



ELSEVIER

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Transportation Research Part A

journal homepage: www.elsevier.com/locate/tra

The implications of drivers' ride acceptance decisions on the operations of ride-sourcing platforms

Peyman Ashkrof^{a,*}, Farnoud Ghasemi^{b,1}, Rafał Kucharski^b, Gonçalo Homem de Almeida Correia^c, Oded Cats^c, Bart van Arem^c

^a Department of Transport and Planning, Delft University of Technology, Stevinweg 1, Delft, Netherlands

^b Faculty of Mathematics and Computer Science, Jagiellonian University, Golebia 24, Krakow, Poland

^c Department of Transport and Planning, Delft University of Technology, Delft, Netherlands

ARTICLE INFO

Keywords:

Ride-sourcing
Ride-hailing
Agent-based simulation
Driver's behaviour
Surge pricing
Ride acceptance

ABSTRACT

As a two-sided digital platform, ride-sourcing has disruptively penetrated the mobility market. Ride-sourcing companies provide door-to-door transport services by connecting passengers with independent service suppliers labelled as "driver-partners". Once a passenger submits a ride request, the platform attempts to match the request with a nearby available driver. Drivers have the freedom to accept or decline ride requests. The consequences of this decision, which is made at the operation level, have remained largely unknown in the literature. Using agent-based simulation modelling on the realistic case study of the city of Amsterdam, the Netherlands, we study the impacts of drivers' ride acceptance behaviour, estimated from unique empirical data, on the ride-sourcing system where the platform applies regular and surge pricing strategies, and riders may revoke their requests and reject the received offers. Furthermore, we delve into the implications of various supply-demand intensities, a centralised fleet (i.e., mandatory acceptance on each ride request) versus a decentralised fleet (i.e., ride acceptance decision by each driver), ride acceptance rates, and surge pricing settings. We find that the ride acceptance decision of ride-sourcing drivers has far-reaching consequences for system performance in terms of passengers' waiting time, driver's revenue, operating costs, and profit, all of which are highly dependent on the ratio between demand and supply. As the system undergoes a transition from undersupplied (i.e., real-time demand locally exceeds available drivers) to balanced and then oversupplied state (i.e., more available drivers than real-time demand), ride acceptance decisions result in higher income inequality. A high acceptance rate among drivers may lead to more rides, but it does not necessarily increase their profit. Surge pricing is found to be asymmetrically in favour of all the parties despite adverse effects on the demand side due to higher trip fare. This study offers insights into both the aggregated and disaggregated levels of ride-sourcing system operations and outlines a series of transport policy and practice implications in cities that offer such ride-sourcing systems.

* Corresponding author.

E-mail addresses: p.ashkrof@tudelft.nl (P. Ashkrof), farnoud.ghasemi@doctoral.uj.edu.pl (F. Ghasemi), rafal.kucharski@uj.edu.pl (R. Kucharski), g.correia@tudelft.nl (G. Homem de Almeida Correia), o.cats@tudelft.nl (O. Cats).

¹ Peyman Ashkrof and Farnoud Ghasemi contributed as co-first authors.

<https://doi.org/10.1016/j.tra.2024.104362>

Available online 4 January 2025

0965-8564/© 2024 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Ride-sourcing platforms are built upon a two-sided business model where drivers, as independent contractors/freelancers, supply passengers, who submit ride requests through dedicated smartphone applications. In such a context, drivers are free to make various decisions from strategic to operational levels (Ashkrof et al. 2020). At the operational level, drivers confront a pivotal choice when a ride request appears, relying on limited information provided by the platform. Notably, in many operational areas, drivers typically lack visibility into the trip fare and final destination before accepting a ride request. Consequently, they base their choice on restricted trip-related information and make ride acceptance decisions which may have profound consequences on system operations. For instance, once a request is rejected by a driver, it is returned to the queue to be matched with a new driver. This leads to longer waiting time for passengers and therefore a lower level of service and higher idle time for that driver.

