Ride-Sharing Efficiency and Level of Service under Alternative Demand, Behavioral and Pricing Settings

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: 6339 words + 3 table(s) × 250 = 7089 words

Submission Date: July 31, 2019

ABSTRACT

Previous studies examining ride-sharing potential assumed that rides can be shared as long as the incurred delay does not exceed a certain threshold. Conversely, we formulate willingness to share as a compensatory cost function at the individual passenger level. The latter considers trade-offs between delays caused by detours, travel discomfort related to sharing a vehicle and a fare discount associated with a shared ride. Next to finding how these behavioral preferences and the offered discount structure affect the efficiency and level of service of ride-sharing services, the effect of directionality in demand is considered. A graph-based approach is applied to perform an efficient assignment of vehicles to requests. We test the model on an experiment representing an urban context. Our findings suggest that service performance is strongly dependent on users' willingness to share and somewhat less strongly on users' tolerance to delays. Implementation of a ride-sharing service is most successful when directionality in demand is low, while ride-specific discounts can be effective in maximizing societal benefits.

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

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. In May 2019, in New York alone nearly 125,000 trips were made using a ride-sharing service (*1*). Users of such platforms allow other travellers to join their ride, even accepting small detours, which offers opportunities for a more efficient utilization of road space and consequently reduced congestion levels, improved air quality and better traffic safety (*2, 3*).

Whether ride-sharing in practice can live up to these expectations is uncertain. Ride-sharing might for example substitute public transit rather than individual rides, leading to more rather than less vehicle kilometers on the roads. Moreover, 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 pick up and drop off other passengers, social issues related to sharing are an important deterrent for potential users. This includes a lack of privacy (*4, 5*), a feeling of dependence and a fear of having negative social interactions with other users (*6, 7*). Social discomfort might help explain why in January 2019 only 25% of Uber's rides in New York were made with its ride-sharing service UberPool (*1*). Also, as ride-sharing efficiency is dependent on the compatibility of trip requests, a ride-sharing trip is not necessarily shared in practice. In fact, a study on the impacts of ride-hailing in Toronto found that in only 18% of all ride-sharing trips, a rider is actually matched to another rider (*8*).

Previous quantitative studies on societal benefits of ride-sharing nevertheless showed promising results. A study by Ma et al. (*9*) for example stated that if the current fleet of taxis in New York allows for shared rides, while users accept a maximum extra travel time of 5 minutes for their ride, 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 concluded that, given the same setting, 32% of the total current vehicle distance of taxis can become unnecessary (*10*). When ride-sharing is executed with high-capacity vehicles of up to ten seats, less than one sixth of the size of the current taxi fleet can serve 98% of the original requests with a maximum delay of 3.5 minutes per passenger (*11*). While the previously mentioned studies focused on ride-sharing in New York, Tachet et al. (*12*) found that ride-sharing potential also exists for cities with a lower density than New York.

A common shortcoming of previous ride-sharing studies is that they account for only one of three elements in the complex trade-off that users typically make between a delay, discomfort and a discounted fare when considering a shared ride. Each of these studies simplify ride-sharing choice by considering only a maximum allowed delay, meaning that they implicitly assume that users are principally willing to ride-share even if it gives them no benefit and a (relatively small) delay. Conversely, in this study we explicitly consider a trade-off of travel attributes by accounting for discomfort stemming from sharing and discounted ride fares. The question to be answered is how the operational efficiency of a ride-sharing service and the level of service that it offers to its users depends on the attitudes of potential users towards delays and the presence of co-riders. At the same time, the incorporation of the cost-benefit trade-off at the individual passenger level allows us to derive first implications for the design of an effective discount structure to boost ride-sharing adoption and consequently reduce the total vehicle distance on the road.

A final 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 how efficiency gains and level of service depend on the extent of directionality in demand.

This paper is structured into four parts. Firstly, a detailed description is given of the methodology that was developed. This is followed by a motivation for the design of the numerical experiment. The results for the different scenarios in the expeiment 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

In order to simulate the operations of a ride-sharing service and determine the total vehicle movement and service quality, several modelling approaches have been taken in previous research for assigning passenger requests to vehicles. For example, one of the earlier studies considered ride-sharing assignment as a Vehicle Routing Problem (VRP) with time windows (*9*). Each incoming request was thereby individually allocated to a vehicle using a greedy algorithm. A study by Santi et al. (*10*) 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 (*11*). A request-group-vehicle (RGV) graph represents memberships of request groups 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 interact (*13, 14*).

Of the different approaches, the RGV-approach is found most suitable given the modeling requirements of this study. First and foremost, it allows to model ride-sharing with more than two passengers per vehicle. Moreover, the assignment can yield an optimal solution. By representing the assignment in a graph form, the problem's computational complexity is minimized, which is useful given our interest in assessing different scenarios 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, requests are assigned to vehicles at fixed intervals, whereby each iteration consists of nine main steps.

