiple requests with a particular vehicle (also Section 2.4.1): by summing the benefits of all passengers / requests in the group-vehicle combination
3. In the objective function when assigning request groups to vehicles (Section 2.4.2): by summing the benefits of the requests and passengers in all chosen group - vehicle combinations

2.4.1 RGV-matching with travel function

Description

The solution space for the group-vehicle assignment is significantly reduced by first checking the feasibility of request, group and vehicle combinations (steps 2-6 in Figure 2.1). For this purpose a travel function is designed that looks for the optimal route that a vehicle v can take to serve a request r or multiple requests R_v while satisfying the relevant constraints, if such a route exists. The travel function is multifunctional in the sense that it can be used in each of the three matching steps, albeit with different input as summarized in Table 2.2. The main difference between the application of the travel function in the pairwise request-matching step (step 2 in Figure 2.1) and in the other two matching steps (steps 3 and 5 in Figure 2.1) is that in the first one a hypothetical vehicle at the origin of one of the two requests ($q_v \leftarrow o_r$) is assumed, not an actual vehicle v in the system with a given location q_v and (possibly) passengers P_v on-board occupying a part of the vehicle capacity. Hence, the first step overestimates the feasibility of a group consisting of these two requests. However, the purpose of this step is merely to improve computational efficiency of the model by narrowing down the initial list of potential groups in G. The other two steps narrow down the list of feasible groups even further to only include request groups R_v that can actually be combined given the actual location q_v of the vehicle and the availability of seats, i.e. residual capacity, throughout the ride.

Algorithm

The procedure to find the feasibility of one or multiple requests to be served with a specific vehicle, is presented in Algorithm 1. The algorithm is explained in more detail below.

There are multiple potential routes S_v in K_v to serve a single or a combination of requests $R_v \in R$ with a particular (possibly hypothetical) vehicle v. For each request r in R_v, the vehicle route must visit two locations: the request origin o_r and the desired destination d_r. Also, if the vehicle has picked up requests in the past, there are passengers

Step	Fig. 2.1	Vehicle	Passengers	Requests
Request-matching	(2)	Hypothetical	None	Two
Request-vehicle matching	(3)	Actual	Possibly	One
Group-vehicle matching	(5)	Actual	Possibly	One or multiple

Table 2.2: Input for travel function in each of the matching steps

Algorithm 1 Feasibility of a group-vehicle combination R_v

for $S_v \in K_v$ **do**
  if $\max(L_v) \leq v$ **then** ▷ Capacity constraint
    if $\forall r \in R_v \cup P_v$, $\pi_r \cdot c_r \geq (t_r^d - t_r^*) \cdot \beta_r - (t_r^{pu} - t_r^r) \cdot \alpha_r - \gamma_r$ **then** ▷ LoS constraint
      $b_{S_v} \leftarrow \sum_{r \in R_v \cup P_v} (\pi_r \cdot c_r - (t_r^d - t_r^*) \cdot \beta_r - (t_r^{pu} - t_r^r) \cdot \alpha_r - \gamma_r)$ ▷ Total benefit
    else
      $b_{S_v} \leftarrow$ "Invalid"
    end if
  else
    $b_{S_v} \leftarrow$ "Invalid"
  end if
end for
if $\exists S_v \in K_v$ such that $b_{S_v} \neq$ "Invalid" **then**
  $b_{R_v} \leftarrow \max(b_{S_v})$ ▷ Route with maximum benefit
end if

$P_v = \{p_1, \ldots, p_{l_v}\}$ on-board the vehicle destined for their respective destinations. Each route has to start in the current location q_v of the vehicle, meaning that each route consists of $2 \cdot |R_v| + |P_v| + 1$ stops in S_v. Potentially some of these locations are duplicates that can be removed from the sequence, reducing the total number of locations visited. The order in which locations are visited can vary as long as the origin o_r of each request proceeds its destination d_r. Between two consecutive stops s_1 and s_2 in a route, vehicles always take the shortest path (i.e. the path with the minimum travel time). It is thereby assumed that the travel time tt_e on all links in the network is independent of the routes that vehicles take (i.e. the impact of fleet mileage on congestion is negligible) and deterministic. Consequently, vehicles always follow the same path between any pair of nodes. Therefore, the travel time tt_{q_i,q_j} between two stop locations q_i and q_j, and consequently also between two stops along route S_v, can be extracted from a predefined lookup table that contains the travel time for all node combinations. A route S_v is therefore defined by the sequence of origins and destinations rather than all nodes that are visited along the route.

The optimal route S_v^* for a vehicle to serve a set of requests R_v is the route with the maximum total benefit, comprising for each request and passenger of the difference between the fare discount dsc_r offered (i.e. the discount rate π_r multiplied by the ride fare c_r) and the sum of the monetized disbenefits following from sharing a ride. As explained in Chapter 1, the potential disbenefits of ride-sharing include an additional waiting time wt_r and a longer total delay del_r. The total disbenefit following from the extra waiting time and total delay can be determined by multiplying the waiting time wt_r and total delay del_r for that request or passenger respectively by the waiting time aversion α_r and delay aversion β_r. Another disbenefit is found in the loss of privacy and comfort resulting from the presence of co-riders in the vehicle, defined as the reluctance to share γ_z or the compensation needed to choose a shared ride when there is a guarantee of no delay. This can also be perceived as an equivalent to the in-vehicle multiplier applied in public transport route choice models when there is some level of crowding in the vehicle (Yap, Cats, & van Arem, 2018). For simplicity's sake, the reluctance to share in this case is determined based only on the maximum number of co-riders throughout the journey, rather than the relative time spent on board with each number of co-riders. Based on the previous, the net benefit for each request r is defined as:

$$b_r = \underbrace{\pi_r \cdot c_r}_{\text{fare discount}} - \underbrace{del_r \cdot \beta_r}_{\text{delay disbenefit}} - \underbrace{wt_r \cdot \alpha_r}_{\text{waiting disbenefit}} - \underbrace{\gamma_r}_{\text{sharing disbenefit}} \tag{2.1}$$

The delay of a request (or passenger) del_r is defined as the difference between the earliest possible arrival time t_r^* (equal to the request time t_r^r plus the direct travel time

RQ	$\alpha_\mathbf{r}, \beta_\mathbf{r}$	$\gamma_\mathbf{r}$	$\pi_\mathbf{r}$
1	Vary α, β	Vary γ	
2	Request-specific α_r, β_r	Request-specific γ_r	
3		$\gamma_r(n_p^{comax})$	
4			$\pi_r(n_p^{comax})$

Table 2.3: Summary of the configuration of Algorithm 1 for different sub research questions

tt_{o_r,d_r} between origin and destination) and the arrival time at the destination t_r^d if route S_v is followed. The waiting time wt_r consists of the difference between pick-up time t_r^{pu} and request time t_r^r. The net benefit for each request r can therefore be rewritten to:

$$b_r = \pi_r \cdot c_r - (t_r^d - t_r^*) \cdot \beta_r - (t_r^{pu} - t_r^r) \cdot \alpha_r - \gamma_r \tag{2.2}$$

At the same time, every route is restricted by two constraints. Firstly, the vehicle capacity κ cannot be exceeded by the number of passengers l_v on-board between any of the stops in K_v (line 2 in Algorithm 1). From the demand perspective, the route should, for all users, offer a sufficiently attractive alternative to the direct route that would be offered by an individual ride-hailing system. If this is not the case, potential users of the system will not opt for a shared ride. In other words, the user behaviour in this case is modeled with a decision rule that considers route S_v feasible only if all requests R_v and passengers P_v experience a net benefit from sharing the ride (line 2, with the net benefit as defined in Equation 2.2). If no feasible route exists, the output of the travel function, S_v^*, is "invalid" (lines 6 and 9). Otherwise, the optimal route S_v^* in K_v is defined as the feasible route with the maximum sum of the benefits for the requests r and passengers in group-vehicle combination R_v (lines 4 and 13).

Configurations

To be able to answer any of the sub research questions, some different configurations of Algorithm 1 are required, as illustrated by Table 2.3. When investigating the effect of homogeneous behavioral preferences regarding delays and sharing (*RQ 1*), only the values of the waiting time aversion α_z, delay aversion β_z and reluctance to share γ_z have to be varied. Heterogeneity in individual attributes (*RQ 2*) can be incorporated in the function by assigning a request-specific waiting time aversion α_r, delay aversion β_r and reluctance to share γ_r to each new request, rather than one generic value of α_r, β_r and γ_r to the whole set of requests. A function describing the reluctance to share a vehicle with other passengers γ_r as a function of the number of co-riders (*RQ 3*) in the vehicle can also be incorporated in the travel function. It means that for each passenger in P_v and for each new request in R_v, the maximum number of co-riders $n_p^{comax} = max(L_{v,p}) - 1$ has to be determined given route S_v. The reluctance to share for passenger p is then determined with a function that describes the relation between the maximum number of co-riders n_p^{comax} and the level of discomfort experienced: $\gamma(n_p^{comax})$. Finally, the validity of the model can be improved by incorporating that the fare discount π_r in a dynamic ride-sharing system is not fixed for all rides but dependent on the maximum number of co-riders throughout the journey (*RQ 4*): $\pi_r(n_p^{comax})$.

2.4.2 Group-vehicle assignment

Description

The matching steps (2, 3 and 5 in Figure 2.1) have shed a light on the feasibility of different group-vehicle combinations. The RGV-graph (step 6 in Figure 2.1) includes all request

groups that can be executed with at least one vehicle within the constraints (i.e. at least one R_v is not 'invalid' if all vehicles are considered to execute that request group). Next to the set of feasible request groups G, the RGV-graph includes all vehicles V and requests R. The edges between G and V store the the total user benefit b_{gv} of the optimal route S_v^* to satisfy this group-vehicle combination ($R_v \leftarrow g$) if such a route exists, as found in the RGV-matching step. The edges between R and G store information on whether a request r is part of a request group g or not, $a_{rg} = [0, 1]$. In this step, request groups are assigned to vehicles to maximize, primarily, the number of requests assigned $|R^a|$ and, secondarily, the total benefit b_{gv} of the chosen group-vehicle combinations. The algorithm described in this subsection corresponds to step 7 in Figure 2.1.

Algorithm

The group-vehicle assignment is treated as an Integer Linear Problem (ILP) with binary decision variables x_{gv} indicating whether a group-vehicle combination with total benefit b_{gv} is chosen or not. The ILP is defined as follows:

$$\max \quad \sum_{g \in G} \sum_{v \in V} (b_{gv} + M \cdot \sum_{r \in R} a_{rg}) \cdot x_{gv} \tag{2.3}$$

$$\text{s.t.} \quad \sum_{g \in G} x_{gv} \leq 1, \ \forall v \in V \tag{2.4}$$

$$\sum_{g \in G} \sum_{v \in V} a_{rg} \cdot x_{gv} \leq 1, \ \forall r \in R \tag{2.5}$$

$$x_{gv} = [0, 1], \ \forall g \in G, v \in V \tag{2.6}$$

The objective function aims at a maximum total benefit for accepted requests and passengers (first term in Equation 2.3), but prioritizes the acceptance of a maximum number of requests by adding a very large reward M for every request in an assigned request group (second term in Equation 2.3). Binary variable a_{rg} indicates whether a request r is part of request group g, meaning that $\sum_{r \in R} a_{rg}$ equals the size of a request group. Thereby, the total benefit of a group-vehicle combination gv consists of the summed net benefit for all requests and passengers in this group plus a large reward M for each request that is part of the group. The objective of the assignment, to maximize the total benefit of all group-vehicle combinations, is thus defined as in Equation 2.3.

The group-vehicle assignment may initially include group-vehicle combinations that are infeasible (as all request groups are included in G that are feasible with at least one vehicle). To prevent that an infeasible group-vehicle combination is actually chosen, a very large disbenefit is attached to such a combination. To be precise, the 'benefit' b_{gv} attached to such a combination is assumed $-M^2$, to compensate for the big M reward that is given to every assigned request in the objective function.

The Integer Linear Problem contains three types of constraints guaranteeing (1) a maximum assignment of one request group g to every vehicle v (Equation 2.4), (2) that each request r is not part of multiple assigned request groups in G (Equation 2.5), and (3) that each decision variable is binary (Equation 2.6).

The number of decision variables in this ILP is equal to the number of group-vehicle combinations, which is, at worst, the number of request groups $|G|$ multiplied by the number of vehicles m in the fleet. The number of constraints (disregarding the integer constraint for each decision variable) on the other hand is equal to the fleet size m (constraints of type in Equation 2.4) plus the number of available requests $|R|$ (constraints of type in Equation 2.5).

2.4.3 Rebalancing empty vehicles

Description

The result of the group-vehicle assignment (step 7 in Figure 2.1) may include a set of unassigned requests R^u, vehicles V^u, or both. Vehicles V^u can be left unassigned because they are already carrying passengers (V^p) and do not have much or any spare capacity to pick-up new requests in R. However, there can also be empty unassigned vehicles V^{np} left, indicating that either demand is lower than supply or that demand and supply are not properly matched. The latter is the case when at the same time there are also unassigned requests R^u in the system. It means that there is a spatial imbalance in supply and demand. The available set of vehicles V^{np} cannot currently serve any of the unassigned requests R^u, nor in the future, as else they would have been immediately assigned to serve those requests. However, these vehicles in V^{np} are able to serve future requests appearing in proximity to unassigned requests in R^u, an area where currently no vehicles or no vehicles with spare capacity are available. Empty vehicles V^{np} can thus move to areas with unserved demand to anticipate future requests. At such vehicle's current location q_v there are likely more other vehicles around, a possible reason for the fact that the vehicle has been left unassigned during the group-vehicle assignment. This subsection explains the vehicle rebalancing step, corresponding to step 8 in Figure 2.1, in which is determined which vehicles in V^{np} rebalance and what the new location of those vehicles V^{reb} will be.

Algorithm

Rebalancing efficiency is an issue as rebalancing by definition leads to extra vehicle mileage. To guarantee efficient rebalancing of empty vehicles, a similar procedure as in the study of Alonso-Mora et al. (2017) is applied. Hereby each idle vehicle in V^{reb} moves in the direction of a unique unassigned request in R^u, creating a request-vehicle assignment problem like the one in the previous step (in which B is defined as the list of potential RV-combinations). Again, this allows the assignment to be solved as an ILP, with x_{rv} as decision variable to indicate whether vehicle $v \in V^{np}$ is or is not assigned to move in the direction of unassigned request $r \in R^u$. The ILP is defined as follows:

$$\min \quad \sum_{v \in V^{np}} \sum_{r \in R^u} x_{rv} \cdot d_{rv} \tag{2.7}$$

$$\text{s.t.} \quad \sum_{r \in R^u} x_{rv} \leq 1, \ \forall v \in V^{np} \tag{2.8}$$

$$\sum_{v \in V^{np}} x_{rv} \leq 1, \ \forall r \in R^u \tag{2.9}$$

$$\sum_{v \in V^{np}} \sum_{r \in R^u} x_{rv} = \min(|R^u|, |V^{np}|) \tag{2.10}$$

$$x_{rv} = [0, 1], \ \forall r \in R^u, v \in V^{np} \tag{2.11}$$

The goal of the ILP is to minimize the total distance resulting from the rebalancing procedure, while ensuring a maximum number of vehicles to be assigned to rebalance. The objective function (Equation 2.7) aims at a minimum total rebalancing distance by minimizing the sum of the distances d_{rv} of all chosen ($x_{rv} = 1$) request-vehicle combinations. To guarantee a maximum number of assigned vehicles, an additional constraint is introduced that guarantees that either all requests or all vehicles are assigned, depending on whether more vehicles or requests exist in the assignment problem (Equation 2.10). In other words, there should be no more feasible request-vehicle combinations left after the assignment.

There are three other types of constraints that need to be considered in this ILP: (1) each vehicle can only move in the direction of one unassigned request (Equation 2.8), (2) it is not allowed for multiple vehicles to rebalance to the same unassigned request (Equation 2.9), and (3) all decision variables are binary (Equation 2.11).

The total number of decision variables in this optimization problem is equal to the number of unassigned requests $|R^u|$ multiplied by the number of empty unassigned vehicles $|V^{np}|$. The number of constraints (disregarding binary constraints) is equal to the sum of the number of vehicles (constraints of type in Equation 2.8) and the number of requests (constraints of type in Equation 2.9) plus 1 for the extra constraint maximizing the number of assigned request-vehicle combinations (Equation 2.10). So in total the rebalancing assignment problem has $|R^u| + |V^{np}| + 1$ constraints.

2.4.4 Update vehicle status and request pool

Description

After group-vehicle assignment is completed and idle vehicles have or have not been assigned to rebalance (steps 7 and 8 in Figure 2.1 respectively), a new assignment period has to be prepared. This means updating the location of vehicles V^a, V^p and V^{reb} as well as the passengers P_v on-board the first two sets of vehicles (step 1 in Figure 2.1). First, vehicle schedules are created in step 9 of Figure 2.1 based on the outcome of the group-vehicle and rebalancing assignment. This subsection describes the algorithm to update the schedules of vehicles and to determine the consequences for the vehicle status and pool of pending requests at the next assignment.

Algorithm

Algorithm 2 summarizes the procedure that is followed after each assignment to update the status of vehicles and the pool of requests for the next assignment. The following paragraphs explain the main steps of this algorithm.

The status of an assigned vehicle v^a is updated based on the request group it was assigned to ($x_{gv} = 1$) and the optimal route S_v^* that was found for that set of requests. Relevant for the next assignment are the locations $S_v{}' \in S_v^*$ that are reached by this vehicle by the time the next assignment commences. If at the time of the new assignment $t := t + \tau^{int}$ a vehicle v is not located at a stop location in Q, the first stop location that is reached after t will form the new q_v of the vehicle, together with the time t_v that this stop is reached. For this purpose, the sequence of stop locations U_v (i.e. nodes in the network) along route S_v is considered, assuming the shortest path between all individual stops $s \in S_v$.

A request r^a that is part of a request group g assigned to vehicle v is not necessarily directly picked-up by that vehicle. Such a request $r^{anp} \in R^{anp}$ will remain in the pool of requests R available for the next assignment. However, the latest pick-up time t_r^{pl} of such a request is updated to the expected pick-up based on route S_v of the vehicle that the request is assigned to. This prevents a repeated assignment of already assigned requests that have not yet been picked up, while at the same time allowing an improved assignment for requests if found. If a vehicle $v \in V^p$ is not assigned but still active, meaning that there are still passengers P_v on-board the vehicle waiting to reach their drop-off location d_p, route S_v is followed, which has been given to the vehicle the last time it was assigned. Alternatively, if the vehicle is idle, it will follow route S_v that followed from the rebalancing procedure or it will remain idle if not assigned in the rebalancing step.

After servicing location q_v, the time that this location is reached t_v and the passengers on-board vehicle P_v are updated (thereby noting that the model assumes that only existing passengers P_v can be dropped-off and not newly picked-up requests in R^a), a new pool

Algorithm 2 Update vehicle status and request batch

$t \leftarrow t + \tau^{int}$ ▷ Update time to next assignment
for all $v \in V$ **do**
  if $P_v \neq \emptyset$ **then** ▷ There are passengers on-board of v
    $U_v{}' = \emptyset$
    $t_u \leftarrow t_v$; $i \leftarrow 1$
    while $t_u \leq t$ **do**
      $u \leftarrow U(i)$
      $U_v{}' \leftarrow (u)$; $i \leftarrow i + 1$ ▷ Find stop locations reached before t
    end while
    if $t_u = t$ **then**
      $q_v \leftarrow u$; $t_v \leftarrow t$ ▷ Update vehicle location and time
      if $u \in S_v$ **then**
        $h_v \leftarrow 1$ ▷ Update whether vehicle is halting
      else
        $u \leftarrow U(i+1)$ ▷ Determine first stop location in U after t
        $q_v \leftarrow u; t_v \leftarrow t$; $h_v \leftarrow 0$ ▷ Update vehicle status
        $U_v{}' \leftarrow u$
      end if
      for all $p \in P_v$ **do**
        if $d_p \in U_v{}'$ **then**
          $P_v\ p$; $l_v \leftarrow l_v - 1$ ▷ Remove passengers
        end if
      end for
      if $v \in V^a$ **then** ▷ Vehicle is assigned to R_v in last assignment
        for all $r \in R_v$ **do**
          if $o_r \in \sigma_v{}'$ **then**
            $P_v \leftarrow \{r\}$ ▷ Pick-up new requests
            $\{o_p, d_p, t_p^r, c_p, \alpha_p, \beta_p, \gamma_p\} \leftarrow \{o_r, d_r, t_r^r, c_r, \alpha_r, \beta_r, \gamma_r\}$
            $t_p^{pu} \leftarrow t_{o_r}$; $l_v \leftarrow l_v + 1$ ▷ Update status of P_v and v
          end if
        end for
      end if
    end if
  else
    if $v \in V^{reb}$ **then** ▷ Vehicle rebalances
      $t_w \leftarrow t_v$; $i \leftarrow 1$
      while $t_w \leq t$ **do**
        $w \leftarrow W(i)$; $i \leftarrow i + 1$
      end while
      $q_v \leftarrow w$; $t_v \leftarrow t_w$ ▷ Update vehicle location and time
    else
      $t_v \leftarrow t$ ▷ Update only time
    end if
  end if
  $R \leftarrow R^{new} \cup R^{anp} \cup R^u(r|t_r^{pl} > t)$ ▷ Create new request pool
end for

of requests R is created. This pool includes requests that have been assigned but not yet picked-up R^{anp} and requests R^u that were left unassigned and are still feasible for assignment in the next phase. Unassigned requests are not feasible for assignment if they cannot reach their destination d_r following the shortest path between o_r and d_r, even when they are picked-up immediately after the next assignment. Filtering requests based on feasibility thus considers the trade-off between ride-sharing benefits and disbenefits for a specific request, as described in Section 2.4.1. New requests $R^{new} \subseteq Z$ that have been collected in the time window between the previous and new assignment are added to the pool of requests. The updated vehicle data and request pool form the input for a new assignment phase.