Ride-sourcing platforms implement varying degrees of limitations regarding the disclosure of trip-related information to drivers. For instance, in the case of Uber in Amsterdam, drivers are provided with the pick-up point address, the distance, the estimated travel time from their location, and the rider's rating. However, the trip fare and final destination are not disclosed initially. Drivers only see the destination once they have accepted the request and are approaching the rider. Uber cites passenger privacy protection as the primary reason for limiting trip-related information (Uber, 2022c). Nevertheless, this practice may also be intended to reduce the likelihood of drivers cancelling requests after acceptance, as the driver has already committed some distance to pick up the passenger (Ashkrof et al. 2020).

In the supply-side literature, three main research streams are identified. The first research stream assumes that drivers are fully compliant with platform strategies (Wang, Fu, and Ye 2018; Cachon, Daniels, and Lobel 2017; Zha, Yin, and Yang 2016). Under this assumption, drivers make no independent decisions and completely follow what they are instructed to do by the platform app. The second stream focuses on an automated fleet centrally operated by the platform (Ruch et al. 2018, Winter et al. 2016, Zhang et al. 2016, Levin 2017, Liang et al. 2018, Wang et al. 2022). In both cases, the assumptions made on fleet characteristics appear not to be in line with current operations and simply ignore the impacts of driver's behaviour on system performance.

A third, recently emerging research stream is devoted to the study of labour behaviour and preferences of the ride-sourcing supply side (Ashkrof et al. 2020; Ramezani et al. 2022; Ashkrof, Homem de Almeida Correia, et al. 2023). Applying a data-driven method, Xu et al. (2018) mention that nearly 40 % of ride requests in the Didi Chuxing platform on 23 January 2017 in Shanghai, China, received no response. They also report that surge pricing significantly increases the probability of accepting ride requests. Analysing 576 responses collected through a stated choice survey from Uber and Lyft drivers in the US, Ashkrof et al. (2022) estimate the effects of various factors influencing ride acceptance behaviour. They suggest that part-time and beginning drivers tend to have higher acceptance rates. Based on choice model estimations, they conclude surge pricing has a higher marginal utility than trip fare per monetary unit.

On the platform side, a large body of research is focused on the design of matching and pricing strategies (Cachon, Daniels, and Lobel 2017; Zha, Yin, and Du 2018; Zha, Yin, and Xu 2018; Li, Jiang, and Lo 2022; Son 2023). One of the focal research topics in this area is surge pricing. Surge pricing, a commonly applied strategy in ride-sourcing practices, offers a monetary bonus for drivers when the demand is excessively higher than the supply. It is often considered a black-box algorithm (Chen, Mislove, and Wilson 2015) due to the lack of transparency and the limited information shared by ride-sourcing platforms. In one of the earliest studies in this area, Chen et al. (2015) reverse-engineered surge pricing by tracking the Uber application operations on 43 mobile phones distributed through San Francisco and Manhattan. They found that surge pricing has a significant negative effect on riders and a minor positive impact on drivers. Guda and Subramaniana (2019) argue that surge pricing should be applied strategically. They suggest that if there is a demand surge in a given zone, the platform should apply strategic surge pricing in a proactive manner.

Most of the abovementioned studies address a specific problem, while the interaction and dynamics between the involved parties have remained largely neglected. In order to understand and analyse the complex nature of ride-sourcing operations, the behaviour and interactions of both drivers and riders with the platform need to be considered simultaneously. Bokányi and Hannák (2020) performed agent-based simulation experiments to model a ride-hailing system at the microscopic level. They found that income inequality among drivers is higher in an oversupplied system (i.e., the number of drivers is excessively higher than what would be needed to satisfy the incoming requests) and the extent of which depends on the spatial characteristics of requests, drivers' repositioning strategies, and the platform matching algorithm. Developing a dynamic framework that takes into account the daily participation of drivers in the platform operations, de Ruijter et al. (2022) model the emergence of substantial income disparities amongst drivers.

Given that the information on the supply side is strictly limited due to the reluctance of ride-sourcing companies to share revealed preference data and the fact that stated preference data collection is costly owing to the highly specific target group, drivers' behaviour, particularly at the operational level, is typically neglected or underestimated in simulation studies. Some studies included driver's relocation choice in the analysis of two-sided ride-sourcing platforms by directly adopting past results for taxi drivers' repositioning choices (e.g. Nahmias-Biran et al., 2019). Nonetheless, the other operational decision – ride acceptance behaviour – is rather unique to the ride-sourcing context, and there is a lack of knowledge of its implications.