Group-vehicle feasibility

Key to the approach is the way in which is determined whether a specific vehicle v in fleet V, with passengers P_v on-board, can serve a group of pending requests Z. 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 a user constraint. The vehicle constraint 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, similar to the approach of Alonso-Mora et al. (*11*).

The user constraint in this study is more complex and considers for each individual request r in Z or P_v, as has been explained in the introduction, a trade-off between ride-sharing benefits and costs. The ride-sharing benefit of r consists of the total fare discount and is thus 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

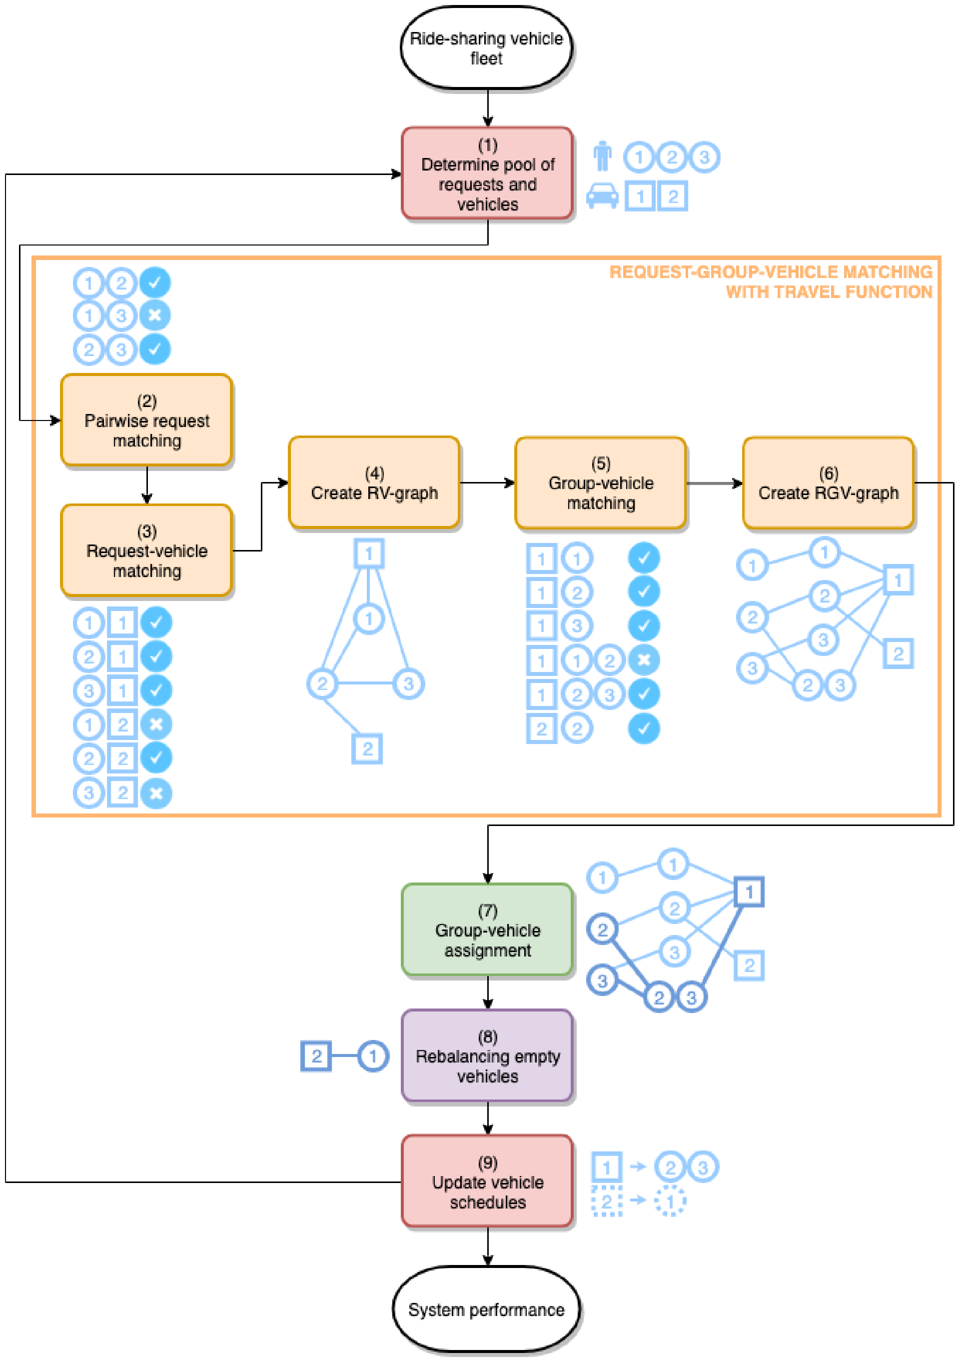

FIGURE 1 Overview of the methodology, including an example (in blue) with three requests and two (empty) vehicles