2.5 Key performance indicators

The potential of dynamic ride-sharing services is measured with a few Key Performance Indicators (KPIs), representing both the level of service offered to users and the efficiency for operators and authorities. This section introduces the main KPIs considered in this research.

2.5.1 Level of service

Users or potential users of public transport perceive reliability, frequency, comfort, travel time and fare level as important attributes for service quality (Bates, Polak, Jones, & Cook, 2001, Edvardsson, 1998, Hensher, Stopher, & Bullock, 2003, König & Axhausen, 2002, Friman & Gärling, 2001). A ride-sharing service resembles public transport in the sense that it also involves sharing a vehicle with strangers, a potential extra (and possibly unreliable) travel time when compared to the direct ride from a person's origin to destination and a fare paid to a service provider. Consequently, except for frequency, all attributes found important by (potential) public transport users might also be relevant for users of a ride-sharing service. The KPIs to measure Level of Service (LoS) therefore have to cover each of these attributes.

More specifically, users will desire reliability in a ride-sharing service not only in the form of reliable travel times but also in the form of a guaranteed or very probable pick-up. Ride comfort in a ride-sharing service can be an issue due to the presence of intermediate stops as well as the fact that the vehicle is shared with strangers. Both service indicators have to be measured separately. Table 2.4 gives the definitions and mathematical formulations of the KPIs that jointly give an indication of the Level of Service perceived by users of a dynamic ride-sharing service.

2.5.2 Efficiency

For service providers and authorities, improved efficiency is the main motive when deploying a ride-sharing system. When passenger share rides, an operator can utilize its fleet more efficiently, meaning that fewer vehicles are needed and less kilometers are driven. This leads to reduced investment costs as well as a reduction in operating costs following from a lower fuel consumption. From the authority's perspective, a potential reduction in total vehicle mileage is most beneficial as it can lead to reduced congestion levels, better traffic safety (Hennessy & Wiesenthal, 1999) and an improvement of the air quality (Pant & Harrison, 2013).

Two of the selected KPIs give an indication of the distance savings from ride-sharing. The vehicle movement distance (VMD) measures the total absolute distance driven by all vehicles in the fleet. The effective vehicle transportation distance ratio (EVTDR)

KPI	Attribute	Definition	Formulation
Acceptance rate	Reliability	Percentage of fulfilled (either shared or non-shared) requests out of the total demand (Ehsani et al., 2018)	$\frac{\sum_{a \in A} \lvert R^{ap} \rvert}{\lvert Z \rvert}$
Delay	Reliability, travel time	Extra travel time experienced by a user compared to the direct travel time between origin and destination	$t_r^d - t_r^*$
Relative delay	Reliability, travel time	Delay as a percentage of the direct travel time from origin to destination	$\frac{t_r^d - t_r^*}{tt_{o_r,d_r}}$
On-board load	Comfort	Maximum simultaneous number of co-riders with which a vehicle is shared	$l_v - 1$
Stop frequency	Comfort	Number of stops (pick-ups and drop-offs) experienced by a passenger while on-board of a vehicle	$\lvert S_r \rvert$

Table 2.4: KPIs representing Level of Service of a dynamic ride-sharing system

compares the minimum vehicle distance that is needed to fulfill all requests when ride-sharing is not allowed with the total distance driven with ride-sharing. A value of less than 1 indicates that a ride-sharing system leads to a reduction in the number of vehicle kilometers in the system. If the value is more than 1, more vehicle kilometers are produced with the ride-sharing service in place. It should be noted however that this indicator does not consider deadheading in a ride-hailing service and is therefore a conservative indicator of the distance saved by a ride-sharing service. In terms of the utilization of the vehicle fleet, the ratio of the distance driven by vehicles when carrying passengers to the total vehicle distance distance, the vehicle transportation distance ratio (VTDR), is a relevant indicator. Also, the ratio of the vehicle distance that empty vehicles spent rebalancing (empty vehicle rebalancing distance ratio, EVRDR) gives an indication of the efficiency of the ride-sharing service. Finally, the vehicle occupancy (VO) is measured as an indicator for efficiency.

Table 2.5 presents the definitions and pseudo-mathematical formulations of the main KPIs that jointly give an indication of the efficiency gains that a dynamic ride-sharing system can offer to operators and authorities compared to individual ride-hailing services.

2.6 Implementation

The simulation model is implemented in Python, using the open-source library Numpy (http://www.numpy.org/index.html) to enable efficient operations of large data structures in the model, such as creating and storing the edges of RGV-graphs when many requests and vehicles are considered. The Networkx package (https://networkx.github.io/) is used to compute shortest paths between two locations in the road network. The optimization problems that are part of the group-vehicle assignment and rebalancing procedure are solved with the MOSEK Optimizer API (https://docs.mosek.com/8.1/pythonapi/index.html).

KPI	Attribute	Definition	Formulation
VMD	Vehicle mileage	Total distance driven by all vehicles in the fleet (Ehsani et al., 2018)	$\sum_{v \in V} distance_v$
VTDR	Fleet utilization	Ratio of the distance driven by all vehicles in the fleet while carrying passengers to the total vehicle movement distance (Ehsani et al., 2018)	$\frac{\sum_{v \in V} (dist\ with\ pax)_v}{VMD}$
EVTDR	Vehicle mileage	Ratio of the minimum vehicle distance needed to fulfill all requests without ride-sharing (i.e. sum of the shortest path between all accepted requests' origins and destinations) to the total vehicle movement distance (Ehsani et al., 2018)	$\frac{\sum_{r \in R^a} distance_{o_r,d_r}}{VMD}$
EVRDR	Fleet utilization	Ratio of total empty vehicle distance that vehicles spent rebalancing, as opposed to deadheading to new requests	$\frac{\sum_{v \in V^{reb}} distance_v}{\sum_{v \in V^{np}} distance_v}$
VO	Fleet utilization	Number of passengers on-board of a vehicle while in revenue service (Ehsani et al., 2018)	$\frac{\sum_{k \in \{1,\dots,\kappa\}} time_k \cdot k}{\sum_{k \in \{1,\dots,\kappa\}} time_k}$

Table 2.5: KPIs representing the efficiency of a dynamic ride-sharing system

Chapter 3

Experimental design

In this research, the model is applied to an artificial rather than to a real-world case. There are two reasons for this. Firstly, it gives us the freedom to test the effect of different spatial distributions of demand, a requirement for answering sub research question 5, and secondly, by simplifying the network and total demand, the computational complexity of this study is limited. A downside of this choice is that the results of this research cannot directly be compared to the results of previous studies to find out how the incorporation of a user cost-benefit trade-off in ride-sharing choice improves the estimated potential of ride-sharing services.

Section 3.1 describes and motivates the general set-up of the designed experiment, while in Section 3.2 fifteen different scenarios are introduced for which the model is run. The scenarios vary based on the five different variables that are tested in this study: the distribution of demand, waiting time aversion α_z, delay aversion β_z, reluctance to share γ_z and discount percentage π_z. By comparing the results of the different scenarios, answers to the different sub research questions can be formulated.

3.1 Set-up

In Table 3.1, the main set-up of the experiment in terms of the simulation characteristics, total demand, fleet properties, ride fares and network characteristics that apply for all scenarios, is summarized. This section describes the set-up of the experiment in more detail and explains why was chosen for this specific set-up.

The total duration of the simulation is 2 hours. This time window does not require an excessive computational effort, but is still long enough for vehicles to serve a significant number of requests, leading to representative results. Request groups are assigned to

Type	**Parameter**	**Value**	**Unit**
Simulation time	τ^{sim}	2	h
	τ^{int}	1	min
Network	$\|Q\|$	121	nodes
	d_e	500	m
	tt_e	50	s
Demand	E($\|Z\|$)	1210	requests/h
Fleet	μ	150	vehicles
	κ	3	seats
Ride fare	c_{base}	3	euro
	c_{km}	2	euro

Table 3.1: Parameter values in numerical experiment

vehicles every minute, thus 120 times in total. An additional warm-up period of 15 minutes applies to make sure the impact of the starting locations of vehicles, the initial absence of passengers in vehicles and any other starting conditions do not influence system output too much.

Ride-sharing in this study is considered predominantly as an alternative to individual ride-hailing. As the adoption of ride-hailing services is so far limited mainly to urban areas (Clewlow & Mishra, 2017), it is sensible to simulate the potential of ride-sharing services in such a context. Therefore, the numerical experiment to which the model is applied will represent an urban context. More specifically, a grid network of size 10 x 10 is assumed. Request locations (both origins and destinations) are limited to the 121 nodes (i.e. intersections) in the network. To mimic the scale of a moderately sized city, the distance of all edges has been set to 500 meters. This means the maximum travel distance in the grid is 10 kilometers (in other words the total length of the roads comprising the circumference of the grid is 20 kilometers) as well as that the grid covers a surface area of 25 km^2. With that size, the spatial scale of the numerical experiment experiment is somewhat similar to the area inside of the Ring Road of Amsterdam, with a surface area of 27 km^2, or the area inside the Inner Ring of Berlin, with a surface area of 21 km^2 (determined with https://www.daftlogic.com/projects-google-maps-area-calculator-tool.htm#).

It should be noted that the number of nodes in the network is small. A similar study on ride-sharing potential in Manhattan by Alonso-Mora et al., 2017, for example, considered the complete set of nodes in the respective road network: a total of 4.092 (distributed over a surface area that is approximately twice as large as Amsterdam). The ride-sharing system that is investigated in this study is thus one that does not allow for pick-ups at all intersections in the road network, but rather one where pick-ups are limited to dedicated pick-up locations at some of the intersections. It is assumed in the model that there is no additional travel time related to moving / walking to such a pick-up location, as users can do so while the vehicle itself is moving to the stop location. In other words, the waiting time wt_z for a user also comprises the stop access time.

As an urban environment is modeled, traffic speeds on the road network are relatively low. In practice, the average traffic speed of major European cities in 2017 varies mostly between 20 and 40 km/h, with Milan as one of the main positive outliers (55 km/h) and Lisbon as one of the negative outliers (15 km/h) (Kfzteile24, 2019). For the sake of simplicity, a traffic speed of 36 km/h or 10 m/s is assumed in the numerical experiment, meaning the travel time on every edge in the network is 50 seconds exactly. Compared to the average European city, the average traffic speed in the numerical experiment is thus relatively (but not extremely) high.

The total demand for trips is assumed to be 1,210 requests per hour, an average of 10 requests per hour per node. The way demand is spatially distributed over the stop locations is scenario-specific (as clarified in Section 3.2). Every request z with an origin o_z and destination d_z is associated with a request time t_z^r, which is generated by sampling from an exponential distribution based on the expected interval λ between two requests with this specific origin and destination, which followed from the demand distribution.

The fleet of the investigated ride-sharing service consists of 150 vehicles. This is based on the average number of requests per vehicle per hour that is applied in the study of Alonso-Mora et al., 2017. This work shows that approximately 92% of the requests can be shared given 2.000 vehicles with a capacity of 4. As demand per hour in New York in the studied week is approximately 15 times the value assumed in this study, the corresponding number of vehicles to serve the demand in this study is roughly 130. In this study, however, vehicle capacity is limited to 3 to minimize the total computational effort needed to generate the KPIs for each of the scenarios. To compensate for the fact that vehicles have a capacity of 3 instead of 4, the fleet size is set at 150. Vehicles are initialized at 25

evenly distributed locations in the network.

Ride fares are determined for each request based on a base fee and an additional kilometer fee. In the numerical experiment, the regulated maximum taxi fare for 2019 for the city of Amsterdam is used (rounded down): a base fee of €3 and a kilometer fee of €2 (Gemeente Amsterdam, 2019).

3.2 Scenario design

To find how behavioral attributes, service design and demand distribution influence ride-sharing efficiency and level of service, different scenarios are run with different model inputs. In terms of the demand, a uniform, moderately directed and strongly directed distribution is tested. Three behavioral attributes are varied: waiting time aversion α_z, delay aversion β_z and reluctance to share γ_z. Finally, the effect of the discount percentage π_z is tested with two scenarios. This section describes the specification of each of the scenarios, starting with a 'base' scenario (scenario 1 in Table 3.2) that contains the base value for each parameter value as well as the demand distribution that is used throughout most scenarios. The base scenario forms a reference for the other scenarios (scenarios 2-15 in Table 3.2), allowing to check the effect of a change in a single variable (parameter values or demand distribution) on the different ride-sharing KPIs that have been defined. Table 3.2 summarizes the scenario design and points out which scenarios are needed to answer which sub research question. Research question 6 considers an overarching analysis of all results and is therefore not included in the table.

#	Demand	β_z (€/h)	γ_z (€)	π_z (%)	Acronym	RQ
1	Uniform	30	3	50	U_30_3_50	1-5
2	Uniform	18	3	50	U_18_3_50	1
3	Uniform	24	3	50	U_24_3_50	1
4	Uniform	36	3	50	U_36_3_50	1
5	Uniform	42	3	50	U_42_3_50	1
6	Uniform	30	1	50	U_30_1_50	1
7	Uniform	30	2	50	U_30_2_50	1
8	Uniform	30	4	50	U_30_4_50	1
9	Uniform	30	5	50	U_30_5_50	1
10	Uniform	N(30,10)	3	50	U_H_3_50	2
11	Uniform	30	N(3,2)	50	U_30_H_50	2
12	Uniform	30	$2 + n_z^{comax}$	50	U_30_D_50	3
13	Uniform	30	5	$50 + 7.5 \cdot n_z^{comax}$	U_30_5_D	4
14	Moderately directed	30	5	50	MD_30_5_50	5
15	Strongly directed	30	5	50	SD_30_5_50	5

Table 3.2: Specification of all fifteen scenarios (with α_z always half of β_z) and the sub research questions they help to answer

3.2.1 Base scenario

First of all, the base values for the behavioral attributes α_z, β_z and γ_z are specified. For the delay aversion β_z, the value of travel time reliability (VOR) of a transit service rather than the value of travel time savings (VOT) might be a good estimate. With delay being defined as the travel time beyond what is planned, it is perceived more negatively than standard in-vehicle time (Cats, West, & Eliasson, 2016). An empirical study on the variability of travel time by Börjesson, Eliasson, and Franklin (2012) even suggests a delay time multiplier of 3.5. Another study on travel time reliability for commuters in Barcelona (Asensio & Matas, 2008) found a VOR of €34.4/h, while a similar study

for Australia concluded a mean value for travel time reliability of approximately €33/h (Li, Hensher, & Rose, 2010). Based on those values, in the base scenario of this numerical experiment, a (rounded) delay aversion β_z of €30/h is assumed. It is likely that users with a higher delay aversion also have a higher additional waiting time aversion. Therefore, like the study on ride-sharing potential with high-capacity vehicles in New York (Alonso-Mora et al., 2017), this research assumes that the additional waiting time aversion is always half the delay aversion, or in other words $a_z = \frac{1}{2} \cdot \beta_z$. This applies to all scenarios. The idea behind this assumption is that waiting time is valued more negatively than in-vehicle time, as a result of a lack of comfort, the inability to be productive at a stop location and psychological burdens related to waiting (Wardman, 2004).

In literature little is known about users' reluctance (or willingness) to share a car with other users, and how this attribute depends on the vehicle occupancy. Studies on crowding are limited to larger transit vehicles, such as buses, metros and trams. The focus in these studies is commonly on the effect of seating probability and the density of standees in the vehicle. This is not useful for the ride-sharing case, where seats are guaranteed and discomfort follows mainly from the social interaction between users. The base value allocated to the reluctance to share γ_z is set to €3, independent of the number of passengers in the vehicle. In this way, the reluctance to share in the base scenario describes the probability of having to share rather than the fact that a vehicle is actually shared. In terms of the disbenefit, a reluctance to share γ_z of €3 compares to an in-vehicle delay of six minutes given an average delay aversion $\beta_z = €30/h$ or to a waiting time of four minutes. If immediately picked-up and served with a direct route, users obviously require a fare discount of at least €3 to accept a shared ride given the assumed (base scenario) willingness to share. Like the willingness to share, the fare discount offered to ride-sharing users is assumed dependent on the risk of sharing instead of on the actual experience in the vehicle. To keep it simple, the base scenario assumes that every request accepted by a vehicle in the ride-sharing system gets a 50% fare discount, i.e. using a shared service is half the price of using the individual service.

Ride-sharing in this study is considered as an alternative mode of transport to individual ride-hailing in an urban context. A demand pattern in which such a service might thrive is one where demand is uniformly distributed over the network, as conventional pubic transit is relatively unattractive in such a case (Daganzo, 2010). The base scenario of the numerical experiment will therefore mimic a pattern in which trip request origins as well as destinations are spread uniformly over the stop locations in the network (i.e. the expected production and attraction in a node is equal, in this case 10 requests per node per hour). OD-pairs are then created based on the shortest travel distance for that pair. Trips with a ride distance of 2 kilometers or shorter are excluded as such rides are uncommon (Liu, Krishnakumari, & Cats, 2019) as well as undesirable in the context of a ride-sharing service. The gravity model (Erlander & Stewart, 1990) is applied:

$$T_{ij} = K_i K_j T_i T_j f(C_{ij})$$

The function $f(C_{ij})$ is applied in the form of a simple distance decay factor: $1/dist_{ij}$, in which $dist_{ij}$ is the distance of the shortest path between origin i and destination j. K_i and K_j are balancing factors to equalize total trip production T_i and attraction T_j if all OD-pairs are considered. As described in Section 3.1, requests are generated by drawing origin-destination combinations from an exponential distribution defined by the expected value between two requests with the respective origin and destination. This is done once, meaning that all scenarios with a specific demand distribution assess the shareability potential based on the same list of requests.

3.2.2 Homogeneous preferences

To find the extent to which homogeneous behavioral preferences influence ride-sharing efficiency and level of service (*RQ 1*), four values of β_z and γ_z (and consequently also α_z) are specified different than the base value. For β_z (scenario 2-5 in Table 3.2), the values range equally between 18 and 42 €/h, centered around the base value of €30/h. To account for the uncertainty in users' reluctance to share in a ride-sharing system, given the lack of empirical underpinning, the γ_z values that are tested (scenarios 6-9) cover a large range: from €1 to €5 with an average base value of €3.

3.2.3 Heterogeneous preferences

Taste heterogeneity (*RQ 2*) is introduced by assigning different values of α_z, β_z and γ_z to users / requests. As mentioned previously, research findings on such values are scarce, especially if heterogeneity across the population also needs to be incorporated. One of the few studies that is devoted to this shows that for Shanghai commuters that have an average Value of Travel Time Reliability (VOR) of 88.3 Chinese yuan per hour, the standard deviation of the VOR across respondents measures 23.9 yuan/h (Gao, Sun, Tu, & Li, 2018). In other words, the standard deviation is approximately one third of the average VOR. We thereby adopt the ratio from the Shanghai case to model heterogeneity in users' delay aversion. Therefore, heterogeneity in delay aversion (scenario 10 in Table 3.2) is captured by a normal distribution with mean €30/h and standard deviation €10/h: $N(30, 10)$.

When it comes to preferences related to fellow passengers in ride-sharing, it has been found that heterogeneity is more pronounced than in traditional preferences such as travel time (Zhang & Zhao, 2018). A possible explanation for this is that some users enjoy the social interactions that come along with sharing a ride while others prefer to ride individually. It means that for some users in fact the reluctance to share γ_z might be negative, or in other words that the fact that a vehicle is shared induces an additional benefit next to a reduction in cost. Again, a normal distribution is assumed to capture heterogeneity. As heterogeneity in reluctance to share γ_z was found to be (likely) larger than heterogeneity in delay aversion β_z (with its assumed standard deviation of one third of the mean), the standard deviation relative to the mean to describe heterogeneity in the reluctance to share γ_z is higher. With a mean reluctance to share of €3, the standard deviation assumed is €2: $N(3, 2)$.

3.2.4 Occupancy-dependent reluctance to share

Another scenarios involves the dependency of a user's reluctance to share γ_z on the maximum number of simultaneous co-riders with whom the user has to share its ride (*RQ 3*), as opposed to a fixed value for γ_z in the base scenario. Scenario 12 assumes a linear relationship in the additional discomfort, meaning that the addition of an extra co-rider in the vehicle always yields the same additional disbenefit independent of the number of passengers currently in the vehicle. As, compared to the base scenario, the actual ride rather than the risk of having to share is perceived as more important, users have a lower reluctance to share, γ_z = €2, if the ride is executed fully privately. The reluctance to share given rides with a maximum of 1 and 2 co-riders is set to respectively €3 and €4 (Figure 3.1).