The identified gap calls for developing a simulation model that incorporates the stochastic dynamic ride acceptance decision of ride-sourcing drivers into system operations while accounting for the decisions made simultaneously by other parties. This modelling approach enables devising new strategies to improve system operational management and analysing the performance consequences thereof. To the best of our knowledge, this is the first study to model an empirically underpinned representation of ride-sourcing drivers' acceptance behaviour into a simulation model. One notable strength of this research lies in deriving the acceptance function from an empirical study conducted by Ashkrof et al. (2022), specifically tailored to discern the variables influencing drivers' ride acceptance decisions. Leveraging the insights gained from this study provides substantial value, especially within the context of

current ride-sourcing system operations that feature restricted information sharing with drivers. This mirrors the existing system operations where drivers make decisions without access to upcoming trip fare and distance details. We also take into account riders' decisions and platform matching and pricing strategies. The main research contributions of this paper are listed as follows:

- Developing a discrete-event agent-based simulation framework for realistically reproducing ride acceptance behaviour of ride-sourcing drivers in a complex dynamic system.
- Extending the framework with characteristics of a centralized (automated) fleet to analyse its implications in comparison to a decentralized fleet.
- Designing a surge pricing model based on actual platform operation and seamlessly incorporating it into the proposed agent-based framework.
- Insights and implications from applying the model to the realistic case study of the city of Amsterdam, the Netherlands

The analysis consists of several layers: First, we contrast the performance of a decentralised fleet where individual drivers make ride acceptance decisions, and a centralised fleet, where drivers are fully compliant with the system in a way that all the requests are accepted (automated fleet) and also analyse the effects of supply–demand intensity on the system-level behaviour. Second, the implications of drivers' ride acceptance rate are analysed. Third, the impacts of surge pricing on system operations are discussed. In the next section, the methodology applied in this research is elaborated.

2. Methodology

2.1. Agent-based simulation model

In this research, we adopt a bottom-up approach for modelling emergent phenomena by means of developing and using an agent-based simulation model in which drivers, riders, and the platform are defined as the model agents. Given the immense complexity arising from the real-time interactions between the ride-sourcing actors, an agent-based simulation model allows capturing the ride-sourcing system dynamics, offers flexibility in adding the system agents with their behaviour, and provides the ability to change levels of detail as well as aggregation and simulate the complex relationship between riders, drivers, and the platform. This modelling approach has been used extensively to study demand-responsive transport modes in urban areas (Martínez et al., 2017).

To study the implications of ride-sourcing drivers' ride acceptance decisions on system performance, we adopt and adapt a within-day discrete event agent-based simulation model in which drivers are individual decision-makers on ride requests, interacting with riders and the platform and integrated into the MaaSSim simulator (Kucharski and Cats, 2022). For the needs of this study, the simulation framework is extended by introducing the drivers' ride acceptance and surge pricing functionalities, as detailed below.

2.1.1. MaaSSim

MaaSSim (Kucharski and Cats, 2022) is an open-source, agent-based simulator in Python, which reproduces the dynamics of ride-sourcing systems. It simulates, on the urban road network, the behaviour and interactions of two classes of agents: (i) travellers, requesting a ride to travel from their origin to their destination at a given time, (ii) drivers, offering rides for travellers to supply their travel needs. The spatiotemporal interactions between the two types of agents are mediated by the platform, matching travellers to the closest idle driver. Both demand and supply, as well as their interactions, are modelled at a disaggregate level. This relates to the explicit representation of single vehicle agents and their movements in time and space for supply. For demand, this pertains to the precise trip request time and destinations defined at the graph node level. MaaSSim does not account for the traffic congestion, assuming a constant speed for ride-sourcing vehicles and that their contribution to traffic is not enough to impact travel times, which are predominantly influenced by private cars.