fixed or be inverse with the level of service of the experienced ride. Disbenefits on the other hand follow from extra travel time and additional discomfort associated with sharing a vehicle. The total disbenefit of a ride thereby depends on how users perceive both attributes, in this study expressed as delay aversion β_r and reluctance to share γ_r. These two parameters indicate what fare discount users require for an hour of delay and for sharing a vehicle with other riders (all other factors being the same), respectively. A variable α_r is added to indicate how much more negatively users experience the waiting time before pick-up compared to the in-vehicle delay. As in (*11*), the waiting time of a request wt_r is calculated 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 an immediate pick-up and a direct route with travel time tt_{o_r,d_r} between origin o_r and destination d_r. The total net benefit of r is consequently specified as:

$$b_r = \pi_r \cdot c_r - (t_r^d - t_r^*) \cdot \beta_r - (t_r^{pu} - t_r^r) \cdot \alpha_r - \gamma_r \quad (1)$$

Route S_v is assumed to satisfy the user level of service constraint only if the net benefit b_r of each request in Z and P_v is positive. If there exists at least one feasible route $S_v \in K_v$ to serve Z, Z and v form a feasible match.

RGV-matching

The set of available requests for assignment R comprises of 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 testing all group-vehicle combinations for feasibility, the process is divided into three successive matching steps, similar to the approach of Alonso-Mora et al. (*11*).

The purpose of the first of these three steps (step 2 in Figure 1) is to find whether two requests in R can share a ride, given 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 current location and available seats. The result of both steps can be combined and stored in 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. Step 5 checks whether a feasible route S_v to satisfy a group-vehicle combination within user and capacity constraints actually exists. 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 for the optimal route S_v^*. If a group-vehicle combination $g - v$ is not feasible, b_{gv} is assigned a very large penalty (so-called big M), to ensure that this combination is not chosen during 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:

$$\max \quad \sum_{g\in G}\sum_{v\in V}(b_{gv} + \sqrt{M}\cdot\sum_{r\in R} a_{rg})\cdot x_{gv} \tag{2}$$

$$\text{s.t.} \quad \sum_{g\in G} x_{gv} \leq 1,\ \forall v \in V \tag{3}$$

$$\sum_{g\in G}\sum_{v\in V} a_{rg}\cdot x_{gv} \leq 1,\ \forall r \in R \tag{4}$$

$$x_{gv} = [0,1],\ \forall g \in G, v \in V \tag{5}$$

The objective function (Equation 2) 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 for each request r in an assigned request group g. The sum of those rewards should, however, never be so large that it can overpass the big M penalty assigned to infeasible group-vehicle combinations in the objective function. Therefore, the reward per request group is set to $\sqrt{M}$. The total benefit of a group-vehicle combination $g-v$ thus consists of the summed net benefit for all requests and passengers in this group plus a large reward $\sqrt{M}$ for each request that is part of this group.

The Integer Linear Problem contains three types of constraints guaranteeing respectively a maximum assignment of one request group g to each vehicle v (Equation 3), that each request r is not part of multiple assigned request groups in G (Equation 4), and that each decision variable is binary (Equation 5).

Rebalancing

Unassigned vehicles are assigned to move in the direction of unassigned requests to anticipate on new requests appearing in areas that currently have undersupply (step 8 in Figure 1). In this study, the rebalancing procedure of Alonso-Mora et al. (*11*) is used. Its objective is to minimize the total empty vehicle rebalancing distance while ensuring a maximum number of vehicles to be assigned to rebalance.

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).

KPIs

The performance of a ride-sharing service is measured using several Key Performance Indicators (KPIs) capturing both its level of service (LoS) towards users and its operational efficiency for authorities and service 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 fares in this case are not considered as KPI, since they are directly dependent on π_r and are thus model input. The main LoS KPIs in this study include the acceptance rate (i.e. the percentage of fulfilled requests out of the total demand, thereby an indicator for reliability), the delay as percentage of the direct travel time (indicating travel time), the average number of stops per passenger, and the share of passenger time with a specific number of co-riders on-board (indicating comfort).