3.2.5 Occupancy-dependent discount

It is likely that ride-sharing with fellow passengers becomes less attractive when the reluctance to share increases with the (maximum) number of co-riders in the vehicle. This is

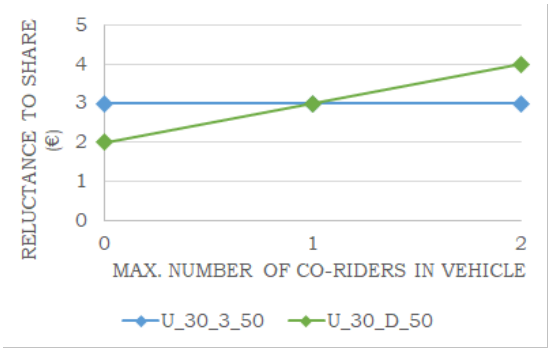


Figure 3.1: Reluctance to share γ_z as function of the number of co-riders

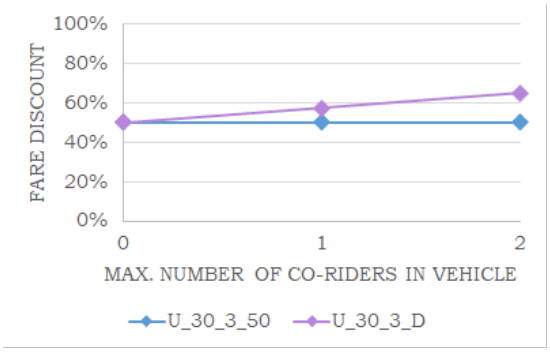


Figure 3.2: Discount percentage π_z as function of the number of co-riders

relevant as it may lead to a ride-sharing service being used very inefficiently with limited sharing and a low average vehicle occupancy. The question is also what happens to the adoption of such a service when the opposite applies: if there is an additional stimulant to ride-share with (as many as possible) other passengers in the vehicle (*RQ 4*). Scenario 13 is designed to model such a case by incorporating a fare discount π_z that is dependent on the maximum number of simultaneous co-riders as experienced by a passenger. The base scenario assumes a 50% fare discount for all users of a ride-sharing service. Alternative scenario 13 offers a similar base discount but also an additional 7.5% discount on the fare for every co-rider on-board of the vehicle in the ride segment with the largest number of co-riders (Figure 3.2).

3.2.6 Demand distribution

All of the previous scenarios assume a demand distribution that is uniform. Such a demand distribution, in which production and attraction at each node is equal, is uncommon as specific travel patterns follow from the spatial distribution of activities in a geographical area. The final two scenarios test a concentrated demand pattern, with a relatively high production in the outer nodes in the network and a relatively high attraction in the central nodes (*RQ 5*). This mimics a morning peak in which commuters travel from suburbs to the center of a city. Ride-sharing benefits can potentially be larger as users can be matched more efficiently to users moving in the same direction. On the other hand, more rebalancing might be required as vehicles will drop-off passengers mostly in areas with a low demand. Scenario 14 is designed to give insights into the ride-sharing efficiency and level of service when the direction in demand is moderate, scenario 15 when the direction in demand is strong. In Appendix B the magnitude of attraction and production in each of the nodes in the network is specified for both scenarios. Like the base scenario, the expected total number of requests per hour in scenarios 14 and 15 is 1,210.

Chapter 4

Results

This chapter includes an analysis of the indicators that follow from the application of the model to different scenarios in the experiment. The analysis is divided into seven subsections, of which the first six are dedicated to each one sub research question, while the seventh subsection focuses on the computational complexity of the methodology. The analyses related to the first five sub research questions are split into three parts, based on the main performance indicators that apply: the acceptance rate of a ride-sharing service, the average delay for users (both indicating level of service) and the total distance of vehicles in the fleet (indicating efficiency).

Specifically, in Section 4.1 the impact of a variation in ride-sharing attitudes β_z and γ_z is investigated given that those attitudes are homogeneous (*RQ 1*). Section 4.2 sheds a light on ride-sharing potential if the latter assumption is dropped and delay aversion β_z and reluctance to share γ_z are assumed request dependent (*RQ 2*). The findings in the first two sections are subject to the assumption that γ_z is not dependent on the vehicle occupancy. Section 4.3 elaborates on how ride-sharing potential is affected if users are found to be less willing to share if the vehicle occupancy is high (*RQ 3*). Section 4.4 considers the effect of an occupancy-dependent discount π_z rather than a fixed discount for all rides in a ride-sharing scheme (*RQ 4*). Section 4.5 clarifies how ride-sharing efficiency and level of service depend on the spatial distribution of demand in the network (*RQ 5*) by comparing three examples: a perfectly uniform demand pattern, one with moderately concentrated flows and one with strongly concentrated flows . Section 4.6 contains a scenario-overarching analysis that examines the correlation between different level of service and efficiency KPIs (*RQ 6*) to find if different objectives are competitive or whether a system optimal scenario exists. Finally, Section 4.7 gives an insight in the computational effort of the followed approach. Note that the comprehensive list of results is found in Appendix A.

4.1 Homogeneous preferences

This subsection shows the effect of delay aversion β_z and reluctance to share γ_z on ride-sharing efficiency and level of service if is assumed that both attributes are homogeneous across the set of potential users. Hence, the results can help to formulate an answer to sub research question 1.

Acceptance rate

A clear relationship is found between the delay aversion β_z of potential ride-sharing users and the acceptance rate of such a service (Figure 4.1a). The acceptance rate decreases approximately linearly with an increasing value of β_z, albeit gradually with an acceptance rate of 88.2% in the 'best-case' scenario ($\beta_z = €18/h$) and 64.2% in the 'worst-case'

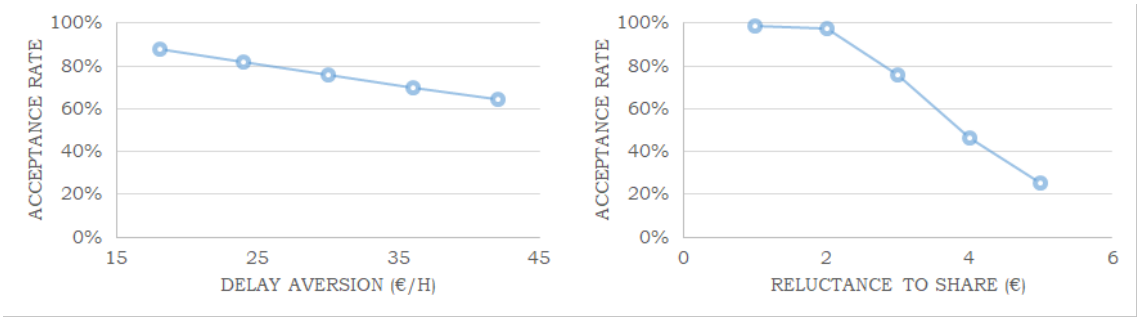


Figure 4.1: Acceptance rate as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z

scenario ($\beta_z = €42$). The relation between the reluctance to share γ_z and the acceptance rate is different from the relationship between the delay aversion β_z and the acceptance rate in two aspects: (1) the acceptance rate is more sensitive to high values of γ_z than to high values of β_z (for $\gamma_z = €5$ the adoption rate drops to 25.4%) and (2) for low values of γ_z the relation is not linear but flattens out to reach the 100%-asymptote. It means that if the willingness to share is high, nearly all users find a ride in which the fare discount outweighs the sharing inconveniences. This is the case for a reluctance to share γ_z of €2 or lower.

User delay

A negative relation exists between the delay aversion of ride-sharing users (if homogeneous) and the average experienced delay. This effect is found to be less prominent when the delay aversion β_z is relatively high, as shown in Figure 4.2a. In this case, the average in-vehicle delay is very limited and delays are mostly caused by a waiting time before pick-up. A lack of in-vehicle delay is an indicator for either infrequent or very efficient detouring in the sharing system. The findings in Figure 4.3a suggest that the former is true, as the ratio of passenger time with at least one co-rider is limited to less than 40% when $\beta_z = €42$. Only 7.8% of the passenger time is spent with the maximum number of two co-riders. The average vehicle occupancy (Figure 4.4a) in such a case does not exceed 1.2. In other words, it can be stated that if the delay aversion is high, a large percentage of passenger time is spent travelling alone, which means the shortest path between this passenger's origin and destination can be followed and users do not suffer from an in-vehicle delay. If the delay tolerance β_z is found to be higher, the average vehicle occupancy increases (Figure 4.4a), which leads to more detouring and a larger total delay (Figure 4.2a). At the same time, the average waiting time per passenger increases almost equally to the in-vehicle delay when the delay aversion β_z decreases (Figure 4.2a). Next to more detouring, it thus also takes more time on average for a vehicle with spare capacity to reach a certain request location if the delay aversion is low. This can also be explained by an increase in the average vehicle occupancy, as it is more likely that passengers first have to be dropped off before a vehicle can pick up new requests. Another explanation for a long average waiting time is a lack of rebalancing when the delay aversion β_z is low (Figure 4.8a), which means that vehicles are less likely to be located in the surrounding of an incoming request.

When considering the effect of the reluctance to share γ_z instead of the delay aversion β_z, similar effects are found, although a bit more pronounced. For the most negative sharing setting, one with a reluctance to share γ_z of €5, the average vehicle occupancy (considering only the vehicle time with passengers on-board) is even lower than in the worst-case scenario for β_z. When vehicles are in revenue mode (i.e. not deadheading) the average occupancy of vehicles is only 1.14 (Figure 4.4b), and 87% of this time is spent

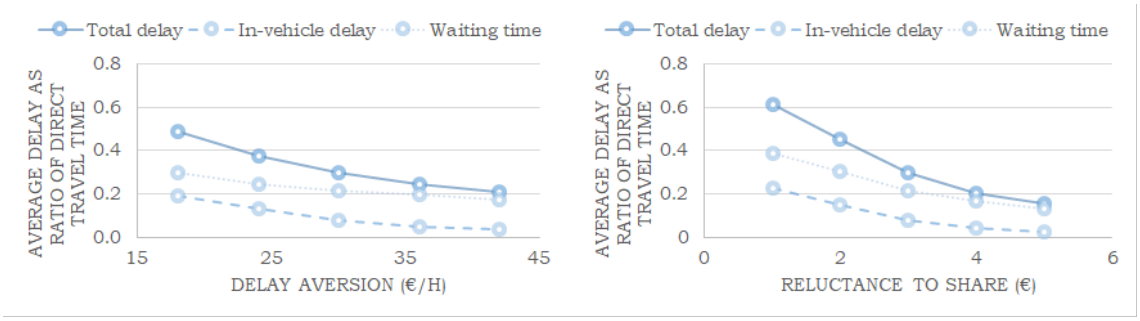


Figure 4.2: Average delay as ratio of the direct travel time, as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z

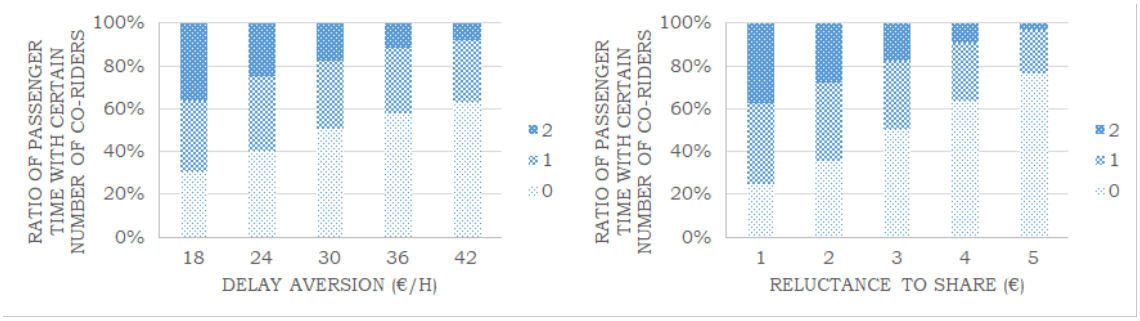


Figure 4.3: Ratio of passenger time with specific number of co-riders on-board as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z

with only one passenger on-board. In fact, only 1.1% of the (revenue) time vehicles are driving at full capacity. Such a scenario resembles a ride-hailing service in which direct transport is offered and the vehicle occupancy is close to 1. On the other hand, if users' reluctance to share appears to be much lower, the average vehicle occupancy can become as high as 1.78 ($\gamma_z = 1$). As shown in Figure 4.3b, 75.1% of the total in-vehicle passenger time is then spent with at least one other passenger, of which about half of it is spent at full capacity, i.e. with two co-riders.

The interesting part here is not merely the direction of the relation, but mainly the magnitude at which these effects take place. While the acceptance rate as a function of β_z or γ_z seems to follow a linear pattern, in which the marginal increase in acceptance rate is independent of the current value of the delay aversion or reluctance to share (unless the acceptance rate is close to 100%), the effect of behavioral preferences on passenger delays and vehicle occupancy is non-linear. The average in-vehicle delay, average waiting

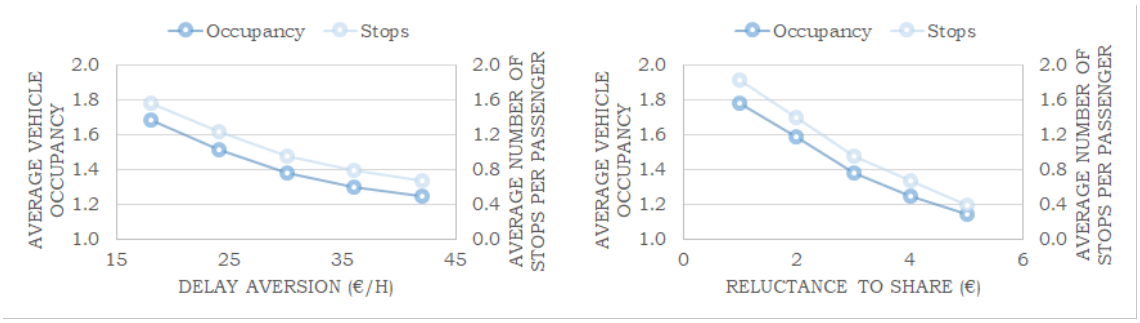


Figure 4.4: Average vehicle occupancy and number of intermediate stops per passenger as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z

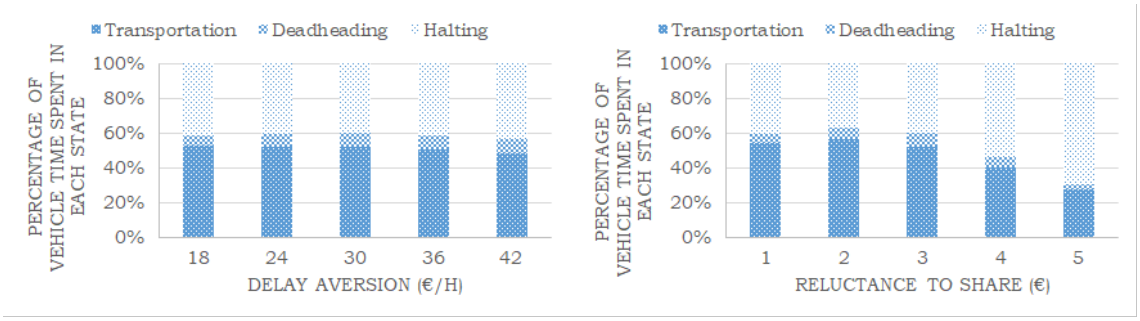


Figure 4.5: Distribution of vehicle time in each state as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z

time (both represented in Figure 4.2), the ratio of passenger time at full capacity (Figure 4.3) and the average vehicle occupancy (Figure 4.4) are all more sensitive to a change in user ride-sharing preferences when those preferences are favorable and acceptance rates are high. In such a situation, the share of empty / halting vehicles is relatively low (Figure 4.5b) and additional requests are more likely to be picked-up by a vehicle that is already carrying passengers. Contrarily, when the acceptance rate is low, additional requests can be picked-up by otherwise empty vehicles that are available throughout the network (Figure 4.5). To summarize, an increase in the number of accepted requests generally leads to an increase in the number of vehicles in revenue mode when the acceptance rate was low, and to a higher utilization of vehicles in revenue mode when the acceptance rate was already high. The average passenger delays (Figure 4.2) and number of intermediate stops as a function of the delay aversion β_z and reluctance to share γ_z (Figure 4.4) also follow the trend of the vehicle occupancy. A higher vehicle occupancy means a vehicle is shared with more co-riders that need to be picked-up and dropped-off as well. The average number of stops ranges from 0.67 per passenger to 1.56 for the range of input values of β_z and from 0.39 to 1.83 for γ_z.

Vehicle distance

One of the main potential societal benefits of a ride-sharing service is a reduction in the total vehicle distance, as indeed shown by Figure 4.6. A clear difference is noticed here between the effect of the delay aversion β_z and the reluctance to share γ_z. For example, depending on users' reluctance to share, the total distance driven by all vehicles in the fleet (VMD in Figure 4.6) can range between 3,321 ($\gamma_z = €5$) and 6,853 kilometers ($\gamma_z = €2$). Interestingly, the maximum total distance driven by vehicles in the fleet is not found at the lowest value for γ_z, but at $\gamma_z = €2$. When the reluctance to share is higher than €2, fewer requests are assigned and a larger proportion of the vehicle time is spent halting (Figure 4.5). Consequently, the total vehicle movement distance decreases. When the reluctance to share is lower (€1) on the other hand, slightly more requests are assigned, but still the total vehicle distance drops to 6,389 kilometers (Figure 4.6). What can be deduced is that if the reluctance to share is very low, requests can be assigned more efficiently to vehicles. This results from the fact that more feasible assignment outcomes are possible when users are flexible in terms of other passengers joining their ride.

Given a reluctance to share a ride of only €1, every ride-sharing vehicle kilometer effectively satisfies 1.36 (direct) passenger kilometers (Figure 4.7), thereby leading to a total distance saving of 2,277 kilometers compared to an individual ride-hailing service. The actual distance saving will be even larger as individual ride-hailing, like ride-sharing, will also require deadheading. Therefore, it is possibly more representative to consider the ride-sharing vehicle transportation distance (VTD), in which deadheading is excluded. A

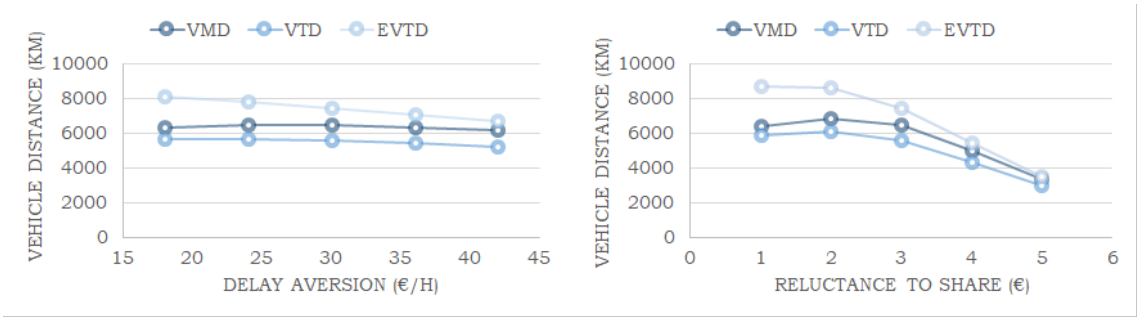


Figure 4.6: Total vehicle movement, transportation movement and effective transportation movement as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z

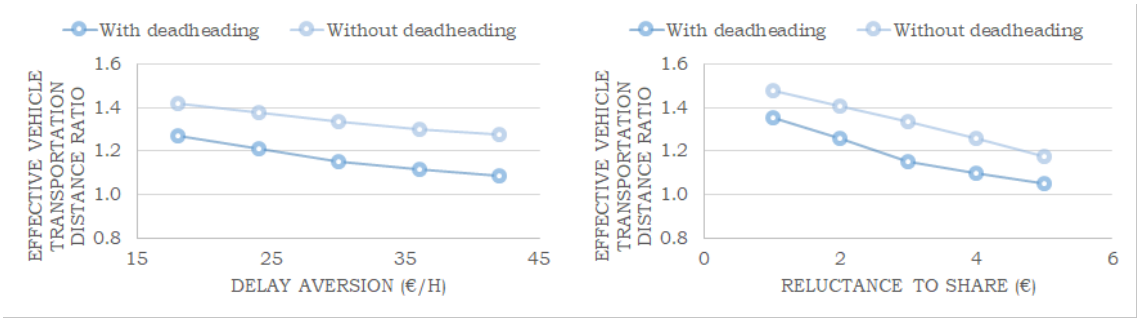


Figure 4.7: Effective vehicle transportation distance as ratio of the total vehicle distance, including and excluding deadheading, as a function of (a) delay aversion β_z and (b) reluctance to share γ_z

transportation kilometer in a ride-sharing service is then found to be equivalent to 1.48 kilometers in a ride-hailing system (saving nearly 2,800 vehicle kilometers). The most relevant finding is that the effective vehicle transportation distance ratio (*EVTDR*, Figure 4.7), an indicator for the distance that can be saved, is strongly dependent on users' ride-sharing tolerancy. For the highest input value of γ_z the ratio drops to 1.05, meaning there is only a minimal reduction in the number of vehicle kilometers (160).

The effect of the delay aversion β_z on vehicle distance (savings) is similar as the effect of the reluctance to share γ_z, although somewhat less apparent. The effective transportation distance ratio ranges from 1.09 (530 saved kilometers) for the lowest delay tolerance ($\beta_z = €42/h$) to 1.27 (1,719 saved kilometers) for the highest delay tolerance ($\beta_z = €18/h$). Again, given that a conventional ride-hailing service also requires deadheading, the actual distance savings will be higher.

Figure 4.8 shows that vehicles deadhead mostly to access the pick-up location of requests that they are assigned to. When examining the effect of γ_z on the total access distance (Figure 4.8b), it is found that the access distance is highest for an average value of γ_z. If the reluctance to share is higher, there are simply fewer requests to be picked-up. However, if potential users are more tolerant and there are more assigned requests, the total access distance also drops. This indicates that a higher acceptance rate can not only lead to more efficient routing, as demonstrated earlier, but also to more efficient pick-ups. A plausible explanation for this is that requests are more likely to be picked-up by 'active' vehicles on their way to dropping off previously accepted requests. There are simply fewer empty vehicles that will deadhead to pick-up new requests.