2.1.2. Matching and ride acceptance

In MaaSSim, a rider agent submits a ride request (r) with a specific origin and destination. Next, the platform agent attempts to match the requests queued in the system (Q_r) with nearby available drivers (d) in the queue (Q_d) based on their distances. Considering our constant speed assumption, nearby drivers can be alternatively defined by pickup times (i.e., travel time between the driver's location and the rider's location (P_{rd})). Thus, by matching passengers with the closest drivers, the possible shortest pick-up times are ensured. The match between rider and driver (M_{rd}) is calculated based on the matching function presented in Eq. (1). Note that the MaaSSim matching function is reactivated upon any changes in either the requests' queue or the drivers' queue, i.e., when a new ride request is made or a driver becomes idle.

$$M_{rd} = \operatorname{argmin}_{P_{rd}, r \in Q_r, d \in Q_d} \quad (1)$$

On the supply side, once a driver agent receives a request, they accept or decline it based on a logit model using a utility function. The model specification is based on the model estimation results reported in the study by Ashkrof et al. (2022b), specifically designed to analyse the ride acceptance behaviour of ride-sourcing drivers. Ride-sourcing platforms adopt various information-sharing policies leading to partial disclosure of information about ride requests. In most cases, ride-sourcing drivers make ride-acceptance decisions based on limited information. For instance, trip fare and final destination are typically not shown to drivers before ride acceptance. Therefore, we only include the typical factors related to the current system operations. To this end, the utility function consists of the

driver's idle time (I_d), which is the duration between the last drop-off and the incoming request, pickup time (P_{rd}), alternative-specific constant (ASC), and the error term (ε) representing the model stochasticity in the base scenario in which the platform applies regular pricing. Eq. (2) and Eq. (3) represent the acceptance utility function specification and the corresponding probability according to a binary logit function since the driver is not able to see and assess several requests at the same time in the same platform.

$$U_{acceptance} = ASC + \beta_{I_d} * I_d + \beta_{P_{rd}} * P_{rd} + \varepsilon \quad (2)$$

$$P_{acceptance} = \frac{e^{U_{acceptance}}}{1 + e^{U_{acceptance}}} \quad (3)$$

If a driver agent accepts a request, the following events take place in the simulation: heading to the passenger's location, picking them up, driving to the destination, dropping them off, staying at the drop-off point, and waiting for the next ride (it is assumed that drivers stay at the drop-off location and make no repositioning decision) or ending the working shift. If the request is rejected, driver and rider agents return to the queue and wait for a possible next match.

2.1.3. Driver's profit

Once a ride is completed (i.e., the rider arrives at the destination), the corresponding driver and the platform receive their share of the trip fare based on the platform pricing strategy. When the platform uses regular pricing, the driver's revenue (R_d) is calculated based on the base fare (f_{base}), fare per kilometre (f_{km}), trip distance in kilometres (td_{km}), and the platform commission fee (α). The operating costs (OC_d) depend on the total driving distance (pickup distance plus trip distance) in kilometres (dd_{km}) and the total costs per kilometre (tc_{km}). These expenses consist of gas, tax, insurance, repair, maintenance, depreciation, or lease payments. The subtraction of operating costs from revenue yields the driver's profit (P_d). Eq. (4) – Eq. (6) present the driver's revenue, operating costs, and profit, respectively.

$$R_d = (f_{base} + f_{km} * td_{km}) * (1 - \alpha) \quad (4)$$

$$OC_d = tc_{km} * dd_{km} \quad (5)$$

$$P_d = R_d - OC_d \quad (6)$$

2.1.4. Surge pricing

Once riders locally outnumber drivers, a bunch of ride requests may be left unserved, potentially harming the platform's reputation. Hence, the platform strives to balance supply and demand using multiple approaches such as surge pricing. In this study, the platform agent applies surge pricing (also known as dynamic pricing) to adjust trip fares based on the real-time ratio between demand and supply. In case surge pricing is applicable, an additional positive term associated with surge pricing (S) is added to the ride acceptance utility function (Eq. (7)).

$$U_{acceptance} = ASC + \beta_{I_d} * I_d + \beta_{P_{rd}} * P_{rd} + \beta_S * S + \varepsilon \quad (7)$$

For surge pricing, for example, Uber employs a geospatial indexing system dividing cities into multiple zones where the real-time ratio between demand and supply is calculated (Chen et al., 2015). We mimic Uber's strategy (Uber surge pricing patent, 2013) for implementing surge pricing by dividing the study area into several hexagonal zones using an open-access python package known as H3 developed by Uber (Uber, 2022a). In each zone, we record the real-time ratio between demand and supply. If this ratio is greater than one, then this implies that the system is undersupplied, and surge pricing should be applied.