An authority on the other hand will be most interested in the share of the vehicle distance that can be reduced with ride-sharing. A suitable KPI to express distance efficiency is the gross effective vehicle transportation distance ratio, which is defined as the sum of the shortest OD-distance of accepted requests (= 'effective vehicle distance') divided by the total vehicle movement distance (*20*). The total vehicle movement distance (or vehicle mileage) consists of the transportation distance (the vehicle distance with at least one passenger on-board), the deadheading distance (the total empty vehicle distance to pick-up a new request) and the empty vehicle rebalancing distance. Also, a net effective vehicle transportation distance ratio is defined. This ratio accounts for the fact that the summed shortest distance of accepted requests, which represents the distance needed when sharing is not allowed, excludes deadheading. For a more fair comparison, the deadheading distance is therefore subtracted from the total vehicle movement in the net effective vehicle transportation distance ratio. For operators, the average vehicle occupancy while a vehicle is transporting passengers is an important efficiency KPI.

Implementation

The simulation model is implemented in Python, using the open-source library Numpy 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 is used to compute the shortest path between a pair of locations in the road network, after which the corresponding travel time is stored in a look-up table. The optimization problems that are part of the group-vehicle assignment and rebalancing procedure are solved using the MOSEK Optimizer API.

EXPERIMENTAL DESIGN

An experiment is constructed to test the effect of users' behavioral preferences, the discounting policy and the spatial distribution of demand on ride-sharing performance in an urban context.

Set-up

The assumed grid network consists of 121 nodes with a link distance of 500 meters, thereby 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 or the Inner Ring of Berlin. The intermediate stop distance is relatively large, whereby we implicitly assume 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 all cases 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. A gravity model (*23*) is applied to create a 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 $\upsilon = 3$, initially evenly distributed over the network. Ride fares are set based on 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 each of the starting conditions.

Scenarios

The effect of 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). 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 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, a median value of γ_r = €3 is assumed.

TABLE 1 Scenario design ($\alpha_\mathbf{r} = \mathbf{0.5} \cdot \beta_\mathbf{r}$)

#	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 ride-sharing rides, independent of whether sharing actually occurred throughout the ride. In the alternative scenario, a similar discount of 50% 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 in demand is tested using 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 of the network and more attraction in the central nodes, intended to mimic a morning peak pattern.

RESULTS

This section shows the effect of the preferences of potential ride-sharing users, the applied discount structure and the demand distribution on the level of service and efficiency of a ride-sharing service. The comprehensive list of KPI values is presented in Tables 2 (level of service) and 3 (efficiency). The three subsections that follow each go into detail on the effect of one of the investigated variables.

TABLE 2 Level of service KPI values for each scenario

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%	126.8	92.1	34.7	30%	0.95	50%	32%	18%
U_18_3_50	88%	200.5	122.4	78.1	49%	1.56	31%	33%	36%
U_24_3_50	81%	158.5	103.6	55.0	37%	1.24	40%	35%	25%
U_36_3_50	70%	106.2	84.8	21.4	25%	0.79	58%	31%	11%
U_42_3_50	64%	93.2	77.1	16.0	21%	0.67	63%	29%	8%
U_30_1_50	99%	219.5	138.2	81.3	61%	1.83	25%	38%	37%
U_30_2_50	98%	169.6	114.0	55.6	45%	1.40	35%	37%	28%
U_30_4_50	46%	102.7	82.9	19.8	21%	0.66	64%	27%	9%
U_30_5_50	25%	88.0	74.7	13.2	15%	0.39	76%	21%	3%
U_30_3_D	82%	221.7	118.7	103.0	54%	1.97	23%	32%	45%
MD_30_3_50	68%	117.7	92.5	25.1	29%	0.79	56%	30%	14%
SD_30_3_50	63%	130.7	106.5	24.2	32%	0.80	53%	33%	14%

TABLE 3 Efficiency KPI values for each scenario

Scenario	Total vehicle movement distance (km)	Total vehicle transportation distance (km)	Total deadheading distance (km)	Total rebalancing distance (km)	Gross effective vehicle transportation distance ratio	Net 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	70%	22%	8%
U_18_3_50	6,357	5,690	667	78	1.27	1.42	0.117	1.68	52%	28%	20%
U_24_3_50	6,456	5,657	799	133	1.21	1.38	0.166	1.52	61%	26%	13%
U_36_3_50	6,350	5,453	898	156	1.12	1.30	0.174	1.30	75%	20%	5%
U_42_3_50	6,142	5,218	925	187	1.09	1.28	0.202	1.24	79%	18%	3%
U_30_1_50	6,389	5,867	522	9	1.36	1.48	0.017	1.78	44%	34%	22%
U_30_2_50	6,853	6,133	720	56	1.26	1.41	0.077	1.59	56%	30%	15%
U_30_4_50	4,966	4,330	636	106	1.10	1.26	0.167	1.24	79%	17%	4%
U_30_5_50	3,321	2,958	364	37	1.05	1.18	0.102	1.14	87%	12%	1%
U_30_3_D	5,687	5,186	501	125	1.35	1.48	0.249	1.85	43%	29%	28%
MD_30_3_50	6,000	4,829	1,171	393	1.06	1.31	0.335	1.32	74%	20%	6%
SD_30_3_50	6,288	4,552	1,736	910	0.95	1.31	0.524	1.35	72%	22%	6%