Figure 4.8 also illustrates that the deadheading distance resulting from rebalancing, as opposed to picking up requests, is relatively limited. Although not of the same order, the

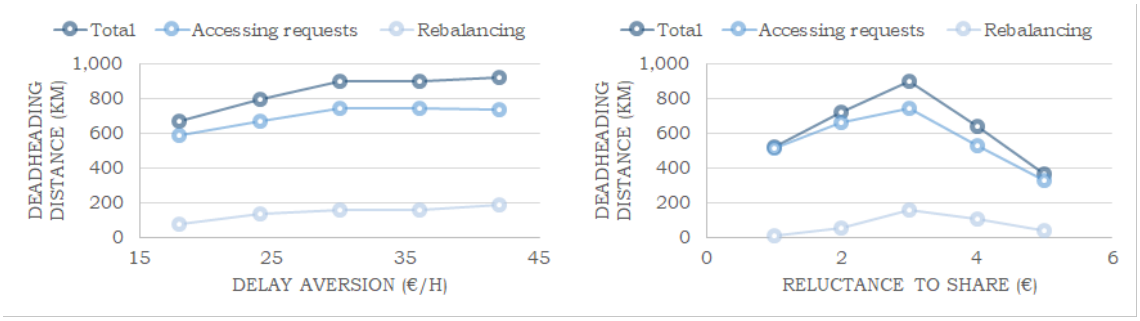


Figure 4.8: Total deadheading distance as a function of a homogeneous (a) delay aversion β_z and (b) reluctance to share γ_z

shape of the rebalancing distance function is fairly similar to the access distance function, with a maximum at $\gamma_z = €3$. For low values of the reluctance to share γ_z, there are simply fewer unassigned requests to rebalance to. For a high reluctance to share on the other hand, there are many unassigned requests left that vehicles can rebalance to as well as many rebalancing vehicles, meaning more rebalancing is expected. However, as illustrated by Figure 4.8, the rebalancing distance is found to decrease. An explanation for this is that when the reluctance to share is high, it is more likely that an unassigned request and an unassigned vehicle are located at the same stop, with the vehicle unable to serve the request even without detouring (as the initial burden of opting for a shared service for a request, a disbenefit of $\gamma_z = €5$, is very high). In other words, there is no spatial imbalance in demand and supply, meaning that unassigned vehicles are more likely to rebalance to requests at or around their own location and the total rebalancing distance will drop. When analyzing the effect of β_z instead (Figure 4.8a), the rebalancing distance is found to actually increase if the tolerance is low. In this case, the reluctance to share is lower and it is less likely that requests at a specific location cannot be served by an empty vehicle at or around that same location.

To summarize, in this section it was found that if delay aversion β_z increases from 18 to 42 €/h:

- The acceptance rate decreases approximately linearly from 88.2% and 64.2%.
- The average vehicle occupancy decreases from 1.68 to 1.24. The decrease is more significant if the delay aversion is low.
- The average user delay decreases corresponding to the vehicle occupancy, from 48.7% of the direct travel time to 21.1%.
- The effective vehicle transportation distance ratio decreases from 1.27 to 1.09.

Furthermore, it was found that if the reluctance to share γ_z increases from 1 to 5 €:

- The acceptance rate decreases from 98.8% and 25.4%. The decrease is approximately linear as long as the acceptance rate is not close to 100%.
- The average vehicle occupancy decreases from 1.78 to 1.14. The decrease is more significant if the delay aversion is low.
- The average user delay decreases corresponding to the vehicle occupancy, from 61.3% to 15.2% of the direct travel time.
- The effective vehicle transportation distance ratio decreases from 1.36 to 1.05.

4.2 Heterogeneous preferences

This section contains the analysis of the effect of heterogeneity in sharing preferences on ride-sharing level of service (also accounting for equity) and efficiency (*RQ 2*), assuming that the objective function is the same as when preferences are homogeneous. The analysis consists of a comparison of the result of three scenarios: one with no heterogeneity (*U_30_3_50*), one with heterogeneity only in delay aversion β_z (*U_H_3_50*) and one with heterogeneity only in the reluctance to share γ_z (*U_30_H_50*).

Acceptance rate

The number of requests that are accepted by the assumed ride-sharing service is very similar for scenarios *U_30_3_50* and *U_H_3_50*: 76.0% versus 75.9% of the total demand (Figure 4.9). However, if heterogeneity in delay aversion exists, the acceptance rate is nearly 10% lower. Figure 4.10 shows that mainly requests with a high reluctance to share are denied, as the share of users with a high reluctance to share in the pool of accepted requests is lower than in the total pool of requests. The drop in acceptance rate can thus be explained by the fact that a relatively large share of users with a reluctance to share γ_z of more than €3 are not accepted. For users with an above average delay aversion β_z (instead of γ_z) this effect is much smaller, as shown by Figure 4.11.

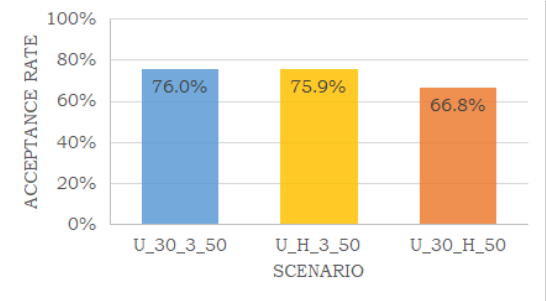


Figure 4.9: Effect of heterogeneity on the acceptance rate

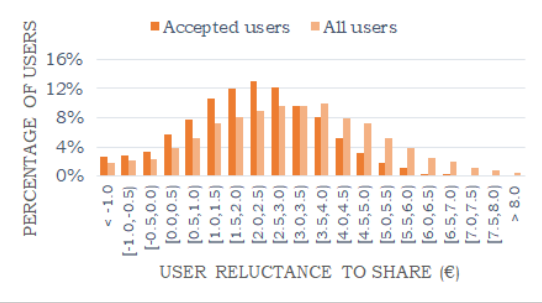


Figure 4.10: Reluctance to share of accepted users versus reluctance to share of all requests in *U_30_H_50*

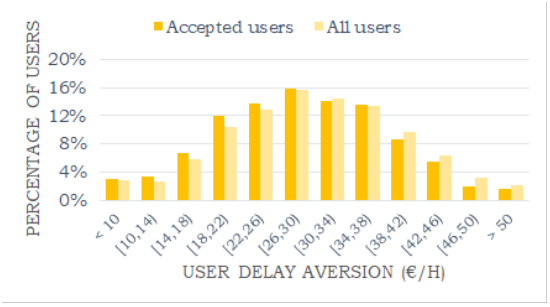


Figure 4.11: Delay aversion of accepted users versus the delay aversion of all requests in *U_H_3_50*

User delay

It is found that both heterogeneity in delay aversion and reluctance to share lead to a higher average delay per passenger (Figure 4.12), respectively 36.4% (154.2 seconds) and 42.1% (162.8 seconds) of the direct travel time, compared to 29.5% for scenario without heterogeneity. Figure 4.13 shows that relatively long delays are more common in both

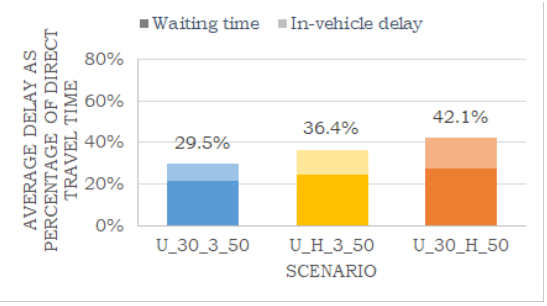


Figure 4.12: Effect of heterogeneity on average delay per passenger

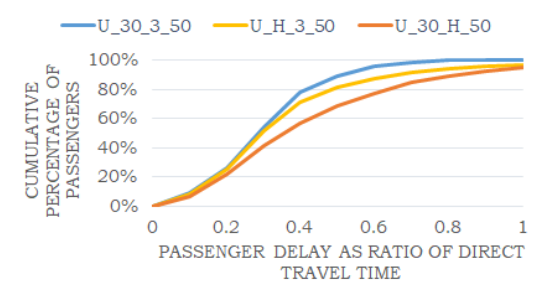


Figure 4.13: Effect of heterogeneity on the distribution of delays

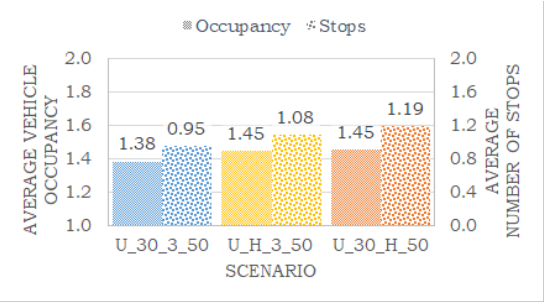


Figure 4.14: Effect of heterogeneity on average vehicle occupancy in revenue mode and average number of intermediate stops

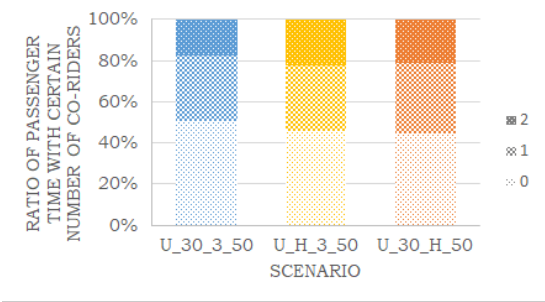


Figure 4.15: Effect of heterogeneity on passenger time with different number of co-riders

heterogeneous scenarios, especially in *U_30_H_50*. If there is no heterogeneity in user preferences, only 1.1% of all users experience a delay of more than 70% of their direct travel time. If β_z is heterogeneous the percentage of users with such a delay rises to 8.0%, for a heterogeneous γ_z even to 15.3%. Although the average waiting time is also slightly higher when heterogeneity among users is present (Figure 4.12), the extra delay compared to a case where there is no heterogeneity consists mainly of in-vehicle time. This can also be noted from the increase in the vehicle occupancy (Figure 4.14) in both alternative scenarios. When there is heterogeneity in β_z, 22.4% of the passenger time is spent with a maximum of two co-riders in the vehicle, 1.2% more than when there is heterogeneity in γ_z and 4.7% more than when user preferences are homogeneous (Figure 4.15).

Next to a higher average passenger in-vehicle delay, an increase in the average number of intermediate stops per passenger (Figure 4.14) is an indicator of less efficient routing for at least part of the travellers when heterogeneity in user preferences exist. The logical explanation for the fact that average delays are higher is that, as mentioned before, the majority of accepted requests in heterogeneous scenarios has an above average tolerance when it comes to delays (54.6%, *U_H_3_50*) or sharing with co-riders (70.1%, *U_30_H_50*). The assignment procedure in the model aims to maximize the number of assigned requests before considering the total net benefit for all users. In this case, the assignment is constrained by a relatively small fraction of requests with a below average delay or sharing tolerance. The assignment procedure will make sure that those requests are assigned within the constraints, which leads to less efficient routes for the other, more tolerant users (but still one that gives users a net benefit). Figures 4.16 and 4.17 show that users with 'positive' sharing preferences indeed tend to have larger delays. As this group forms the majority of requests, the average delay per passenger will increase.

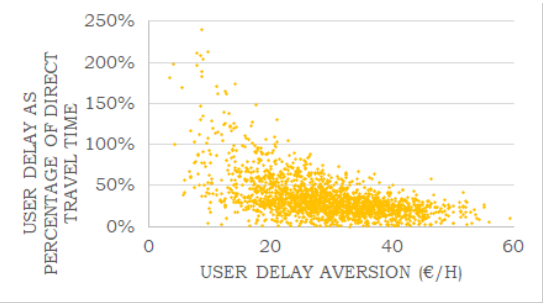


Figure 4.16: Delay per passenger as function of its delay aversion β_z

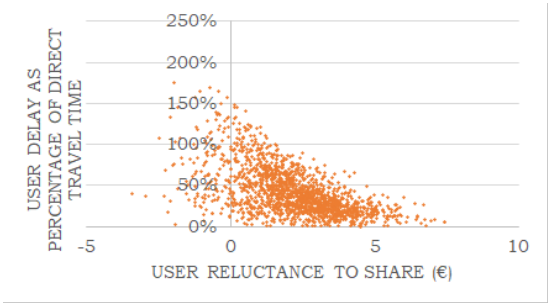


Figure 4.17: Delay per passenger as function of its reluctance to share γ_z

Vehicle distance

Figure 4.18 shows that the distance savings are hardly dependent on whether heterogeneity in user preferences exists. Without heterogeneity, each ride-sharing vehicle kilometer effectively transports 1.15 (direct) passenger kilometers, while with heterogeneity in delay aversion β_z this is 1.16 and with heterogeneity in reluctance to share γ_z 1.17. In absolute terms, ride-sharing generates a distance saving of 983, 989 and 939 kilometers respectively. If accounting for the fact that ride-hailing also requires deadheading, the (corrected) effective vehicle transportation distance ratio is still fairly similar, 1.34, 1.33 and 1.32 respectively.

A notable difference is found in the deadheading distance, that is significantly smaller when heterogeneity in users' reluctance to share exists. Figure 4.19 shows that both the total distance to access requests (related to a lower acceptance rate, Figure 4.9) and the rebalancing distance is lower when there is heterogeneity in the reluctance to share. The latter can be caused by the fact that it is more likely that an empty vehicle is already at the location of an unassigned request as there are more unassigned requests in the network. As many of those unassigned requests have an above average reluctance to share, the vehicles will not be assigned to serve the requests, as even without an in-vehicle delay the request will experience a net disbenefit as a result of its high reluctance to share and any previous waiting time. This phenomenon has been explained earlier in Section 4.1.

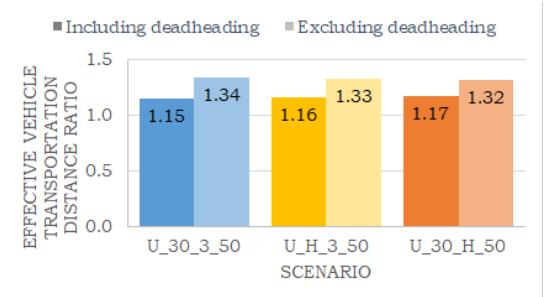


Figure 4.18: Effect of heterogeneity on relative distance savings, including and excluding deadheading

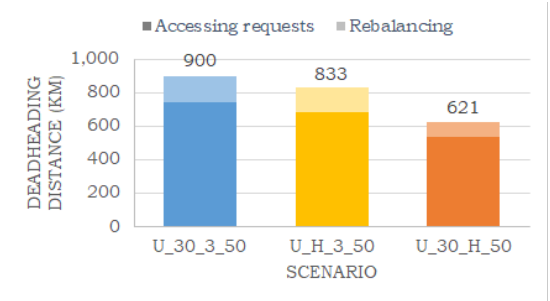


Figure 4.19: Effect of heterogeneity on the total deadheading distance

To summarize, in this section it was found that when heterogeneity exists in delay aversion β_z (i.e. a normal distribution with $\mu = €30$/h and $\sigma = €10$/h) or reluctance to share γ_z (normal distribution with $\mu = €3$ and $\sigma = €2$):

- The acceptance rate drops from 76.0% to 75.9% with heterogeneity in delay aversion β_z and to 66.8% with heterogeneity in reluctance to share γ_z.
- The average vehicle occupancy rises from 1.38 to 1.45 both with heterogeneity in β_z and in γ_z.
- The average user delay rises from 29.5% to 36.4% with heterogeneity in β_z and to 42.1% with heterogeneity in γ_z.
- The effective vehicle transportation distance ratio rises slightly from 1.15 to 1.16 with heterogeneity in β_z and 1.17 with heterogeneity in γ_z.

4.3 Occupancy-dependent reluctance to share

In previous sections it was assumed that the user reluctance to share represents the risk that a vehicle has to be shared with one or more passengers, and is therefore not dependent on the actual execution of the ride. In this section is analyzed how the efficiency and level of service of a ride-sharing system changes if this assumption is dropped and the discomfort of using a ride-sharing service is found to be lower if the ride is not actually shared (*RQ 3*). This section consists of a comparison of the KPIs of two scenarios: *U_30_3_50* and *U_30_D_50*.

Acceptance rate

It strikes that the acceptance rate of a ride-sharing service is much higher when ride-sharing discomfort is found to depend on the (maximum) number of co-riders during a user's ride: 94.1% versus 76.0% (Figure 4.20). Of the accepted requests, 74.3% is served fully privately like in an individual ride-hailing system without sharing, as presented in Figure 4.21. 24.2% of the accepted requests get assigned a ride that is shared at least at some point but never with more than one co-rider. Only for the remaining 1.5% of the requests at least a part of the ride is executed with maximum occupancy. In scenario *U_30_3_50* only 40.8% of the requests is served fully privately, 31.3% of the requests is shared but not at full capacity and for 27.9% of the requests the vehicle is shared with two co-riders at least a part of the ride. On average, in scenario *U_30_D_50* a vehicle in revenue mode is carrying 1.14 as opposed to 1.38 passengers, leading to a respective average of 0.31 and 0.95 intermediate stops (Figure 4.22).

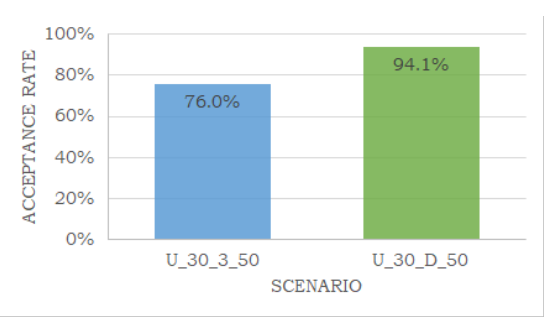


Figure 4.20: Acceptance rate in scenarios *U_30_3_50* and *U_30_D_50*

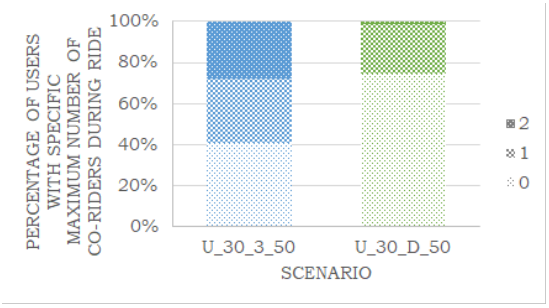


Figure 4.21: User distribution of maximum number of co-riders in ride in scenarios *U_30_3_50* and *U_30_D_50*

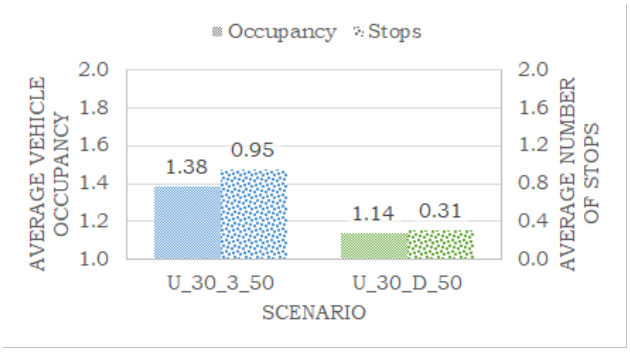


Figure 4.22: Average vehicle occupancy and number of intermediate stops in scenarios *U_30_3_50* and *U_30_D_50*

User delay

A consequence of the limited time that passengers have to share their vehicle is a low average in-vehicle delay (Figure 4.23), as detouring only occurs infrequently. To be specific, the average in-vehicle delay drops from 8.1% of the direct travel time (35 seconds) if is assumed that the willingness to share does not depend on the number of co-riders, to 3.0% (11 seconds) otherwise. Yet, the total average delay is higher with an occupancy-dependent reluctance to share, as shown by Figure 4.23. This figure illustrates also that this increase is caused by an increase in the average waiting time that outweighs the decrease in in-vehicle delay. The increase in waiting time follows from an increase in the number of accepted requests (Figure 4.20) in combination with a lower average vehicle occupancy while transporting passengers (Figure 4.22), which means that more vehicles are transporting passengers and fewer empty vehicles are available. The percentage of time vehicles spent halting in this scenario is 20.2%, almost twice as low as in the base scenario in which the proportion is 39.9% (Figure 4.24). Also, rebalancing in such a situation is less common (Figure 4.25), meaning that the probability of a quick pick-up is lower.