We consider a surge multiplier ranging from 1x to 5x with an interval of 0.1, which is in line with Uber operations (Chen et al.,

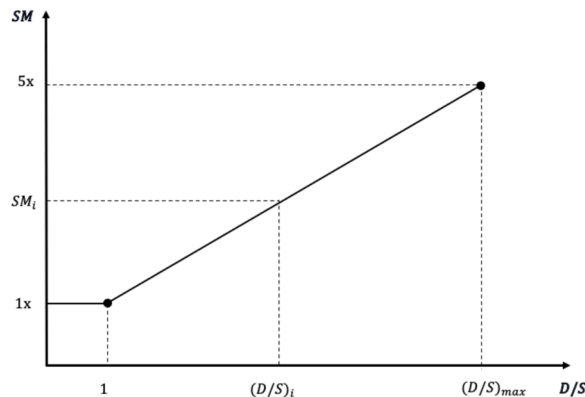


Fig. 1. Surge Multiplier Graph.

2015). The multiplier is determined based on the following surge function that depends on the real-time ratio between demand and supply in zone i ($(D/S)_i$). As shown in Fig. 1, the function is bounded by: (i) a minimum demand–supply ratio of 1 (i.e., perfect local balance between supply and demand) and an absolute maximum ratio between demand and supply (D/S_{max}) in all zones, on the x-axis; (ii) a minimum surge multiplier of 1x (i.e., regular price) and a maximum surge multiplier of 5x, on the y-axis. We specify a linear function for estimating the surge multiplier based on the real-time ratio between supply and demand in any given zone. Eq. (8) mathematically presents the surge function yielding the surge multiplier in zone i (SM_i).

$$SM_i = \begin{cases} 1 + \frac{4 * \left(\left(\frac{D}{S} \right)_i - 1 \right)}{\frac{D}{S_{max}} - 1}, & \left(\frac{D}{S} \right)_i > 1 \\ 1, & \left(\frac{D}{S} \right)_i \leq 1 \end{cases} \quad (8)$$

Once the surge multiplier is determined, the trip fare is updated in the system, and the driver's revenue is calculated based on Eq. (9):

$$R_d = (f_{base} + f_{km} * td_{km}) * SM_i * (1 - \alpha) \quad (9)$$

As mentioned earlier, the surge pricing strategy aims at restoring the balance between demand and supply. This is achieved by retaining the ride requests for which the willingness to pay is the highest. To account for the impacts of surge pricing on the demand side, we adopt the demand elasticity estimated by Cohen et al. (2016), who conducted a regression discontinuity analysis. We obtain the probability of an offer being accepted by riders based on the applied surge multiplier (Fig. 2).

2.2. Experimental design

We apply the developed simulation framework to the city of Amsterdam, the Netherlands, where merely non-shared (UberX) services are available. According to the municipality of Amsterdam (Gemeente Amsterdam, 2019), around 8 million taxi/ride-sourcing trips were made in 2019 in Amsterdam, resulting in an average of 22,000 rides per day. In order to study the within-day ride-sourcing system operations, we set the simulation starting and ending times at 8:00 and 16:00, respectively, to cover a typical working shift of 8 h. Given the non-uniform distribution of demand within a day, we assume 10,000 ride-sourcing rides out of 22,000 daily rides are made during the simulation time. The demand is sampled from a dataset of trips of 3 km or longer obtained from an updated version of the activity-based Albatross model (Arentze and Timmermans, 2004). Given that the matching algorithm of Uber reportedly takes seconds to run (Medium, 2022; Uber, 2022), we assume that riders revoke their requests without making new request once no offer is given within 3 min. Note that MaaSsim does not account for the traffic congestion, thereby assuming a constant speed for ride-sourcing vehicles and that their contribution to traffic is not enough to impact travel times, which are predominantly influenced by private cars.