Effect of behavioral preferences

The acceptance rate (Figure 2a) is found to increase when reluctance to share γ_r 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) increases more than linearly when γ_r decreases, as well as passengers' waiting time and in-vehicle delay (Figure 2c). It is found that rides are hardly shared (i.e. the average vehicle occupancy is 1.14) if users are very sensitive to sharing with other passengers, meaning that the average in-vehicle delay is 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 willingness to share. It can be explained by the finding that the total effective vehicle distance (due to more requests served) increases more than the total vehicle movement distance when users are more flexible (Figure 2d), as a result of a more efficient assignment of vehicles to requests. 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 drop off other passengers. If γ_r = €1 for example, the average effective passenger distance per total vehicle kilometer in the system (including transportation and deadheading) rises to 1.36 kilometer.

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

FIGURE 2 Effect of reluctance to share γ_r on (a) acceptance rate, (b) vehicle occupancy and number of intermediate stops, (c) average passenger delay, and (d) total vehicle movement, total transportation distance and effective transportation distance

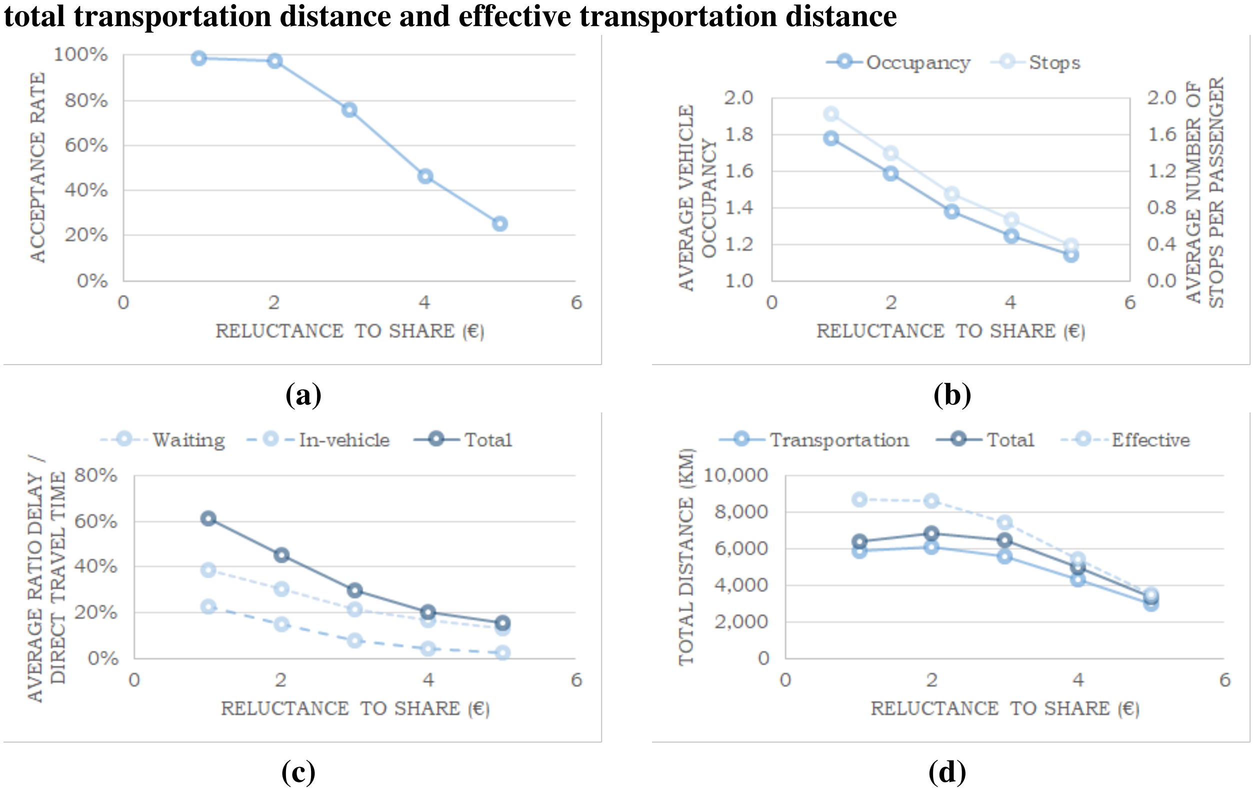

Effect of discount mechanism

As expected, when users receive an additional 7.5% discount per 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 3a), and a similar increase is found in the passenger time in a full vehicle (from 17,7% to 45.1% of the total passenger time). By utilizing the available vehicle capacity more efficiently, the acceptance rate (also Figure 3a) increases from 76.0% to 82.1%, although at the cost of a higher average delay (Figure 3b). 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% vs 8.1% of the direct travel time). Also, the average waiting time of requests is marginally higher 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 (gross) effective vehicle transportation distance ratio increases from 1.15 (*U_30_3_50*) to 1.35 when an additional 7.5% discount is awarded per co-rider (*U_30_3_D*). In combination with a higher acceptance rate, relatively large distance savings (Figure 3c) can thus be achieved with an additional occupancy-dependent discount. The distance that a ride-sharing service can save 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 drop off other passengers.