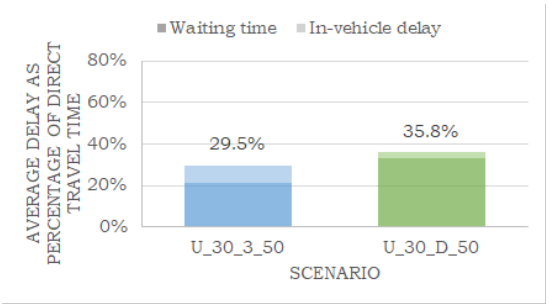


Figure 4.23: Average total delay per passenger in scenarios *U_30_3_50* and *U_30_D_50*

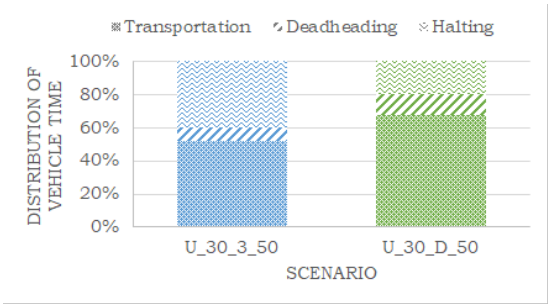


Figure 4.24: Proportion of vehicle time in different states in scenarios *U_30_3_50* and *U_30_D_50*

Vehicle distance

Next to a deterioration of the level of service due to long waiting times, the potential distance benefits that a ride-sharing service can offer also diminishes when reluctance to share increases with the number of co-riders. A reason for this is that the average vehicle occupancy is found to be low, which means that rides are hardly shared. To be precise, the average number of effective passenger kilometers that are satisfied with a ride-sharing

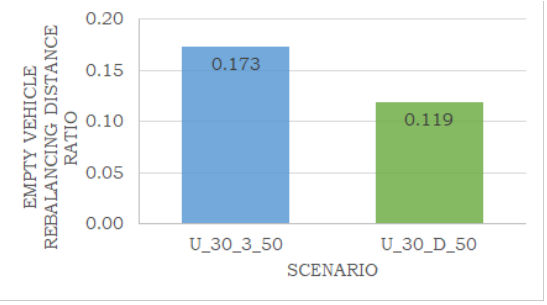


Figure 4.25: Rebalancing distance as ratio of the vehicle transportation distance in scenarios *U_30_3_50* and *U_30_D_50*

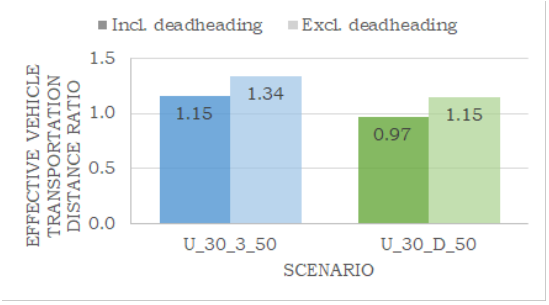


Figure 4.26: Effective vehicle transportation distance ratio in scenarios *U_30_3_50* and *U_30_D_50*

vehicle kilometer drops to 0.97 (Figure 4.26). The fact that this ratio is below 1 means that the total vehicle movement in the service is higher than the sum of the shortest paths between the origins and destinations of each of the accepted requests. Of course, in reality deadheading will also be needed to serve requests in a ride-hailing system and thus it is still likely that a ride-sharing service with the given assumptions induces a (minimal) reduction in total vehicle distance. Figure 4.26 shows that if deadheading is excluded from the vehicle distance, the effective vehicle transportation ratio is 1.15. Clear in any occasion is that the potential distance savings when sharing discomfort increases with the number of co-riders are much lower than when this is not the case.

To summarize, in this section it was found that if a user's reluctance to share is assumed lower when served with a private ride ($\gamma_z = €2$) and higher when (at least partially) served in a full vehicle ($\gamma_z = €4$):

- The acceptance rate increases significantly from 76.0% to 94.1%.
- The average vehicle occupancy drops from 1.38 to 1.14.
- The average user delay increases from 29.5% to 35.8% of the direct travel time.
- The effective vehicle transportation distance ratio drops from 1.15 to 0.97.

4.4 Occupancy-dependent discount

As opposed to the previous section that assumed that the reluctance to share increases with the (maximum) number of co-riders in the vehicle, this section elucidates the change in ride-sharing potential if users are assigned a larger discount if they share their ride with more co-riders (*RQ 4*). In other words, in this section ride-sharing with multiple co-riders is more rather than less attractive compared to the base scenario. To find out what the effects are hereof on the level of service and operational efficiency of a ride-sharing service, the results of scenario *U_30_3_D* are compared to the base scenario *U_30_3_50*. Specifically, the alternative scenario presents users with an additional 7.5% discount for every co-rider present in their vehicle in the busiest segment of the ride. This discount is added on top of the standard discount of 50% when a user chooses a ride-sharing service for their ride, meaning that the discount of an accepted request can either be 50% (when no co-riders join the ride), 57.5% (with a maximum of 1 co-rider) or 65% (when at least part of a user's ride is executed at full capacity).

Acceptance rate

One of the findings is that more requests are accepted by the system if sharing is stimulated by means of an additional discount that depends on vehicle occupancy: 82.1% vs 76.0% (Figure 4.27). Interesting is that 87.3% of the accepted rides are at least partially shared, compared to 59.2% when the discount is fixed (Figure 4.28). 74.1% of all accepted requests are even shared with two co-riders at some point. In the base scenario this is only the case for slightly more than a quarter of the accepted requests (27.9%).

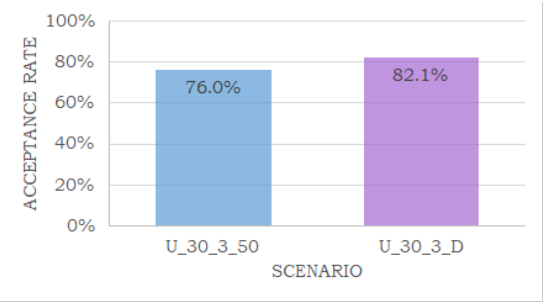


Figure 4.27: Acceptance rate in scenarios *U_30_3_50* and *U_30_3_D*

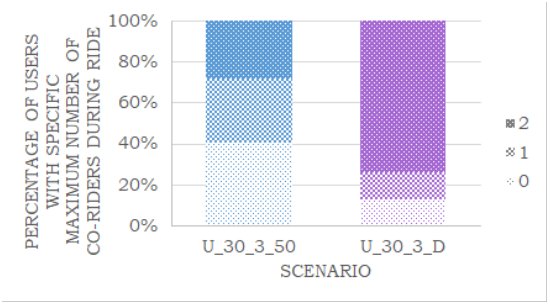


Figure 4.28: Distribution of maximum number of co-riders during ride in scenarios *U_30_3_50* and *U_30_3_D*

User delay

The larger tendency to share that was found when an additional discount is received, is also expressed in the average vehicle occupancy (1.85 vs 1.38, Figure 4.29), the average number of intermediate stops (with 1.97 more than doubled, Figure 4.29) and the passenger time spent with a specific number of (co-)riders on-board (Figure 4.30). It appears that with the extra discount, 45.1% of the passenger time is spent with two co- riders, which is 27.4% more than when no extra discount is given.

As sharing involves detouring, the average total delay per passenger rises when the discount received by users of a sharing service is assumed to be occupancy-dependent. Remarkably, an additional discount of maximum 15% of the fare leads to an increase in the average delay (Figure 4.31) from 29.5% of the direct travel time (127 seconds) in *U_30_3_50* to 54.0% (222 seconds) in *U_30_3_D*. The increase can be attributed mainly to the in-vehicle delay, which rises from 8.1% of the direct travel time to 25.1% as a result of increased detouring following from an increase in the vehicle occupancy. Apparently,

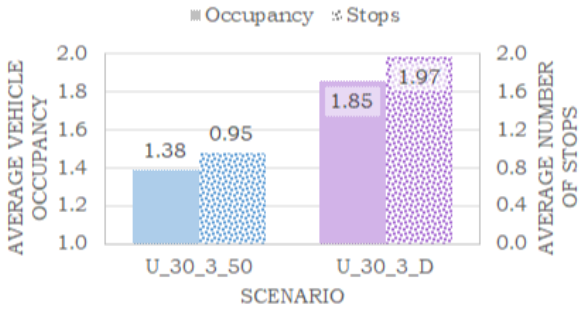


Figure 4.29: Average occupancy of vehicles in revenue mode and average number of intermediate stops in scenarios *U_30_3_50* and *U_30_3_D*

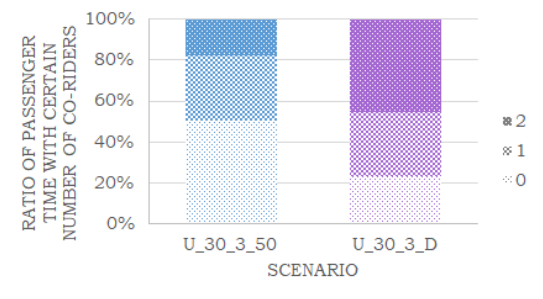


Figure 4.30: Distribution of passenger time with specific number of co-riders in scenarios *U_30_3_50* and *U_30_3_D*

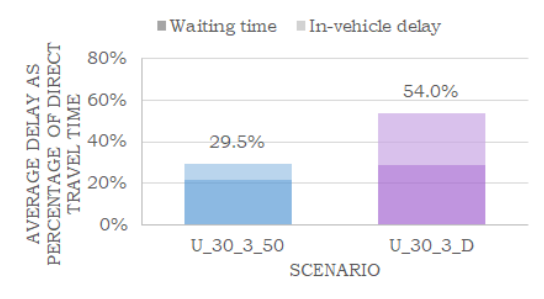


Figure 4.31: Average total delay per passenger in scenarios *U_30_3_50* and *U_30_3_D*

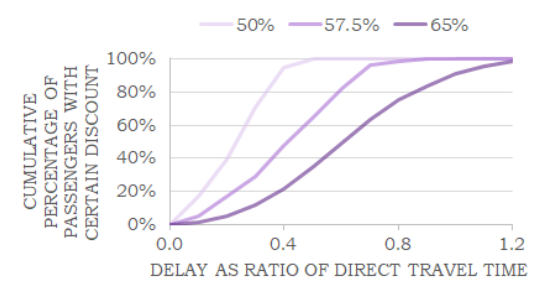


Figure 4.32: Cumulative percentage of passengers with certain delay, based on their discount (*U_30_3_D*)

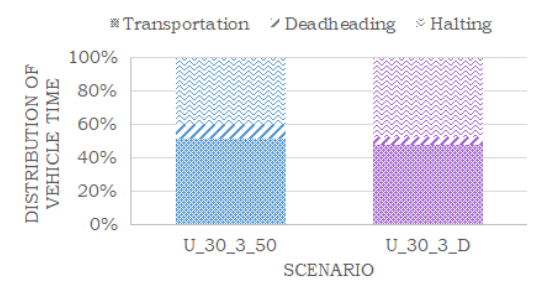


Figure 4.33: Percentage of vehicle time in different states in scenarios *U_30_3_50* and *U_30_3_D*

for those requests the benefit when sharing a vehicle with an additional passenger weighs up against the extra delay that follows from additional detouring to pick-up and drop-off co-riders. Figure 4.32 confirms that users on average indeed experience a larger (relative) delay when they receive a higher discount. Next to a higher in-vehicle time, Figure 4.31 illustrates that there is also a (moderate) increase in the average waiting time. On the one hand this is remarkable as vehicles are found to spend a larger proportion of their time halting, as shown by Figure 4.33. However, a way to explain it is that vehicles in transportation mode are being used very efficiently, meaning that they have a high occupancy. In fact, 27,8% of the time that vehicles are transporting people they do not have any more spare capacity and an additional 29.4% of the time there is only one empty seat. Hence, a new request is less likely to be picked-up by such a vehicle. The result is an increase in the average waiting time for a request.

Vehicle distance

When a higher discount is issued to rides with more co-riders, the number of vehicle kilometers decreases to 5,687 kilometers (Figure 4.34) as a result of more efficient transportation of requests with the available fleet of vehicles. Every kilometer driven by a vehicle in the ride-sharing fleet now effectively satisfies 1.34 passenger kilometers, compared to 1.15 with a flat 50% fare discount (Figure 4.35). In total 1,970 vehicle kilometers can be saved, more than double the amount than without an occupancy-dependent additional discount. Considering that ride-hailing also considers deadheading, total distance savings are potentially larger. If deadheading is excluded from the total vehicle distance, the effective transporation distance ratio increases even to 1.48 (Figure 4.35).

Interesting is that the reduced mileage compared to the base scenario is caused almost

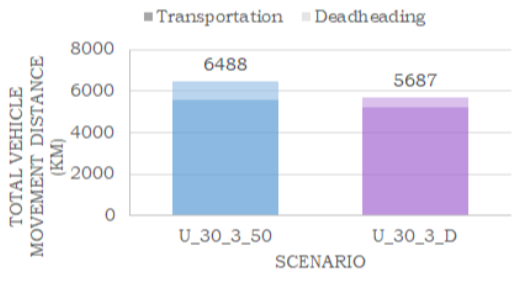


Figure 4.34: Total vehicle movement distance in scenarios *U_30_3_50* and *U_30_3_D*

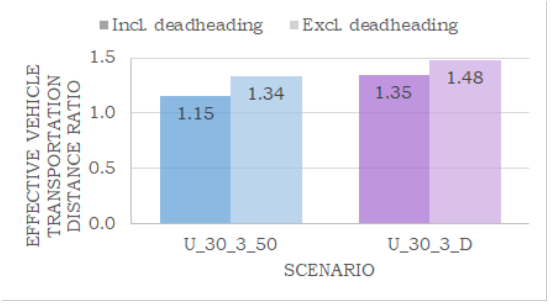


Figure 4.35: Effective vehicle transportation ratio in scenarios *U_30_3_50* and *U_30_3_D*, including and excluding deadheading

equally by a lower transportation distance (402 km) and lower deadheading distance (399 km), as shown also by Figure 4.34. With passengers receiving a compensation for additional detouring and more in-vehicle discomfort, vehicles that are in transportation mode have more flexibility to pick-up new requests along their way. Therefore, fewer requests are picked-up by empty vehicles, minimizing the total access distance. The total rebalancing distance is also lower (125 vs 156 kilometer) when there is an additional 7.5% discount per co-rider. This has to do with the fact that there are fewer unassigned requests in the system that vehicles can rebalance to. Therefore, many vehicles remain idle if they are not assigned to serve one or multiple requests.

To summarize, in this section it was found that if an additional 7.5% discount is awarded to users for each co-rider that a vehicle is shared with during the busiest part of the ride:

- The acceptance rate increases from 76.0% to 82.1%.
- The average vehicle occupancy increases from 1.38 to 1.85.
- The average user delay increases from 29.5% to 54.0% of the direct travel time.
- The effective vehicle transportation distance ratio increases from 1.15 to 1.35.

4.5 Demand distribution

This section contains the analysis of the effect of the spatial distribution of demand on ride-sharing level of service and efficiency, which will help answer sub research question 5. The analysis consists of a comparison of the result of three scenarios: one where demand is uniformly distributed (*U_30_3_50*), one where demand is moderately directed towards the center of the grid (*MD_30_3_50*) and one where demand is strongly directed towards the center (*SD_30_3_50*).

Acceptance rate

A first clear result is that the acceptance rate is significantly lower when the demand is more directed (Figure 4.36). When demand is perfectly uniform, 76.0% of all requests are accepted. When there is a moderate level of concentration, the acceptance rate drops to

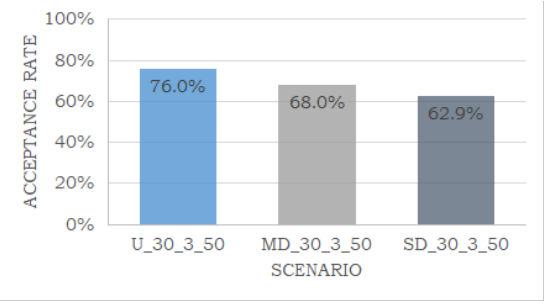


Figure 4.36: Effect of demand distribution on acceptance rate

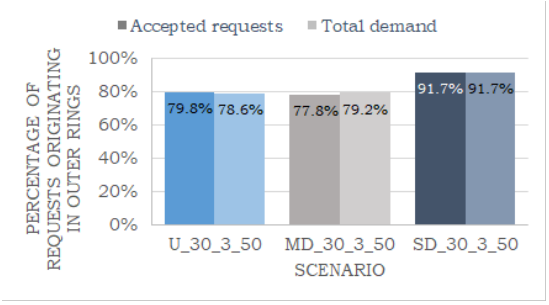


Figure 4.37: The share of requests originating in outer rings for accepted versus total demand, given different distributions of demand

68.0% and when there is strong concentration even to 62.9%. When comparing the origins of accepted requests to the origins of all requests, no clear difference is found (Figure 4.37). In other words, there is an approximately similar drop in the acceptance rate for both requests originating in the center of the grid and in the periphery.

User delay

The average delay per passenger is almost similar for a case with uniform and with moderately directed demand: 29.5% of the direct travel time versus 29.1% (Figure 4.38). When comparing both cases, it is found that with moderate direction in demand, the waiting time is higher and in-vehicle delay lower. This can be explained by a lower average vehicle occupancy, on average 1.32 compared to 1.38, as shown in Figure 4.39. An explanation for a lower vehicle occupancy is that the average direct trip length of requests is lower than in the scenario where demand is uniformly distributed (Figure 4.39). Passengers therefore require fewer time in the vehicle, meaning that, given that the fleet size is fixed, it is less likely that rides are shared, as indeed shown by Figure 4.40.

When demand is more than moderately directed (*SD_30_3_50*), the average user delay rises to 31.9% of the direct travel time between a request's origin and destination. Figure 4.38 illustrates that the extra travel time in this scenario is mostly caused by an additional waiting time. 88.1% of the requests with a pick-up location in one of the three outer rings and a drop-off location in/around the center of the grid ("O → I" in Figure 4.41), do not have an in-vehicle delay at all. However, due to a high waiting time, the total average delay of those requests in such a scenario is still high (Figure 4.42). More waiting time

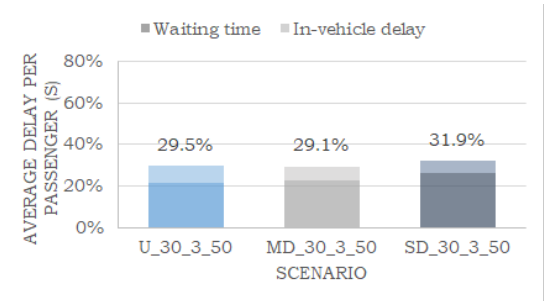


Figure 4.38: Effect of demand distribution on average total delay per passenger

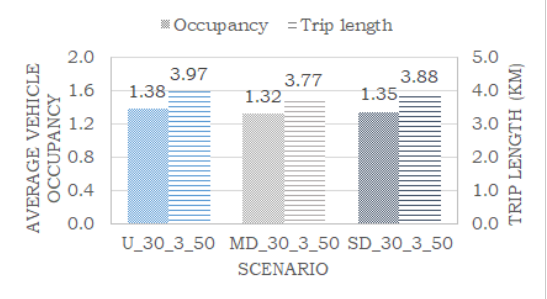


Figure 4.39: Effect of demand distribution on the average vehicle occupancy in revenue mode and average shortest path trip length between a request's origin and destination

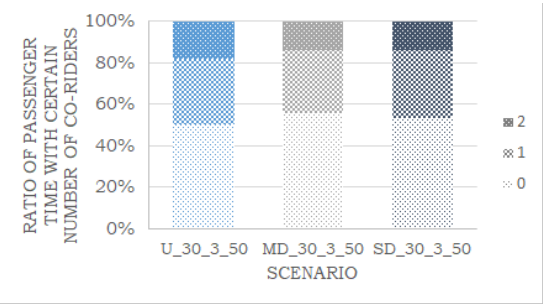


Figure 4.40: Effect of demand distribution on distribution of passenger time with specific number of co-riders

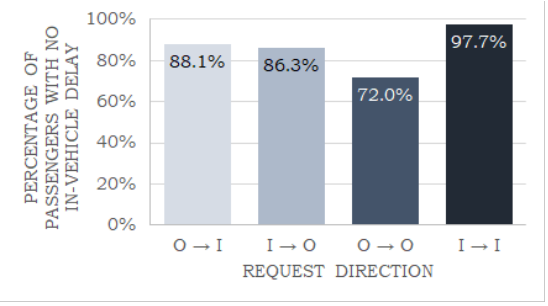


Figure 4.41: Percentage of requests in *SD_30_3_50* with no in-vehicle delay given the trip direction, "O" meaning location is in one of the three outer rings, "I" in any of the other rings

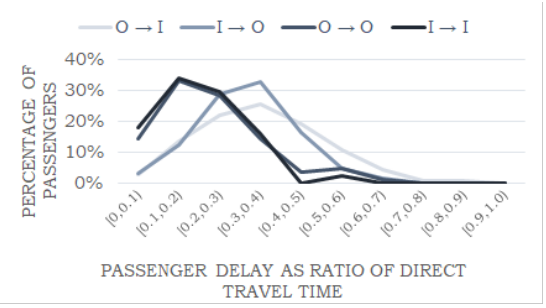


Figure 4.42: Total passenger delay in *SD_30_3_50* given trip direction, "O" meaning location is in one of the three outer rings, "I" in any of the other rings

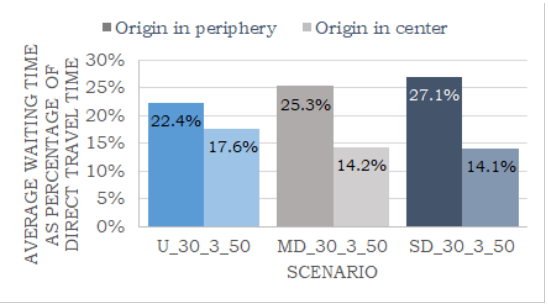


Figure 4.43: Relative user waiting time given request location in scenarios *U_30_3_50*, *MD_30_3_50* and *SD_30_3_50*

for a more directed demand distribution is actually expected, with such a distribution basically consisting of areas with overproduction (more pick-ups than drop-offs) and underproduction (more drop-offs than pick-ups). Users in the first type of area have to wait for empty vehicles to arrive from areas with underproduction in order to be picked-up. Therefore, the waiting time at locations in the periphery, with overproduction, is much higher (nearly twice as high as illustrated by Figure 4.43) than in the center, an area with underproduction.