We perform experiments for fleet sizes varying between 100 and 1000 drivers with an increment step of 100, enabling us to analyse the effects of the supply–demand intensity. Driver agents are geographically generated based on a negative exponential distribution function to ensure more drivers are placed in the city centre where more demand is expected. Drivers' working shift matches the simulation time, i.e., no working shift decision is made. Vehicle speed is assumed to be constant at 36 km/h. We consider two types of fleets: a) a decentralised fleet in which drivers make ride acceptance decisions; and b) a centralised fleet where drivers accept all the requests (fully compliant). Then, the performance of both fleets is compared to provide deeper insights into the implications of drivers' ride acceptance decisions. We assume that only one platform is operating in the city, therefore multi-homing (i.e., driving for multiple platforms) is not applicable.

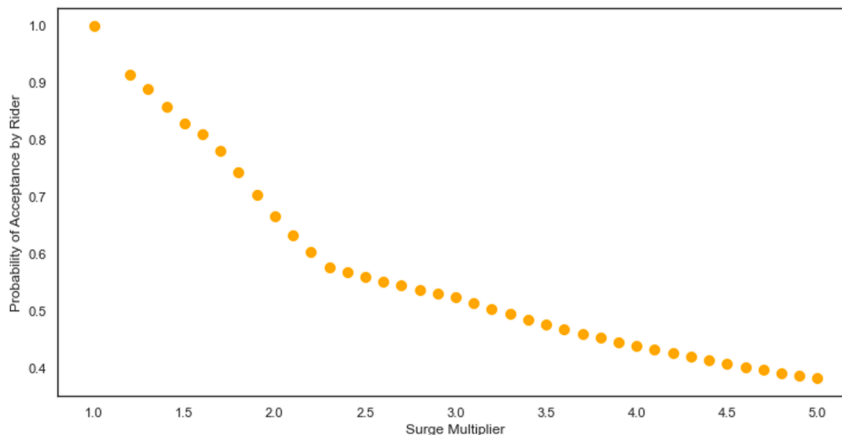


Fig. 2. Demand elasticity as a function of the surge multiplier (Cohen et al. 2016).

Empirical data from the real-world Albatross trip set, which we are using in our study, demonstrates higher demand in the city center (Arentze and Timmermans, 2004) which is in line with the ride-sourcing demand distribution in the literature (Marquet, 2020; Wang et al. 2021). Accordingly, we assume higher supply in the city center based on a negative exponential distribution function to align with expected demand patterns. Such an assumption is in line with previous studies on real-world data, which, based on GPS trajectories, indicate that drivers are distributed throughout the city with a high concentration in the city center, and their number decrease with increasing distance from the city center (Qian and Ukkusuri, 2015; Santi et al., 2015; Peng et al., 2012).

Based on the models estimated by Ashkrof et al. (2022), the parameters of the ride acceptance functions, ASC , β_{1d} , β_{pd} , and β_s , are set to 1.810, -0.017 , -0.050 , 0.101, respectively. Using the Uber price estimator tool in 2022, the base fare (f_{base}), the fare per kilometre (f_{km}), and the minimum fare (f_{min}) are set to €2, €1.2, and €5, respectively. The platform commission rate is 25 % (α), in line with the current Uber operations in Amsterdam. The operating cost of drivers (OC_d) is set to €0.5 per kilometre (Standard mileage rates (2021)). Given the stochasticity of the simulation process, several experiment replications are needed to attain statistically robust results. The number of required replication runs is set to five as it is found to obtain results with a confidence level of 95 %. We use the average values to conduct the analysis at the aggregate level.

2.3. Results

The results of the experiments are presented and analysed along four dimensions: the impacts of the supply–demand intensity on the performance of a centralised (automated) fleet versus a decentralised fleet operated by drivers making their decisions on ride requests (Section 3.1), the effects of ride acceptance rate on drivers' performance (Section 3.2) and an analysis of applying surge pricing and its consequences for system and drivers' performance (Section 3.3).

2.4. Centralised versus decentralised fleet

A widespread assumption in most of the literature concerned with ride-hailing is that a central operator makes all the decisions regarding a match between a driver and a client. This in reality implies that the fleet consists of either AVs or fully-compliant drivers who do not make independent choices, none of which reflects the current state of affairs. Rather, drivers are independent decision-makers who choose what they perceive to be in the best of their interest. In this sub-section, we analyse and compare the performance of centralised and decentralised fleets.