FIGURE 3 Effect of discount structure on (a) acceptance rate and average vehicle occupancy, (b) average passenger delay, and (c) total vehicle movement distance and effective transportation distance; the difference indicating distance savings

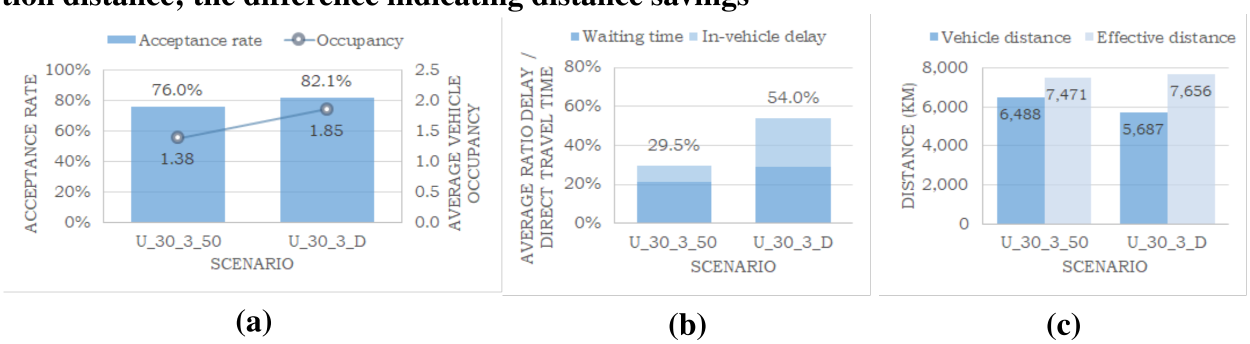

Effect of demand distribution

More directionality in demand leads to more requests being rejected by the ride-sharing service (37.1% when demand is strongly directed versus 24.0% when demand is perfectly uniform, as shown by Figure 4a). If demand is perfectly uniform, the average vehicle occupancy of vehicles in revenue mode (also Figure 4a) is 1.38 and the average passenger delay (Figure 4b) is 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 in 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 relatively short in the center, where attraction exceeds production (Figure 4c), compared to the nodes in the periphery of the network. Since only the majority of requests originates here, the average waiting time is mainly determined by requests originating outside the center, where production exceeds attraction. In these nodes, the average waiting time is nearly twice as high (27.1% versus 14.1%).

Deadheading to solve inequality in supply and demand is responsible for 27.6% of all vehicle kilometers in a scenario in which demand is strongly directed, compared to 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 high, the total vehicle distance can be longer than the effective transportation distance (Figure 4d).

FIGURE 4 Effect of directionality in demand on (a) acceptance rate and average vehicle occupancy, (b) average passenger delay, (c) location-based average waiting time, and (d) total vehicle movement distance and effective transportation distance; the difference indicating distance savings