Vehicle distance

The effective vehicle transportation ratio, an indicator for operational efficiency or potential distance savings, is found to drop when the magnitude of direction in demand increases (Figure 4.44). With a perfectly uniform demand distribution, every vehicle kilometer of a ride-sharing service effectively serves 1.15 passenger kilometers, with a moderate level of direction 1.06 passenger kilometers and with strongly directed demand only 0.95 kilometers, meaning that the latter will actually induce an increase rather than a reduction in vehicle kilometers on the road. As stated before, this ratio assumes that no deadheading is needed to serve requests individually, while the total ride-sharing distance does include deadheading. When deadheading is excluded for ride-sharing as well, the ratios are much more similar (respectively 1.34, 1.31 and 1.31, Figure 4.44). What can be deduced is that the lower efficiency when there is direction in demand is caused mainly by an increase in the deadheading distance. In fact, 27.6% of the vehicle distance is generated without any

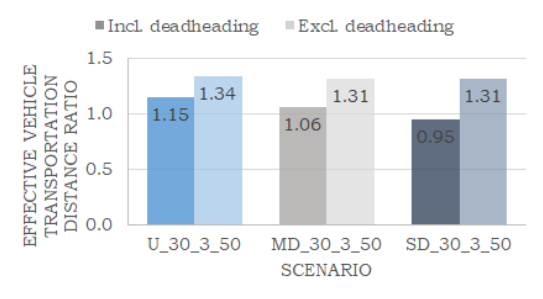


Figure 4.44: Effect of demand distribution on effective vehicle transportation distance ratio

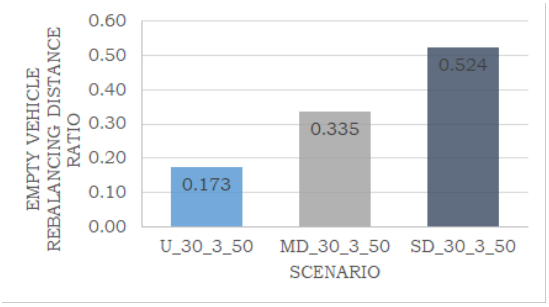


Figure 4.45: Effect of the demand distribution on rebalancing distance ratio

passengers on-board when demand is strongly directed, as opposed to only 13.9% with a perfectly uniform demand. These vehicle kilometers are principally undesired as they do not directly contribute to satisfying passenger needs for transportation.

As rebalancing can help solve spatial imbalances in demand, which exist when demand is at least to some extent directed, it is not surprising that a large share of the deadheading distance results from rebalancing vehicles (Figure 4.45). In fact, in a scenario with strongly directed demand, more than 50% of the total deadheading distance is produced by rebalancing vehicles.

To summarize, in this section it was found that if the level of direction in demand increases:

- The acceptance rate decreases, from 76.0% when demand is uniform to 62.9% when demand is strongly directed.
- The average vehicle occupancy initially decreases and later increases, with 1.38 when demand is perfectly uniform, 1.32 with moderately directed demand and 1.35 with strongly directed demand.
- The average user delay rises, with on average a delay of 31.9% of the direct travel as compared to 29.5% and 29.1% when demand is perfectly uniform or moderately directed.
- The effective vehicle transportation distance ratio decreases, from 1.15 when demand is perfectly uniform to 0.95 when demand is strongly directed.

4.6 Correspondence of objectives

Different actors have different interests in a ride-sharing system. An average potential user has a preference for a minimal delay, prefers a ride with a limited number of co-riders and wants to receive a sufficient financial compensation for the inconveniences of sharing a ride. Authorities and operators are more interested in the operational efficiency that ride-sharing can deliver. The question is if these objectives can be unitedly achieved or whether a trade-off is always required. This section considers the relationship between three KPIs: the acceptance rate, average delay as ratio of the direct travel time (relevant for users) and the effective vehicle transportation distance ratio, i.e. the total effective passenger kilometers as a ratio of the total vehicle distance (relevant for operators and

authorities).

In Figure 4.46, the average user delay as ratio of the direct travel time is plotted against the acceptance rate of a ride-sharing service. Generally, it can be said that there is a positive correlation between the two outputs, meaning that an increase in the acceptance rate generally comes along with an increase in the average user delay. This means a system controller will need to trade-off the number of accepted requests versus the average delay of those requests. However, there also exist Pareto-inefficient scenarios that perform worse on both the acceptance rate (i.e. it is lower) and average user delay (higher) compared to at least one other scenario. Such scenarios include both scenarios that capture heterogeneity and those in which demand is non-uniform. From a user perspective, these scenarios are thus not desired. Note that the Pareto frontier that is included in Figure 4.46 would be more relevant when comparing alternative designs, or in other words, when the service provider can influence the decision variables that determine the outcomes of the ride-sharing service. In this study, most of the scenarios consider a variation in external variables.

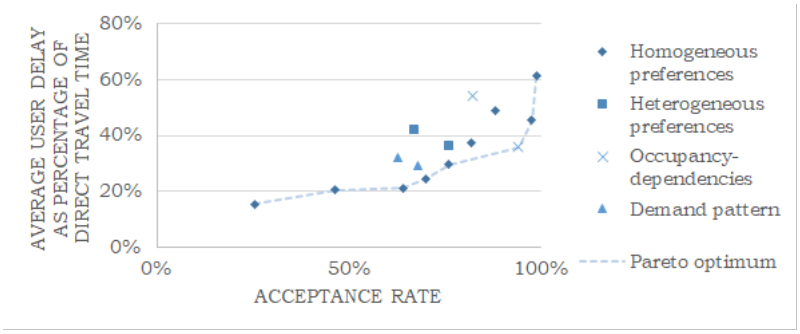


Figure 4.46: Relation between the acceptance rate and average user delay for all scenarios, including the Pareto frontier

Interestingly, the acceptance rate and effective vehicle transportation distance ratio seem to be positively correlated (Figure 4.47), meaning that a higher acceptance rate is likely related to larger distance savings. Therefore, both objectives do not have to be traded-off and aiming for the highest acceptance rate will most likely also yield the highest operational efficiency (or vice versa). For both objectives, the optimum scenario is the one in which the reluctance to share is only €1.

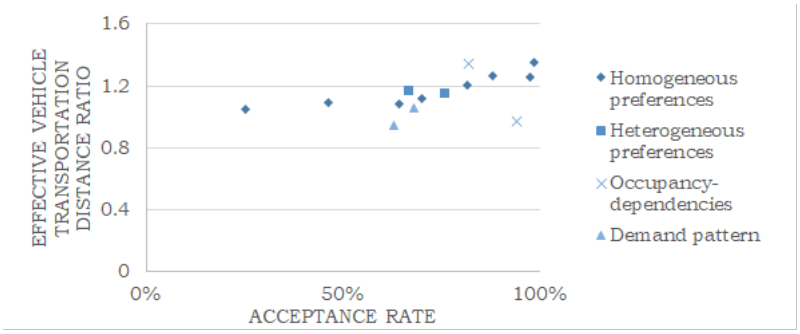


Figure 4.47: Relation between the acceptance rate and total distance savings for all scenarios

There exists no scenario in which both the average delay for users and the operational efficiency are optimal. In other words, a Pareto frontier is found consisting of all scenarios that are Pareto optimum, as illustrated in Figure 4.48. This frontier is made up of most of the scenarios in which user tolerances are homogeneous, demand is uniform and neither

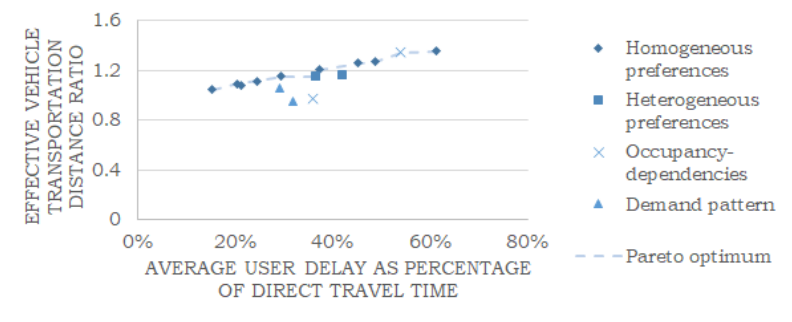


Figure 4.48: Relation between the average user delay and total distance savings for all scenarios, including the Pareto frontier

the fare discount nor the reluctance to share is occupancy-dependent. One of the two scenarios in which the delay aversion is heterogeneous is also Pareto optimum, as well as the scenario in which the fare discount increases with vehicle occupancy. The Pareto front shows that there is competition between the level of service for a user and the operational efficiency for an operator / authority, as well as that scenarios with directed demand or an occupancy-dependent reluctance to share are ideal for neither.

To summarize, in this section it was found that:

- The acceptance rate, effective vehicle transportation distance ratio and average user delay as ratio of the direct travel time are positively correlated.

4.7 Computational complexity

Computational efficiency is one of the main reasons to use the RGV-approach to assign requests to vehicles. However, the computational effort can still be significant if the number of variables in the system is large. In this study, to minimize the computational effort, it was decided to limit the vehicle capacity to 3, while at the same time applying the model on a small-scale network with relatively few requests.

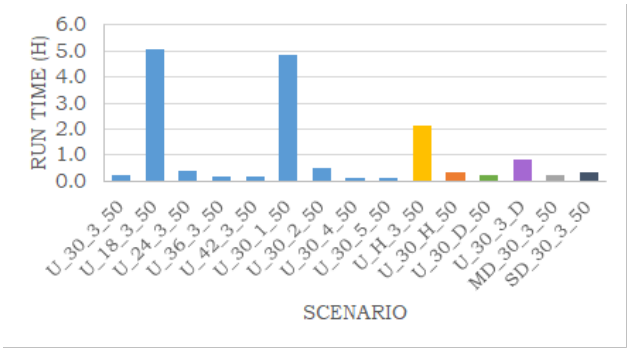


Figure 4.49: Model run time of all scenarios with a single-core 2.30GHz processor

When comparing the computational effort of different scenarios, expressed as the model run time with a computer with a single-core 2.30GHz processor (Figure 4.49), it stands out that two scenarios are much more computationally heavy than the others: *U_18_30_50* and *U_30_1_50*. The two scenarios concerned have respectively the lowest delay aversion and reluctance to share of all scenarios. It means that requests can be fulfilled with

large delays and potentially large request groups can be assigned to a single vehicle. It requires significant computational time to test request groups consisting of a large number of requests as the set of sequences of origins and destinations of those request group - vehicle combinations is significantly (i.e. more than exponentially) larger than for small request groups.

Another scenario that requires a relatively large computational effort is the one with heterogeneity in delay aversion β_z (*U_H_30_3*). The reason for this is that a large fraction of the accepted users in this scenario was found to have a below average delay aversion, meaning that, again, large request groups can be composed that have to be checked for feasibility against the user (net benefit) and vehicle (capacity) constraints.

Finally, and again with a similar reason, the scenario in which users receive a discount for the maximum number of co-riders with whom they share the vehicle (*U_30_3_D*), is computationally complex. Large request groups are more likely to be feasible as the benefit of each of the requests in such a group is higher and therefore more detouring is accepted.

To summarize, in this section it was found that:

- The computational complexity is (by far) largest for scenarios in which user tolerances to delays and co-riders are very high and homogeneous.

Chapter 5

Conclusions & recommendations

This chapter wraps up the research by summarizing the main findings (Section 5.1) and presenting the conclusions that can be drawn from those findings (Section 5.2). Section 5.3 includes recommendations for authorities and service providers that are interested in the introduction of ride-sharing services. Finally, Section 5.4 reflects on the assumptions and results of this study and concludes with the according recommendations for future research in this field.

It should be kept in mind that the results of this study are subject to numerous assumptions. The conclusions drawn in this research may thus not be valid for alternative fleet designs, network properties, demand magnitudes and driver behavior.

5.1 Key findings

In this section the main findings are discussed based on the six sub research questions that were formulated for this research. In each part, the sub question is first mentioned and then answered based on the results that were presented in Chapter 4. The different parts are then linked together to answer the main overarching research question.

1. *How do efficiency and level of service of a dynamic ride-sharing service depend on the delay tolerance, waiting time tolerance and willingness to share, assuming uniformity across users?*

For the dynamic ride-sharing service proposed in the numerical experiment, the acceptance rate is found to increase approximately linearly with the delay and sharing tolerance of potential users when is assumed that those preferences are homogeneous. If the willingness to share is for example very high, meaning that users demand only a 1 or 2 euro compensation for sharing a vehicle, more than 98% of all requests are accepted, as opposed to 25.4% when a compensation of €5 is required. The linear increase in accepted requests leads to a more than linear increase in the average vehicle occupancy. One of the consequences of this is that both waiting time and extra in-vehicle time increase more than linearly as well. Not only is there more detouring when users are more open to delays and sharing a vehicle, also pick-ups take longer as vehicles on average have less spare capacity to pick-up new requests. On the other hand, if users are very sensitive to delays or the fact that the vehicle has to be shared, sharing is found to be very limited with the average vehicle occupancy dropping from 1.38 for average sharing tolerances to 1.24 for the lowest delay tolerance and 1.14 for the lowest willingness to share. In either case, in-vehicle delays are very limited.

In such a situation, the number of effective passenger kilometers per vehicle kilometer can be as low as 1.05, given a reluctance to share of €5. Based on the numerical experiment, this ratio is found to increase approximately linearly with an increase in the

delay and sharing tolerance. If users are very tolerant to sharing a vehicle with co-riders (requiring only a €1 compensation for a shared ride without delays), the total effective passenger kilometers to be served with one vehicle kilometer rises to 1.36.

If users are more tolerant to delays and sharing with co-riders, there is more flexibility in the assignment of requests to vehicles, which results in a more efficient operation of the respective requests in terms of the vehicle distance needed to serve those requests. Deadheading is also more uncommon when users' sharing and delay tolerances are high, as new requests can be picked-up by vehicles on their way to dropping off other passengers. However, the deadheading distance was also found to be low in case of negative sharing attitudes, given that relatively few requests are picked-up. Like deadheading, rebalancing in this study is found to lead to the largest vehicle distance when the attitude towards sharing is neutral. When users have a relatively positive sharing attitude, more requests are served, meaning that there are fewer empty vehicles that can rebalance. On the contrary, if the average willingness to share is low, it is more probable that a vehicle cannot serve a request of which the pick-up location is at or near the vehicle's location, even if the vehicle follows the direct route to serve that request. Therefore, in such a situation it is more likely that a vehicle in the rebalancing step is assigned to a request at its own location or one that is close-by. This explains why, based on the numerical experiment, the total rebalancing distance decreases when the sharing tolerance is below average. A final remark should be made that all of the above mentioned findings are more pronounced for the willingness to share than for the delay tolerance. Apparently, the level of service and operational efficiency are more sensitive to the tested values of the reluctance to share a vehicle with co-riders than to the tested values of users' aversion to delay.

2. *What is the effect of heterogeneity in the delay tolerance, waiting time tolerance and willingness to share on the efficiency and level of service of a dynamic ride-sharing service?*

In a scenario with heterogeneity in the willingness to share, the acceptance rate of the assumed ride-sharing service drops from 76.0% to 66.8%, with many users with a below average willingness to share not served. If there is heterogeneity in the delay tolerance instead, the acceptance rate (75.9%) hardly drops. An explanation for this difference is that users with a below average delay tolerance can still be satisfied by serving them with the shortest or a relatively short path to minimize their delay. This happens at the cost of other more flexible passengers that will get served with a delayed pick-up and/or a less direct route. With the majority of accepted requests having an above average delay tolerance (since it is easier to find a ride for those requests), the average delay experienced by users in the system, 36.4% of the direct travel time, is higher than when the delay tolerance is assumed homogeneous (29.5%). When considering a scenario with heterogeneity in the willingness to share, it is found that users with a below average willingness to share are much more likely to reject a ride-sharing service even if offered a direct ride, leading to a lower total acceptance rate. With many unassigned vehicles and many unaccepted requests, it is relatively common for an unassigned vehicle to be located at or near an unassigned request. The reason that such an unassigned vehicle cannot serve a nearby unassigned request is that unaccepted requests in this scenario tend to have a relatively high reluctance to share, meaning that even if the request is picked-up quickly and served without an in-vehicle delay, it will experience a net disbenefit. As there are thus relatively many unassigned requests around unassigned vehicles, the total rebalancing distance is limited. The effective transporation distance ratio ranges between 1.15 and 1.17 for the three scenarios in this part of the experiment, and thus does not seem to be significantly dependent on whether heterogeneity in ride-sharing tolerances exists. It can be said that, based on the numerical experiment, the operational efficiency in terms

of the distance that a ride-sharing system can save is not very dependent on the variation of sharing preferences over the population.

3. *How do efficiency and level of service of a dynamic ride-sharing service depend on how willingness to share relates to the number of co-riders in the vehicle?*

In previous research questions, the willingness to share actually represented the willingness to risk having to share, or in other words the willingness to use a ride-sharing service disregarding of the service quality that is offered. When the willingness to share is assumed to be dependent on the (maximum) number of (simultaneous) co-riders in the vehicle, it is found that many more requests can be accepted by the system (94.1%). However, as sharing a vehicle with co-riders now induces an additional disbenefit, only 25.7% of the accepted rides is actually partially shared. With 74.3% of all rides served fully privately, the distance savings are limited. In fact, the number of effective passenger kilometers per vehicle kilometer in this scenario is only 0.97. It means that the ride-sharing service needs more kilometers to serve the accepted requests than the sum of the shortest paths between the origins and destinations of those requests. It is still possible that in reality there is a distance reduction as the latter metric assumes that no deadheading is needed to serve requests. Anyhow, it is clear that operational efficiency is lower when willingness to share is assumed to be dependent on the vehicle occupancy.

It is remarkable that although sharing is limited in the system, meaning that usually the shortest path between origin and destination is followed, the total average delay of a user in the system is in fact higher than without an occupancy-dependent willingness to share: 35.8% versus 29.5% of the direct travel time. It was found that this is caused by a significant increase in the average waiting time of passengers, which follows from the fact that there are very few empty vehicles in the system to pick-up new requests (only 20% of the vehicle time is spent halting) and that non-empty vehicles are unlikely to pick-up new requests when passengers' willingness to share drops with the number of co-riders in the vehicle. Rebalancing is also limited as there are both few unassigned requests and unassigned vehicles. To summarize, a ride-sharing service in this context basically resembles an individual ride-hailing system with a poor level of service, with requests being served mostly with a direct route but also long waiting times resulting from a lack of available vehicles, with the fleet of ride-sharing vehicles in the numerical experiment being smaller than the fleet of traditional taxis in a ride-hailing system without sharing.

4. *How do efficiency and level of service of a dynamic ride-sharing service change when the fare discount for a shared ride depends on the number of co-riders?*

As expected, when an additional financial reward is given to users based on the number of co-riders they shared a ride with, the average vehicle occupancy (from 1.38 to 1.85) and the relative time passengers spent in a vehicle at full capacity (from 17,7% to 45.1%) increases. By utilizing the available vehicle capacity more efficiently, the number of accepted request increases from 76.0% to 82.1%, although at the cost of a higher average delay. A higher vehicle occupancy will burden a passenger with larger detours and consequently an in-vehicle delay that is more than three times as high as when no additional discount is offered (25.1% versus 8.1% of the direct travel time). Also, the waiting time of requests is marginally larger in the scenario with an occupancy-dependent discount with pick-ups being complicated by the fact that many vehicles are driving around fully occupied. Users in the system get compensated for a longer travel time by the additional discount on the ride fare that was offered to them.

It was found that such an (maximum 15%) additional discount can potentially pay off from an operational perspective, with the effective passenger kilometers per ride-sharing

kilometer in the experiment increasing from 1.15 to 1.35. In combination with a higher acceptance rate, a relatively large distance savings can be achieved with the additional discount. The distance that a ride-sharing service can save in this way originates not only from more efficient transportation of requests but also from a reduction of the deadheading distance to access new requests, as requests are being picked-up by non-empty vehicles on their way to dropping off other passengers.

5. *How do efficiency and level of service of a dynamic ride-sharing service depend on the spatial distribution of demand over the network?*

Based on the numerical experiment, it can be concluded that more direction in demand leads to more requests being rejected by a ride-sharing service (37.1% compared to 24.0% when demand is perfectly uniform). The probability of being rejected for a request is found to be independent of whether the requested pick-up location is located in an area with overproduction (in this case in the periphery) or underproduction (center of the grid). If demand is perfectly uniform, the average vehicle occupancy of vehicles in revenue mode is 1.38 and the average passenger delay 29.5% of the direct travel time. The drop in the number of accepted requests when there is a moderate level of direction in demand leads to a drop of the vehicle occupancy (1.32) and the average delay (29.1% of direct travel time) accordingly. It is intriguing that if the level of direction increases further in the experiment, the average delay starts to increase again (to 31.9% of the direct travel time in a scenario where demand is strongly directed). With a larger spatial inequality in pick-ups and drop-offs, in some locations there is an excess in vehicles and in other areas an excess in requests. In the latter area, many requests are competing for the few vehicles that are near, meaning that the average waiting time is high for those requests. In areas with a surplus of vehicles that have just dropped off their remaining passengers, average waiting times are low, however only the minority of requests is originating there. Therefore, the average waiting time is mainly determined by requests in areas with overproduction that have to wait for empty vehicles to deadhead from areas with underproduction. At the same time, deadheading is responsible for 27.6% of the vehicle kilometers in the scenario in which demand is strongly directed, while responsible for only 13.9% of the mileage when demand is uniform. Both deadheading to access requests and to rebalance are common to spatially match supply and demand in the former situation. The effective passenger kilometers per ride-sharing vehicle kilometer drops from 1.15 when demand is perfectly uniform to 0.95 when demand is strongly directed. It is thus fair to conclude that, based on the numerical experiment, the operational efficiency is found to be strongly dependent on the spatial distribution of demand.