Fig. 3 presents the distribution of passenger waiting time with centralised (CF) and decentralised fleets (DF) using a violin plot. It can be seen that passenger waiting time in the DF, where drivers can reject ride requests, is not necessarily higher than in the CF. It is observed that there is no substantial difference between the waiting time of passengers in both fleets in case the system is under-supplied with 300 drivers or fewer. This is because the number of drivers in this state is not sufficient to serve all requests. Therefore, even if a request is rejected in the DF, another request (not far away from the other) is immediately offered to the driver which, if accepted, results in a shorter waiting time for another passenger. While the rejected passenger probably loses their patience and leaves the system. Once the system is in a balanced/oversupplied situation (400 drivers and more), the average and standard deviation of passenger waiting time are higher in the DF than the CF, i.e., drivers' acceptance behaviour results in system efficiency loss. This is because some requests get rejected in the DF by the closest available driver, and then a farther away driver may get assigned to the request which in turn might get rejected again.

Across fleet, the general pattern indicates that the average passenger waiting time does not necessarily decrease for larger fleet sizes which might be counter-intuitive. For instance, the average waiting times are 190, 217, and 303 s for fleet sizes of 100, 200, and 300, respectively. This can be explained by the trade-off between match rate (i.e., the percentage of rides matched within a specific time interval) and the match quality (i.e., the attractiveness of a ride such as pickup distance which is the distance between the driver's location and the pickup point). When the fleet size is 100, 6,224 out of 10,000 requests are rejected which results in a total match rate of 38 % while the average pickup distance is 1.32 km. When 100 more drivers are added to the system – resulting in a fleet size of 200 – the match rate rises to 71 %, whereas the match quality decreases with an average pickup distance of 1.72 km leading to a higher waiting time for the passengers. The average waiting time peaks with 300 drivers, in which the total match rate reaches 92 % and the pickup distance levels up to 2.81 km. With a fleet size of 400, the system seems to reach a balance between supply and demand given that all the requests during the simulation period are served (i.e., 100 % match rate). Once the number of drivers exceeds this

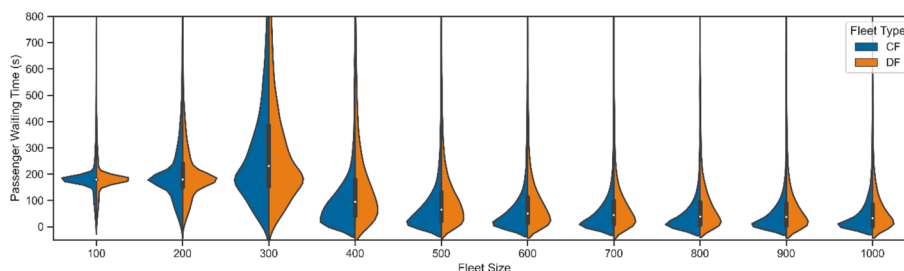


Fig. 3. Passenger waiting time in the centralised and decentralised platforms.

threshold, a sharp drop in the passengers' average waiting time is observed. From this point onward, the waiting time pattern changes and the average value constantly decreases with the growing fleet size. This is because the number of drivers is (more than) sufficient to cover all requests with shorter pickup distances.

As the system state changes from balanced to oversupplied (around 500 drivers to 800 driver), the likelihood that drivers are available in the vicinity of requests increases leading to a decrease in the average passengers' waiting time. From 800 drivers onward, average passengers' waiting time remains roughly constant. The threshold of 800 drivers emerges from experiments with different fleet sizes in the context of Amsterdam, a city with a population of approximately 920,000 and a total of around 5000 taxi drivers, including Uber drivers (ROA Rapport, 2023; Gemeente Amsterdam, 2019). Beyond this point, we observed that passenger waiting times plateaued, indicating no significant improvement in system performance. The threshold is mainly determined by demand level, trip distances, and pick-up distances. In Amsterdam, ride-sourcing demand is relatively low (around 2 %), likely due to the high share of trips completed by bikes. Additionally, the city has relatively small and dense road network leading to shorter trip and pick-up distances. While short trips increase the drivers availability resulting in low matching times, short pick-up distances means short pick-up times, ultimately decreasing passenger waiting times. In larger cities with higher demand and/or longer travel distances (e.g. due to urban sprawl), thresholds are expected to be higher, and conversely, lower in smaller and more compact cities.