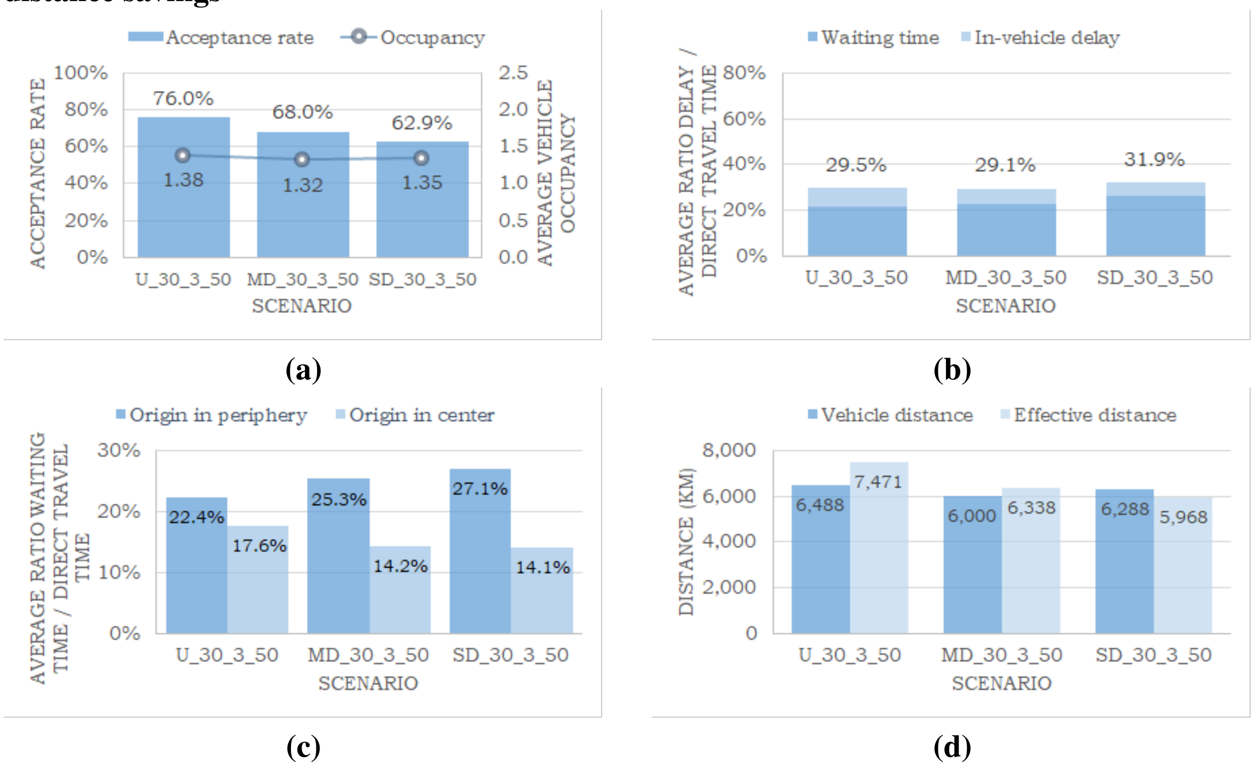

Computational complexity

With run times of approximately five hours using a single-core 2.30GHz processor, the scenarios with 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, hence increasing the solution space. 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 consequently the computational complexity of this scenario is also relatively high compared to most other scenarios (i.e. a run time of nearly one hour with the same processor).

CONCLUSIONS

This work is the first study to consider ride-sharing potential while accounting for the trade-off that users are faced with when presented with the option of ride-sharing. Previous studies, such as Santi et al. (*10*) and Alonso-Mora et al. (*11*), 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 discounted 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 sharing a vehicle with strangers induces a disutility, which in literature has been found to be one of the main barriers for a successful implementation of ride-sharing services.

Our results show that both users' tolerance to delays and willingness to share a vehicle with co-riders can have a large impact on ride-sharing potential. Depending on the tolerance of users towards sharing their ride and experiencing delays caused by detours, the acceptance rate of a ride-sharing service varied between 25.4% and 98.8%, the average delay between 15.2% and 61.3% of the direct travel time, and the gross effective vehicle transportation distance ratio between 0.95 and 1.36.

Furthermore, this study has shown that the design of a ride-sharing service, such as its pricing structure, can potentially significantly affect the expected societal benefits and service quality. A relatively 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. At the same time, the percentage of rejected requests drops from 24.0% to 17.9% if such a discount policy is implemented. In return for a discount, users are on average burdened with an extra travel time of 24.5% of the direct travel time. Hence, the pricing structure of alternative scenarios can have substantial consequences for both service performance and related externalities.

This study also shows that the potential of a ride-sharing system can be greatly dependent on external variables, such as the spatial 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 directed towards the center of the grid, typical for a morning peak, the performance of a ride-sharing system is relatively poor, both in terms of the level of service and efficiency. In such a case, only 62.9% of all requests can be served, compared to 76.0% when demand is perfectly uniform. The (gross) effective vehicle transportation distance ratio can even drop below 1 when the directionality in demand is high. Level of service of ride-sharing users is then found to be low too, following from long waiting times before pick-up and consequently a relatively large average total delay of 31.9% of the direct travel time. To summarize, we found that directionality in demand negatively affects both level of service and operational efficiency of ride-sharing services.

When representing the choice whether to ride-share or not as a compensatory function between travel attributes (travel time, ride fare and number of co-riders), the potential for reducing the total vehicle mileage is found to be relatively limited. At most 27% of the vehicle kilometers in the network can be removed, which is attained when users have a relatively high willingness to share (i.e. they are willing to pay no more than 1 euro to upgrade a shared ride to an individual one, assuming no change in travel time). In a few scenarios in this study, a ride-sharing system was even found to result in more vehicle kilometers than an equivalent system offering only individual rides would. The above findings suggest that the efficiency benefits of ride-sharing services have been overestimated in previous research. For comparison, a study on ride-sharing in New York by Santi et al. (*10*) found a 40% reduction in total vehicle mileage.