6. *What is the relation between the acceptance rate, the average delay of shared rides and the total reduction in vehicle distance following from a ride-sharing service?*

When it comes to trading off different ride-sharing objectives, it is found that striving for a maximum acceptance rate (a relevant reliability indicator for a potential user) will generally also yield a maximum effective vehicle transportation distance ratio. In that case, more requests are potentially shared, meaning that fewer extra vehicle kilometers are needed for each extra request. In terms of the acceptance rate and potential distance savings, a scenario therefore exists that is optimal for both KPIs. Specifically, it concerns the scenario with the highest (homogeneous) willingness to share. User delays on the other hand have to be traded off against the acceptance rate and total distance savings, meaning that an improvement in the acceptance rate / effective vehicle transportation distance ratio will at the same time induce a larger average user delay, and vice versa. Thus, the design of a ride-sharing service in a specific context will comprise trading-off

user level of service with efficiency of resources. Yet, it can be noted that some of the scenarios are Pareto inefficient, meaning that they are desired from neither a user's nor an operator's perspective. This is generally the case for scenarios in which heterogeneity in sharing preferences exists and where demand is at least moderately directed.

After answering each of the sub research questions, an answer can be formulated to the main research question, that was defined as follows:

> *How do efficiency and level of service of dynamic ride-sharing services depend on behavioral and service attributes and the distribution of demand?*

Both the efficiency and level of service of a ride-sharing service are found to be strongly dependent on delay tolerance and (especially) willingness to share of users, if assumed that behavioural preferences are homogeneous as in some scenarios of the numerical experiment. The acceptance rate of a ride-sharing service can vary between 25.4% and 98.8%, the average delay between 15.2% and 61.3% of the direct travel time, and the number of effective passenger kilometers per vehicle kilometer between 0.95 and 1.36, given different sharing and delay attitudes. Generally, if a ride-sharing service under specific conditions has a high acceptance rate in the experiment, the vehicle distance that is saved is also high. On the other hand, in such a system the average delay for users is also likely to be large. Although heterogeneity in user preferences seems to have a limited effect on potential distance savings, it has negative consequences for the level of service with more rejected requests and a larger average delay. It is found that in such a scenario to serve a few users with a below average sharing tolerance, many of the accepted requests with an above average tolerance get assigned a less efficient ride. If assumed that the willingness to share is not only risk-based but also dependent on the experienced ride, 94.1% of the requests are served but mostly by a fully private ride. Yet, the level of service of accepted requests is mediocre in such a scenario, with a long average waiting time due to the fact that few empty vehicles are available and non-empty vehicles are not likely to be shared. In terms of efficiency, the ride-sharing service in the numerical experiment actually produces more rather than less vehicle kilometers in such a context. The efficiency is much better when an additional discount is offered to users depending on the maximum number of simultaneous riders they are sharing a vehicle with. More requests are accepted (82.1%) and the effective transportation distance ratio rises from 1.15 to 1.35 effective passenger kilometers per vehicle kilometer. In return for a discount, users are burdened with an extra travel time of on average 24.5% of the direct travel time. Finally, collective attributes such as the demand distribution can significantly influence ride-sharing efficiency and level of service. The larger the directionality of demand in the experiment, the fewer requests are accepted and the lower the effective vehicle transportation distance ratio, potentially even below 1 when the concentration of demand is very strong. In such a situation, the level of service of ride-sharing users was also found to be low with long waiting times before pick-up.

The main findings of this research are visually presented in Figure 5.1. This diagram shows the main relations between the individual, collective and service attributes that were tested and the different ride-sharing level of service and efficiency indicators. It is interesting to notice here that sometimes a variable can influence another in multiple, possibly opposite, ways. For example, stronger directionality in demand can principally lead to a longer or shorter average waiting time with (1) large competition for pick-ups in areas with overproduction (direct, positive effect on average waiting time), (2) a better spatial compatibility of requests, meaning that vehicle occupancy is higher and less spare capacity is available to pick-up new requests (indirect, positive effect), and (3) a lower total acceptance rate following from inefficient assignment, meaning that more empty vehicles are available to pick-up new requests (indirect, negative effect). The fact that variables

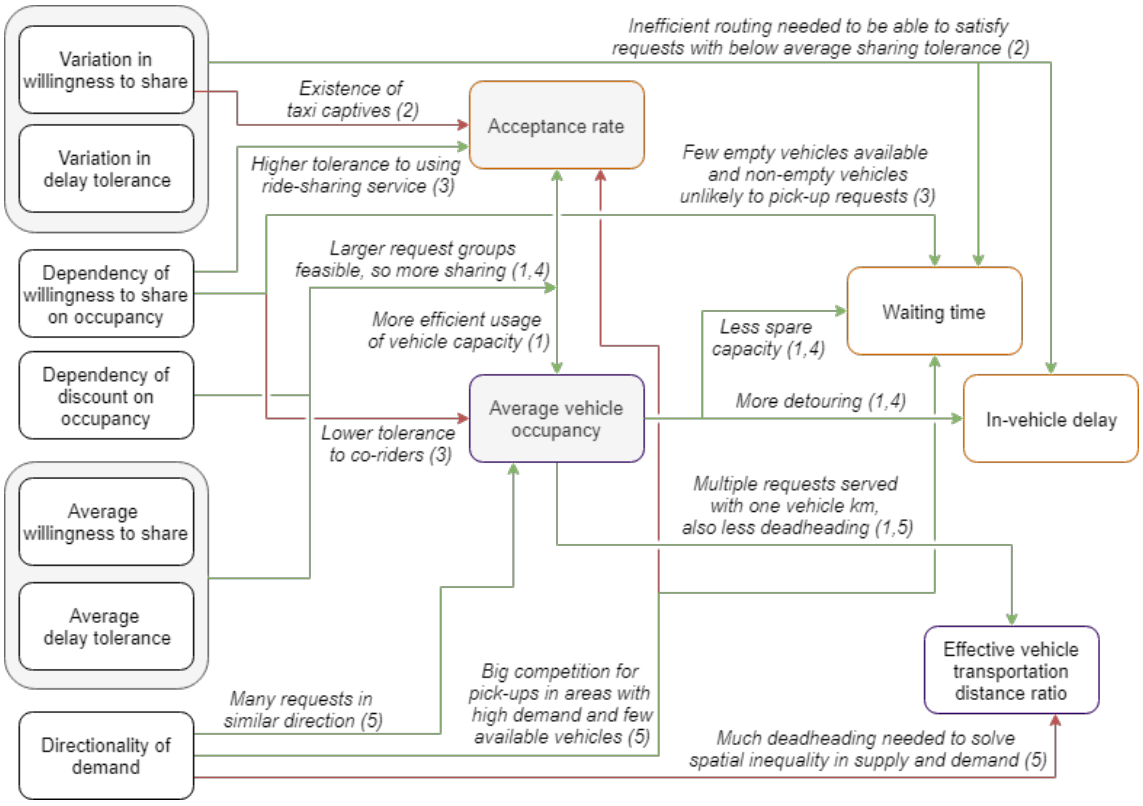


Figure 5.1: Most relevant relations found between individual, collective and service attributes and ride-sharing level of service and efficiency indicators (green = positive relation, red = negative relation, numbers correspond to the research question with which a relation was found)

can relate in different ways, illustrates the complexity of ride-sharing studies. It means also that the direction of the relation between ride-sharing attributes and performance can potentially be context-specific.

Nevertheless, based on the different scenarios that were tested in the experiment, it can be stated that a higher acceptance rate is positively correlated with the efficiency of operations (in terms of the effective vehicle transportation distance ratio), as also shown by the nature of the relations in Figure 5.1. In the majority of the scenarios however, an increase in the acceptance rate and effective vehicle transportation distance ratio comes along with an increase in the average user delay, meaning that a trade-off between the different performance indicators is required.

5.2 Implications

This work is the first study to consider ride-sharing potential based on the trade-off that users are faced with when presented with the option of ride-sharing. Previous studies, such as that of Santi et al. (2014) and Alonso-Mora et al. (2017), assumed that all users are potentially willing to ride-share as long as their waiting time and total delay does not exceed a certain threshold. This is not very realistic as taxi users have no reason to share their ride (and accept a delay) when they do not get a benefit in return. Therefore, a discount on the base fare, which is distance-dependent, is introduced in this study to compensate users for a potential delay. The assumption is that users will only switch to a ride-sharing service if such a choice gives them a net positive utility over a conventional taxi or ride-hailing ride. Also, as opposed to previous studies on ride-sharing potential,

this work accounts for the fact that negative utility can also follow from social discomfort when sharing a vehicle with strangers. It is relevant to include this aspect in studies on ride-sharing potential as it has been found in literature that the fear of negative social interaction can be one of the main barriers for a successful implementation of ride-sharing services. This research shows, in fact, that users' willingness (or actually reluctance) to share seems to have a larger impact on ride-sharing potential than users' attitudes towards delays. This is especially the case when sharing discomfort is found to be dependent on the number of co-riders with whom a vehicle is shared.

This section summarizes the implications of a user trade-off between a delay, sharing discomfort and fare discount for the expected total vehicle mileage on the road network (Subsection 5.2.1). It also offers some first insights into the profitability of ride-sharing services (Subsection 5.2.2) and the equity for its users (Subsection 5.2.3).

5.2.1 Vehicle mileage

With the incorporation of the cost-benefit trade-off, in the most positive scenario of the numerical experiment, 27% of the vehicle kilometers on the road can be removed, as opposed to 40% that was for example found in the study for New York by Santi et al. (2014). This is the scenario in which users are most tolerant to sharing a vehicle with co-riders, assuming no heterogeneity. In that case, nearly all of the existing taxi / ride-hailing users can be incentivized to using a ride-sharing service by means of a 50% discount. When reluctance to share is low on the other hand, fewer users will opt for a ride-sharing service and the distance savings will not be as targeted for. In some of those scenarios it was even found that a ride-sharing system can produce more vehicle kilometers than a ride-hailing system would need (without deadheading). Compared to the findings in previous research, there is a clear indication that with an incorporation of user cost-benefit trade-offs, the expectation for potential distance savings should be negatively adjusted.

Moreover, this study analysed how ride-sharing potential is affected by a variation in the delay and sharing tolerances across the population of potential users. Previous studies assumed that as long as detouring in a ride-sharing system is limited, all taxi requests can possibly be served with a ride-sharing service. It neglects the fact that some taxi or individual ride-hailing users will never consider ride-sharing service as an alternative to their original ride, even if a shared ride is as fast as their normal ride and they get offered a discount compared to their current fare. This study found that ignoring the existence of such 'taxi captives' can indeed lead to an overestimation of ride-sharing potential as heterogeneity in willingness to share was found to lead to a lower adoption of a ride-sharing service as well as to a smaller improvement of the distance efficiency with such a service. Also, it was discovered that the average level of service is lower when heterogeneity exists in terms of the preferences of potential users.

Another contribution of this work is that it shows that the potential of a ride-sharing system can be greatly dependent on external variables, such as the distribution of demand. The proposed ride-sharing system in this study is tested for different demand conditions rather than applied to one case, as is the case in many of the previous studies in which ride-sharing is analyzed based on real-world taxi data. The application of the model on different cases helps to prevent drawing generalized conclusions that are in fact context-specific. In this study it was for example found that ride-sharing potential is strongly dependent on the magnitude of the direction in demand. Based on the numerical experiment, ride-sharing benefits seem most promising when demand is perfectly uniform, meaning that production and attraction are balanced. Demand in reality is more likely to be concentrated due to spatial clustering of activities like work, residency and shopping. This study shows for example that, when most requests are going into the direction of the center of the grid (common in a morning peak), the performance of a ride-sharing system is much lower,

both in terms of the level of service and efficiency. Most likely, this is also the case for (evening) peaks in which demand is directed outward. The potential for improving user level of service and operational efficiency should thus always be seen in the light of a specific demand distribution.

Furthermore, this study has shown that the design of a ride-sharing service can significantly affect the potential societal benefits and user level of service. A small additional discount of 7.5% per co-rider with whom a user shares their ride at the part of the ride with highest occupancy, on top of the standard 50% discount assumed throughout this research, can for example more than double the total reduction in vehicle kilometers.

5.2.2 Profitability

The results of the numerical experiment can also be used to give a first estimation on the profitability of ride-sharing services. When comparing the profit to that of a ride-hailing service without shared rides, three main differences are relevant: (1) a lower total revenue as a result of assigned fare discounts, (2) a lower total revenue due to requests that are rejected by the system and (3) lower operational costs due to a reduction of the total vehicle distance that is needed to serve the assigned requests. The question is thus whether an improvement in the operational efficiency outweighs the loss of revenue. In the base scenario with average sharing tolerances, 76.0% of the total demand is accepted and each of the accepted requests gets a 50% discount. The revenue of a ride-sharing service in this scenario will therefore be approximately $0.50 \times 0.76 = 38\%$ of the revenue of a ride-hailing service. While the revenue is 62% lower, the vehicle distance and the related operational costs in a ride-sharing service are only 13% lower (or slightly more when is considered that ride-hailing also requires deadheading). The difference in revenue loss and cost reduction indicates that profitability of a ride-sharing service can be a serious issue. Obviously, many more variables affect the profitability of a ride-sharing service (e.g. the number of driver hours, size of the fleet etc.), but such a high-level, back-of-the-envelope derivation of the profitability shows that the prospects for a profitable ride-sharing service are not immediately bright. The operation of a ride-sharing service to generate societal gains might thus require significant subsidization.

5.2.3 Equity

The introduction of a ride-sharing service can also affect the equity of the population of users. In essence, ride-sharing can promote equity as it offers users that cannot afford taxi services a cheaper alternative. In this study it was however noted that, if heterogeneity exists in ride-sharing preferences, users with a high sharing or delay tolerance get assigned worse rides than users with low sharing and delay tolerances. The former group, possibly existing of people with a below average income, are discriminated by the service as they receive a ride with a lower level of service (e.g. longer delay) for the same price as a user with a low tolerance to delays and sharing. Equity can thus also be an issue for a ride-sharing service.

5.3 Practical recommendations

With the finding that ride-sharing efficiency and level of service can strongly depend on the behavioral attributes of potential users as well as on external variables, one of the main recommendations for authorities interested in the implementation of ride-sharing services is to carefully analyze the setting in which they wish to implement such a service, in order to find whether the goals they aim for are actually feasible given the specific situation. Feasibility studies should therefore always include an analysis of the potential demand

in an area, in terms of the spatial distribution as well as the behavioral preferences of potential users. Special attention should be given to the variation in preferences across the population, as it has been found that 'taxi captives' may exist that will not, in any circumstance, opt for a ride-sharing service.

This study proves also that not only the environment but also the design of a ride-sharing service is a determinant for the success of a ride-sharing service. Therefore, when considering the feasibility of a ride-sharing service, a recommendation for authorities or transport providers is to try out different service designs to find the general design most effective in reaching certain objectives. One of the main service variables that could be interesting to alter is the pricing mechanism, as has been found that a minor additional benefit per co-rider can significantly impact the operational efficiency of a ride-sharing service. The experienced delay should also be incorporated in the pricing mechanism to better match the financial compensation with ride-sharing discomfort. This can also help to solve equity-related issues as users that are served with a ride with a poor level of service (i.e. high delays and many co-riders) then also pay less for their ride. In this way, there is basically a form of product differentiation, with rides with high and low levels of sharing and prices that are set accordingly. The properties of the fleet, for example the size and capacity of the fleet, is another design aspect that should be investigated when initiating a ride-sharing service.

In general, the model that was developed for this work can be used to identify circumstances / markets in which ride-sharing can be successful, as has been done in this study as well. The model can however also be applied by authorities in feasibility studies for a concrete case. With the expected demand and the anticipated service design as input, it can help them to get a first idea on the societal benefits and service quality that a ride-sharing service can offer for their context. If the outlook is promising, an agent-based model can for example help to get a more detailed insight in the benefits and costs, and to find the service design that is optimal.

5.4 Future research

In this section, some of the limitations of this study are mentioned and recommendations for further research are given accordingly. The discussion is divided into the following themes: application (Subsection 5.4.1), behavioral attributes (Subsection 5.4.2), operations (Subsection 5.4.3), network (Subsection 5.4.4), assignment procedure (Subsection 5.4.5), interaction with public transit (Subsection 5.4.6) and spatial equity (Subsection 5.4.7).

5.4.1 Application

It would be interesting to apply the model to a real-world case to improve the validity of the research. As several previous studies have focused on ride-sharing in Manhattan, it would be wise to apply the model to the same network with demand based on the same taxi data set. This allows for a comparative analysis to point out specifically what the contribution is of a cost-benefit trade-off on the average level of service and the vehicle distance that can be saved with ride-sharing.

5.4.2 Behavioral attributes

Although user level of service in this work is considered more explicitly than in previous work, the cost-benefit trade-off assumed in this study is far from complete. Especially psychological factors related to the route that is followed are missing, such as a (full) vehicle passing by a request that is waiting for a pick-up (or even worse, picking up only a

part of the requests at that location), or direct route reversals (basically U-turns). Also, users do not like to visit a single location twice during their ride. Besides, there are non-route related factors that influence the user experience, such as the preference for drivers and co-riders. To improve the level of service of users, the assignment procedure should consider previously-mentioned factors, which will first require getting better insights into the decision-making of potential ride-sharing users.

When the willingness to share in this study is assumed dependent on the vehicle occupancy, it is assumed dependent only on the number of co-riders at the busiest part of the ride. In other words, it does not matter for a user if there are two other passengers on-board during the complete or only during a small part of the ride. Also, the number of different people that a single ride is shared with and the number of intermediate stops are ignored in the ride-sharing choice. All of the previous is also applicable to the way the occupancy-dependent discount is determined. It is recommended to use a weighted measure for the sharing experience rather than simply the maximum simultaneous number of co-riders during a ride.

Also, attributes that were investigated individually in this study (e.g. heterogeneity in delay tolerance and in willingness to share) can be combined to see how they interact and whether the whole is greater, equal or smaller than the sum of its parts. It is also interesting to reconsider heterogeneity in this study by analyzing the ride-sharing potential for different probability distributions to represent the delay and sharing tolerance across the population. The normal distributions that were used to describe heterogeneity in this study might not be the best way to capture variations in sharing preferences, as it for example does not exclude negative delay aversions, meaning that some (extreme) users actually prefer having a delay. A positive willingness to share is more likely as some users prefer the social interactions that come along with ride-sharing. For the delay tolerance, however, a more skewed distribution might be more suitable to describe taste heterogeneity.

5.4.3 Operations

Although the effect of several individual, collective and service attributes are analyzed in this study, there are many more that can influence the ride-sharing potential. Future research can focus on the effect of service design, like the fleet-size, vehicle capacity or discount structure. The RGV-approach that is used in this study is especially suitable for high-capacity ride-sharing. As shown by Alonso-Mora et al. (2017), ride-sharing with high-capacity vehicles can further improve societal gains of ride-sharing services. It would be interesting to consider high-capacity ride-sharing especially when the willingness to share depends on the number of passengers in the vehicle, as in one of the scenarios in this study. It is for example possible that when there are more passengers are on-board of a vehicle, an extra passenger will induce a smaller marginal disbenefit than when a user is riding alone. The incorporation of willingness to share in this study can thus improve the validity of ride-sharing potential studies focusing on high-capacity vehicles. In the light of heterogeneous sharing preferences, it might also be interesting to investigate the ride-sharing potential if a heterogeneous fleet of vehicles with different capacities is deployed, which might open up new opportunities for product differentiation. This can easily be incorporated in the current model. The effect of other attributes on ride-sharing potential that can be investigated include total demand, size of the road network and the number of stop locations, which are all limited in size in this study. The question to be answered there is how ride-sharing efficiency and level of service scale with network and demand size.

Principally, the model developed is suited for both a driver-based and an autonomous fleet of vehicles. Although in this research non-autonomous vehicles were (implicitly)

assumed, it is also possible to investigate the potential of a ride-sharing service with autonomous vehicles. It should be considered that sharing preferences might be different when there is no driver on-board the vehicle, which means that more research is needed on how users value the absence of a driver. Another aspect is that autonomous ride-sharing without a driver can possibly also create space in the car for another passenger. Dandl and Bogenberger (2018) suggest that fares of a taxi service can be lower when executed with a fleet of autonomous vehicles. This might also apply for a ride-sharing service. When analyzing the potential of a ride-sharing service with a fleet of autonomous vehicles, the current model can thus be applied with a different input for the behavioral preferences, vehicle capacity and ride fares.

5.4.4 Network

A major model enhancement can be made in the way that travel times are considered. In this study, the travel time is fixed for each link in the network and determined by the average speed on the links. Congestion effects are neglected, which can have an impact on the request-vehicle assignment and the potential of a ride-sharing system as a whole. Although the choice for static travel times facilitates the model application by allowing travel times to be extracted from a predefined look-up table, the assumption is unrealistic and can significantly impact ride-sharing performance. Moreover, the model currently does not consider parking availability at stop locations, which can also lead to a higher travel time if such a spot is not available. Incorporation of congestion effects would significantly advance the validity of the results.

5.4.5 Assignment procedure

The objective function of the assignment of requests to vehicles should also be rethought. The current secondary objective (with the primary objective maximizing the number of accepted requests) is to maximize the sum of benefits over all accepted users. In the application of the model it was noted that this objective function can cause passengers to stay on-board of a vehicle after reaching their destination to wait for the pick-up of a new request as to collect a larger fare discount (whereby the additional discount outweighs the additional delay and sharing discomfort). This behaviour is not desired in terms of the efficiency of the system nor realistic. As the discount is transferred from the operator / authority to the user, the question is also whether such an objective makes sense from a system perspective. Accounting for the system optimum, it might be better to maximize or minimize an indicator in which different actors do not have an opposite interest, unlike the fare discount. Possible alternative (primary or secondary) objectives that can be considered in future research include a minimized sum of disbenefits of accepted users or a minimized total vehicle distance. Alternatively, constraints can be added to the optimization problem to prevent undesired behavior, such as passengers staying on-board of the vehicle after reaching their drop-off location with the aim of getting a higher fare discount.