There are other potentially relevant attributes of which the effect can be investigated, besides the ones considered in this work. Future research can focus for example on the effect of fleet properties (capacity and fleet size), the effect of the fares of alternative (single-rider) services, test a more complex discounting mechanism than the one assumed here, or investigate external variables like the number of stop locations in the network. 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 presence or absence of a driver. Also, it might be interesting to find whether ride-sharing efficiency 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. Moreover, the validity of ride-sharing studies can be improved by the incorporation of mode choice, whereby passengers can choose to opt for the ride-sharing service or travel using other means. 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.

ACKNOWLEDGEMENTS

This research was supported by the CriticalMaaS project (grant number 804469), which is financed by the European Research Council and Amsterdam Institute of Advanced Metropolitan Solutions.

CONTRIBUTION

The authors confirm contribution to the paper as follows: study conception and design; de Ruijter, Cats; analysis and interpretation of results: de Ruijter, Cats; draft manuscript preparation: de Ruijter, Cats. All authors reviewed the results and approved the final version of the manuscript.

REFERENCES

1. Schneider, T. *NYC Taxi & Ride-hailing Stats Dashboard*, 2019. https://toddwschneider.com/dashboards/nyc-taxi-ridehailing-uber-lyft-data/. Accessed July 12, 2019.
2. Wang, X., M. Dessouky, and F. Ordonez. A Pickup and Delivery Problem for Ridesharing Considering Congestion. *Transportation Letters*, 2016. 8: 259–269.
3. Cici, B., A. Markopoulou, E. Frías-Martínez, and N. Laoutaris. Quantifying the Potential of Ride-Sharing Using Call Description Records. *Proceedings of the 14th Workshop on Mobile Computing Systems and Applications*, 2013.
4. Dueker, K. J., B. O. Bair, and I. P. Levin. Ride sharing: Psychological factors. *Transportation Engineering Journal of the American Society of Civil Engineers*, 1977. 103: 685–692.
5. Teal, R. Carpooling: Who, How and Why. *Transportation Research Part A: General*, 1987. 21: 203–214.
6. Morales Sarriera, J., G. Escovar Álvarez, 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*, 2017. 2605: 109–117.
7. 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*, 2011. 45: 81–90.
8. Big Data Innovation Team, *The Transportation Impacts of Vehicle-for-Hire in the City of Toronto*. City of Toronto, 2019.
9. Ma, S., Y. Zheng, and O. Wolfson. T-share: A Large-Scale Dynamic Taxi Ridesharing Service. *IEEE 29th International Conference on Data Engineering (ICDE)*, 2013. 410–421.
10. Santi, P., G. Resta, M. Szell, S. Sobolevsky, S. H. 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*, 2014. 111: 13290–13294.
11. 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*, 2017. 114: 462–467.
12. Tachet, R., O. Sagarra, P. Santi, G. Resta, M. Szell, S. H. Strogatz, and C. Ratti. Scaling Law of Urban Ride Sharing. *Scientific Reports*, 2017.
13. Fiedler, D., M. Čertickỳ, J. Alonso-Mora, and M. Čáp. The Impact of Ridesharing in Mobility-on-Demand Systems: Simulation Case Study in Prague. *IEEE International Conference on Intelligent Transportation Systems*, 2018.
14. Winter, K., O. Cats, G. Correia, and B. van Arem. Performance analysis and fleet requirements of automated demand-responsive transport systems as an urban public transport service. *International Journal of Transportation Science and Technology*, 2018. 7: 151-167.
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*, 2001. 37: 191–229.
16. Edvardsson, B. Causes of customer dissatisfaction-studies of public transport by the critical-incident method. *Managing Service Quality: An International Journal*, 1998. 8: 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*, 2003. 37: 499–517.
18. König, A. and K. W. Axhausen. The Reliability of the Transportation System and its Influence on the Choice Behaviour. *2nd Swiss Transport Research Conference*, 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*, 2001. 8: 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. https://www.kfzteile24.de/best-and-worst-cities-to-drive-usd. Accessed Jan. 23rd, 2019.
22. Liu, T., P. Krishnakumari, and O. Cats. Exploring Demand Patterns of a Ride-Sourcing Service using Spatial and Temporal Clustering. *The 6th IEEE International Conference on Models and Technologies for Intelligent Transportation Systems*, 2019. In print.
23. Erlander, S. and N. F. Stewart, *The Gravity Model in Transportation Analysis: Theory and Extensions*. VSP, Utrecht, 1990.
24. Gemeente Amsterdam. *Wat kost een rit in een taxi?*, 2019. https://www.amsterdam.nl/veelgevraagd/?caseid=%7B7DC3ADE8-163A-4201-ABC4-4ACC1830EB88%7D. Accessed Jan. 23rd, 2019.
25. Asensio, J. and A. Matas. Commuters' valuation of travel time variability. *Transportation Research Part E: Logistics and Transportation Review*, 2008. 44: 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*, 2010. 46: 384–403.