The efficiency of a ride-sharing service can potentially also be improved by rejecting certain requests. This is for example the case for requests that are destined for a location far away from where new demand is expected. This prevents vehicles ending up in location without demand, which leads to significant deadheading or rebalancing to pick-up the next request. Thus, it would be interesting to incorporate the 'attractiveness' of a request's destination in the objective function when assigning vehicles to requests. Also, in some cases it might be better to reject passengers with very low sharing tolerances as they will occupy a vehicle but not allow for any other passengers in their vehicle. From a system perspective it might be better to keep the vehicle capacity available for next assignment

iterations, where it can possibly be assigned to serve multiple more flexible users. In other words, the efficiency of a ride-sharing service can potentially be improved by anticipating on future demand, both spatially and in terms of their sharing attitudes. Future research is desired to see this confirmed.

Another issue related to the assignment is that requests are kept in the pool of available requests as long as they can be served by a direct ride. It means that requests potentially wait a long time before finally being notified that they cannot be served by the ride-sharing service, while in reality users demand quick feedback on whether a ride is found or not. In the current assignment procedure there is a penalty for each rejected request, but the waiting time of those rejected requests is not directly considered. In future research it is therefore recommended to allow for a shorter time window of assignment for each incoming request by adding a constraint on the maximum waiting time for assignment (note that this is not equal to the waiting time before pick-up).

5.4.6 Public transit interaction

The validity of the results can be improved furthermore by considering ride-sharing in the context of a wider mode choice for a request seeking to go from A to B. In this research, societal benefits of ride-sharing are determined by considered shared rides as a replacement of (individual) taxi rides only, while it might be possible that ride-sharing will also lure travellers away from public transport, thereby leading to more rather than less vehicle kilometers on the road. It would be interesting to find out more about how ride-sharing positions itself between private and public transit, so that more valid conclusions can be drawn on the environmental impacts of ride-sharing services. For this purpose, it would be of added value to actually incorporate mode choice into the model. Currently, the only choice that users make is whether to opt for a ride-sharing service or not (in which the choice is actually made for them by the system). It is thereby unknown what happens to requests that are rejected by the ride-sharing service.

It should also be considered that 'rebound' effects can take place, whereby a reduction in vehicle kilometers on the road as a result of more shared rides leads to lower travel times in the network and ultimately to an increase of the relative attractiveness of the car compared to public transit (Yin, Liu, Coulombel, & Viguié, 2018).

On the other hand, it is also interesting to consider how ride-sharing can in fact complement public transit by offering cheap and efficient first and last mile transport to and from transit stops. Shared vehicles have also been found to be able to reduce the space needed for access and egress facilities at train stations (Stam, 2019).

5.4.7 Spatial equity

This research did not specifically consider the spatial equity of users, or whether incoming requests in different parts of the network have approximately equal chances of being picked-up as well as comparable average delays. It might for example be the case that requests in the outer ring are not picked-up at all, which would raise issues about the accessibility of the more remote areas in a ride-sharing setting. It would be therefore be interesting to compute performance indicators or metrics specific to the location in the network, for example the average delay per origin.

References

Alonso-Mora, J., Samaranayake, S., Wallar, A., Frazzoli, E., & Rus, D. (2017). On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(3), 462-467. doi: 10.1073/pnas.1611675114

Asensio, J., & Matas, A. (2008). Commuters' valuation of travel time variability. *Transportation Research Part E: Logistics and Transportation Review*, *44*(6), 1074–1085.

Bates, J., Polak, J., Jones, P., & Cook, A. (2001). The valuation of reliability for personal travel. *Transportation Research Part E: Logistics and Transportation Review*, *37*(2-3), 191–229.

Bian, Z., & Liu, X. (2018). A detour-based pricing mechanism for first-mile ridesharing in connection with rail public transit. In *2018 Joint Rail Conference.*

Bimpikis, K., Candogan, O., & Saban, D. (2016). Spatial pricing in ride-sharing networks. doi: 10.2139/ssrn.2868080

Börjesson, M., Eliasson, J., & Franklin, J. P. (2012). Valuations of travel time variability in scheduling versus mean–variance models. *Transportation Research Part B: Methodological*, *46*(7), 855–873.

Cats, O., West, J., & Eliasson, J. (2016). A dynamic stochastic model for evaluating congestion and crowding effects in transit systems. *Transportation Research Part B: Methodological*, *89*, 43–57.

Cheng, S.-F., Nguyen, D. T., & Lau, H. C. (2012). A mechanism for organizing last-mile service using non-dedicated fleet. In *Proceedings of the The 2012 IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technology* (pp. 85–89).

Cici, B., Markopoulou, A., Frías-Martínez, E., & Laoutaris, N. (2013). Quantifying the potential of ride-sharing using call description records. In *Proceedings of the 14th workshop on mobile computing systems and applications* (p. 17).

Clewlow, R. R., & Mishra, G. S. (2017). Disruptive transportation: The adoption, utilization, and impacts of ride-hailing in the united states. *University of California, Davis, Institute of Transportation Studies, Davis, CA, Research Report UCD-ITS-RR-17-07*.

Correia, G., & Viegas, J. (2011). Carpooling and carpool clubs: Clarifying concepts and assessing value enhancement possibilities through a stated preference web survey in Lisbon, Portugal. *Transportation Research Part A: Policy and Practice*, *45*(2), 81-90. doi: 10.1016/j.tra.2010.11.001

Daganzo, C. F. (2010). Structure of competitive transit networks. *Transportation Research Part B: Methodological*, *44*(4), 434–446.

Dandl, F., & Bogenberger, K. (2018). Comparing future autonomous electric taxis with an existing free-floating carsharing system. *IEEE Transactions on Intelligent Transportation Systems*(99), 1–11.

DeLoach, S. B., & Tiemann, T. K. (2012). Not driving alone? American commuting in the twenty-first century. *Transportation*, *39*(3), 521–537.

Dueker, K. J., Bair, B. O., & Levin, I. P. (1977). Ride sharing: Psychological factors. *ASCE Transp Eng J*, *103*(6), 685-692.

Edvardsson, B. (1998). Causes of customer dissatisfaction-studies of public transport by the critical-incident method. *Managing Service Quality: An International Journal*, *8*(3), 189–197.

Ehsani, S., Guglielmetti, R., Laumanns, M., Markov, I., & de Souza, R. (2018). *Simulation-based design and analysis of on-demand mobility services* (Tech. Rep.). Bestmile.

Erlander, S., & Stewart, N. F. (1990). *The gravity model in transportation analysis: theory and extensions* (Vol. 3). VSP.

Fagnant, D., & Kockelman, K. (2014). The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios. *Transportation Research Part C: Emerging Technologies*, *40*, 1-13. doi: 10.1016/j.trc.2013.12.001

Fiedler, D., Čertickỳ, M., Alonso-Mora, J., & Čáp, M. (2018). The impact of ridesharing in mobility-on-demand systems: Simulation case study in prague. In *IEEE International Conference on Intelligent Transportation Systems.*

Friman, M., & Gärling, T. (2001). Frequency of negative critical incidents and satisfaction with public transport services. *Journal of Retailing and Consumer Services*, *8*(2), 105-114.

Gao, K., Sun, L., Tu, H., & Li, H. (2018). Heterogeneity in valuation of travel time reliability and in-vehicle crowding for mode choices in multimodal networks. *Journal of Transportation Engineering, Part A: Systems*, *144*(10).

Gemeente Amsterdam. (2019). *Wat kost een rit in een taxi?* Retrieved January 23, 2019, from https://www.amsterdam.nl/veelgevraagd/?caseid=\%7B7DC3ADE8-163A-4201-ABC4-4ACC1830EB88\%7D

Handke, V., & Jonuschat, H. (2012). *Flexible ridesharing: new opportunities and service concepts for sustainable mobility.* Springer Science & Business Media.

Hennessy, D. A., & Wiesenthal, D. L. (1999). Traffic congestion, driver stress, and driver aggression. *Aggressive Behavior: Official Journal of the International Society for Research on Aggression*, *25*(6), 409–423.

Hensher, D. A., Stopher, P., & Bullock, P. (2003). Service quality - developing a service quality index in the provision of commercial bus contracts. *Transportation Research Part A: Policy and Practice*, *37*(6), 499–517.

Kfzteile24. (2019). *Best and worst cities to drive 2017.* Retrieved January 23, 2019, from https://www.kfzteile24.de/best-and-worst-cities-to-drive-usd

König, A., & Axhausen, K. W. (2002). The reliability of the transportation system and its influence on the choice behaviour. In *Swiss Transport Research Conference (STRC), Monte Verità/Ascona.*

Li, Z., Hensher, D. A., & Rose, J. M. (2010). Willingness to pay for travel time reliability in passenger transport: A review and some new empirical evidence. *Transportation Research Part E: Logistics and Transportation review*, *46*(3), 384–403.

Liu, T., Krishnakumari, P., & Cats, O. (2019). *Exploring demand patterns of a ride-sourcing service using spatial and temporal clustering.* (Unpublished)

Lunden, I. (2016, May). Uber says that 20% of its rides globally are now on UberPool. *Tech Crunch*. Retrieved from https://techcrunch.com/2016/05/10/uber-says-that-20-of-its-rides-globally-are-now-on-uber-pool/?guccounter=1

Ma, S., Zheng, Y., & Wolfson, O. (2013). T-share: A large-scale dynamic taxi ridesharing service. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)* (pp. 410–421).

Morales Sarriera, J., Escovar Álvarez, G., Blynn, K., Alesbury, A., Scully, T., & Zhao, J. (2017). To share or not to share: Investigating the social aspects of dynamic

ridesharing. *Transportation Research Record: Journal of the Transportation Research Board*, *2605*, 109-117.

Pant, P., & Harrison, R. M. (2013). Estimation of the contribution of road traffic emissions to particulate matter concentrations from field measurements: a review. *Atmospheric environment*, *77*, 78–97.

Park, Y., Chen, N., & Akar, G. (2018). Who is interested in carpooling and why: The importance of individual characteristics, role preferences and carpool markets. *Transportation Research Record*. doi: 10.1177/0361198118756883

Santi, P., Resta, G., Szell, M., Sobolevsky, S., Strogatz, S., & Ratti, C. (2014). Quantifying the benefits of vehicle pooling with shareability networks. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(37), 13290-13294. doi: 10.1073/pnas.1403657111

Stam, B. (2019). *Access/egress facilities at railway stations* (Unpublished master's thesis). Delft University of Technology.

Tachet, R., Sagarra, O., Santi, P., Resta, G., Szell, M., Strogatz, S., & Ratti, C. (2017). Scaling law of urban ride sharing. *Scientific Reports*, *7*. doi: 10.1038/srep42868

Teal, R. (1987). Carpooling: Who, how and why. *Transportation Research Part A: General*, *21*(3), 203-214. doi: 10.1016/0191-2607(87)90014-8

Wang, X., Dessouky, M., & Ordonez, F. (2016). A pickup and delivery problem for ridesharing considering congestion. *Transportation letters*, *8*(5), 259–269.

Wardman, M. (2004). Public transport values of time. *Transport policy*, *11*(4), 363–377.

Xiong, C., Hetrakul, P., & Zhang, L. (2014). On ride-sharing: A departure time choice analysis with latent carpooling preference. *Journal of Transportation Engineering*, *140*(8). doi: 10.1061/(ASCE)TE.1943-5436.0000692

Yap, M., Cats, O., & van Arem, B. (2018). Crowding valuation in urban tram and bus transportation based on smart card data. *Transportmetrica A: Transport Science*, 1–20.

Yin, B., Liu, L., Coulombel, N., & Viguié, V. (2018). Appraising the environmental benefits of ride-sharing: The paris region case study. *Journal of Cleaner Production*, *177*, 888–898.

Zhang, H., & Zhao, J. (2018). Mobility sharing as a preference matching problem. *IEEE Transactions on Intelligent Transportation Systems*.

Appendix A

List of results

Scenario	Acceptance rate	Average total delay per passenger (s)	Average waiting time per passenger (s)	Average in-vehicle delay per passenger (s)	Standard deviation of total delay per passenger (s)	Average passenger ratio delay / direct travel time	Standard dev. passenger ratio delay / direct travel time	Average passenger ratio waiting time / total delay	Standard dev. passenger ratio waiting time / total delay	Average number of stops per passenger	Standard deviation of number stops per passenger	Average number of co-riders	Ratio of passenger time with 0 co-riders	Ratio of passenger time with 1 co-rider	Ratio of passenger time with 2 co-riders
U_30_3_50	76.0%	126.8	92.1	34.7	98.9	29.5%	15.3%	86.6%	26.2%	0.95	1.20	0.67	50.5%	31.8%	17.7%
U_18_3_50	88.2%	200.5	122.4	78.1	164.6	48.7%	26.8%	78.5%	30.3%	1.56	1.67	1.05	30.9%	33.0%	36.1%
U_24_3_50	81.8%	158.5	103.6	55.0	132.0	37.3%	19.9%	82.2%	29.2%	1.24	1.38	0.85	40.2%	34.8%	25.0%
U_36_3_50	70.1%	106.2	84.8	21.4	76.6	24.6%	12.2%	90.7%	22.8%	0.79	1.07	0.53	58.0%	30.6%	11.4%
U_42_3_50	64.2%	93.2	77.1	16.0	65.2	21.1%	10.2%	92.3%	21.3%	0.67	0.94	0.45	63.3%	28.9%	7.8%
U_30_1_50	98.8%	219.5	138.2	81.3	136.4	61.3%	26.6%	75.4%	30.5%	1.83	1.65	1.12	24.9%	37.8%	37.3%
U_30_2_50	97.6%	169.6	114.0	55.6	117.6	45.3%	19.1%	81.1%	29.3%	1.40	1.52	0.92	35.3%	37.2%	27.5%
U_30_4_50	46.4%	102.7	82.9	19.8	78.5	20.5%	11.4%	91.7%	22.1%	0.66	0.96	0.45	63.8%	27.4%	8.8%
U_30_5_50	25.4%	88.0	74.7	13.2	63.7	15.2%	8.5%	94.1%	18.5%	0.39	0.70	0.27	76.3%	20.8%	2.9%
U_H_3_50	75.9%	154.2	103.5	50.8	190.4	36.4%	34.7%	86.1%	27.0%	1.08	1.46	0.77	45.9%	31.7%	22.4%
U_30_H_50	66.8%	162.8	106.4	56.4	131.7	42.1%	28.7%	82.0%	29.6%	1.19	1.41	0.77	44.8%	34.0%	21.2%
U_30_D_50	94.1%	127.3	116.8	10.5	76.4	35.8%	15.2%	95.5%	16.3%	0.31	0.68	0.25	75.9%	22.9%	1.2%
U_30_3_D	82.1%	221.7	118.7	103.0	163.0	54.0%	27.4%	70.3%	32.6%	1.97	1.76	1.22	23.1%	31.8%	45.1%
MD_30_3_50	68.0%	117.7	92.54	25.1	86.9	29.1%	14.7%	90.1%	22.6%	0.79	1.06	0.58	55.9%	30.4%	13.7%
SD_30_3_50	62.9%	130.7	106.53	24.2	84.3	31.9%	13.6%	90.3%	22.1%	0.80	1.03	0.61	53.4%	32.6%	14.0%

Table A.1: KPIs representing ride-sharing level of service for all scenarios

Scenario	**Total vehicle movement distance (km)**	**Total vehicle transportation distance (km)**	**Total deadheading distance (km)**	**Total rebalancing distance (km)**	**Total effective vehicle transportation distance (km)**	**Total distance savings (km)**	**Vehicle transportation distance ratio**	**Effective vehicle transportation distance ratio**	**Corrected effective vehicle transportation distance ratio**	**Empty vehicle rebalancing distance ratio**	**Average vehicle occupancy**	**Ratio of non-empty vehicle time with occupancy 1**	**Ratio of non-empty vehicle time with occupancy 2**	**Ratio of non-empty vehicle time with occupancy 3**	**Model run time (h)**
U_30_3_50	6,488	5,588	900	156	7,471	983	0.861	1.15	1.34	0.173	1.38	69.8%	22.0%	8.2%	0.24
U_18_3_50	6,357	5,690	667	78	8,076	1,719	0.895	1.27	1.42	0.117	1.68	52.0%	27.8%	20.3%	5.06
U_24_3_50	6,456	5,657	799	133	7,799	1,343	0.876	1.21	1.38	0.166	1.52	61.0%	26.4%	12.7%	0.40
U_36_3_50	6,350	5,453	898	156	7,093	743	0.859	1.12	1.30	0.174	1.30	75.2%	19.8%	4.9%	0.20
U_42_3_50	6,142	5,218	925	187	6,672	530	0.849	1.09	1.28	0.202	1.24	78.8%	18.0%	3.2%	0.19
U_30_1_50	6,389	5,867	522	9	8,666	2,277	0.918	1.36	1.48	0.017	1.78	44.3%	33.6%	22.1%	4.87
U_30_2_50	6,853	6,133	720	56	8,616	1,763	0.895	1.26	1.41	0.077	1.59	56.0%	29.5%	14.5%	0.52
U_30_4_50	4,966	4,330	636	106	5,445	479	0.872	1.10	1.26	0.167	1.24	79.3%	17.0%	3.7%	0.16
U_30_5_50	3,321	2,958	364	37	3,481	160	0.891	1.05	1.18	0.102	1.14	87.0%	11.9%	1.1%	0.13
U_H_3_50	6,389	5,556	833	152	7,378	989	0.870	1.16	1.33	0.182	1.45	66.3%	22.9%	10.8%	2.16
U_30_H_50	5,570	4,949	621	84	6,508	939	0.889	1.17	1.32	0.135	1.45	65.0%	24.7%	10.3%	0.33
U_30_D_50	8,620	7,313	1,307	156	8,375	-245	0.848	0.97	1.15	0.119	1.14	86.5%	13.1%	0.5%	0.25
U_30_3_D	5,687	5,186	501	125	7,656	1,970	0.912	1.35	1.48	0.249	1.85	42.7%	29.4%	27.8%	0.82
MD_30_3_50	6,000	4,829	1,171	393	6,338	339	0.805	1.06	1.31	0.335	1.32	73.9%	20.1%	6.0%	0.24
SD_30_3_50	6,288	4,552	1,736	910	5,968	-320	0.724	0.95	1.31	0.524	1.35	71.8%	21.9%	6.3%	0.34

Table A.2: KPIs representing ride-sharing efficiency and computational complexity for all scenarios

Appendix B

Demand distribution

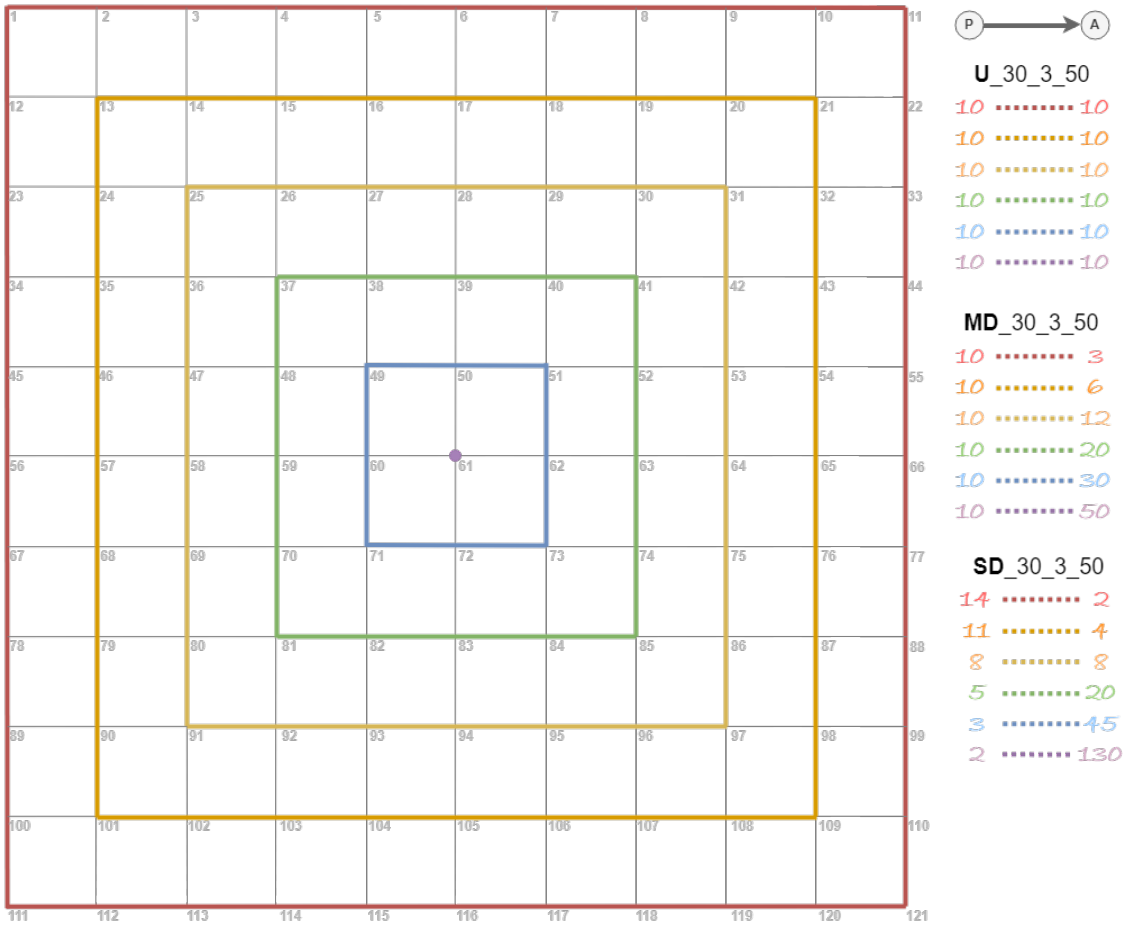


Figure B.1: Production and attraction per hour per node in different scenarios