Next, the performance of the ride-sourcing supply side is analysed starting with the drivers' revenue in Fig. 4. Comparing the drivers' revenue in both fleets, we find a significantly higher income variation between the drivers of the DF in the balanced and oversupplied states. In other words, human intervention causes more income inequality as drivers seek to maximise their income regardless of the system state. That is why the income dispersion is considerably sharper in the DF where more drivers can earn a higher income, in some cases even more than the maximum revenue in the CF. On the other hand, the number of drivers who have a low revenue is higher in the oversupplied DF given the high competition. As expected, the average revenue is nearly the same in both fleet types, especially when all the requests are satisfied (equivalent total revenue). Nevertheless, the revenue is slightly higher in the CF in the undersupplied state given that more rides are served.

Considering the overall trend with respect to fleet size, the average revenue of the drivers decreases as the fleet size becomes larger for a given demand level. We do, however, observe that once the number of drivers exceeds 500, this income rate of reduction due to the joining of more drivers decreases. Another observation concerns the income variation which sharply rises once the system state changes from being undersupplied to balanced and oversupplied. This might be due to higher competition amongst drivers.

Drivers' revenue includes the costs incurred during their operations, known as operating costs. As illustrated in Fig. 5, more variation in the drivers' operating costs in the DF can be found. Such difference is more striking in the balanced and oversupplied states where the disparity in driver's ride acceptance behaviour becomes large. This is because in the DF, drivers are free to decline rides assigned based on the shortest pickup distance. Once a request is rejected, another one, which is likely to be farther away, might be matched with the driver resulting in a longer pickup distance and thereby contributing to increased operating costs. Such a disparity in drivers behaviour results in diverse drivers' revenue and operating costs for DF, i.e., their distributions on Figs. 4 and 5 tend to shift from being unimodal toward a more uniform shape. Thus, we expect a uniform distribution for highly oversupply conditions (fleet size > 1000).

Similar to the overarching revenue pattern concerning fleet size, we observe that the undersupplied system performs strikingly different from the base case. In this range ($< = 300$ drivers), there is no notable average difference in the operating costs of drivers, while a sharp drop in the operating costs occurs in the balanced and oversupplied state. The operating cost is obtained from driving distance consisting of pickup and trip distance (Fig. 6). It should be highlighted that drivers are not paid for picking up passengers. Thus, pickup distance merely incurs costs for these drivers. At the same time, the average trip distance travelled in the shift decreases with an increasing number of drivers.

It is interesting to note that the increase in the average pickup distance is offset by the fall in the average trip distance for fleet sizes of 100, 200, and 300. This results in similar driving distances and consequently nearly equivalent average operating costs for these drivers. Given that the average pickup distance starts decreasing with the rise in fleet size in the balanced state, there is a steep decline in operating costs due to the synergy between the decline in the pickup and trip distances in the balanced and oversupplied ranges. In addition to the mean, the cost variation between drivers grows for fleet sizes of 400 and higher, this is due to the more fierce competition between drivers.

Following the higher variation in revenue and operating costs of drivers in the DF, drivers' profit varies more dramatically when drivers are able to make decisions on ride requests. With larger fleet sizes, the average profit is slightly higher in the CF once the

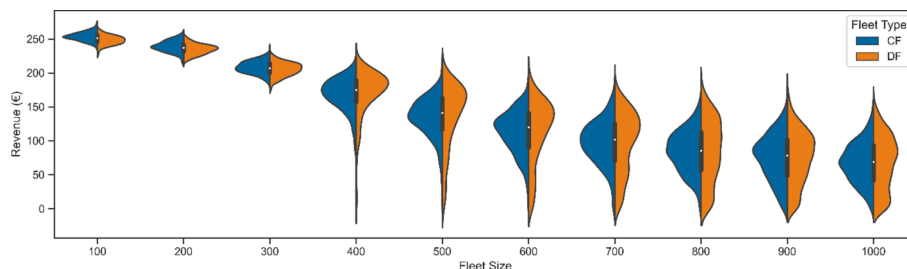


Fig. 4. Driver's revenue in the centralised and decentralised platforms.

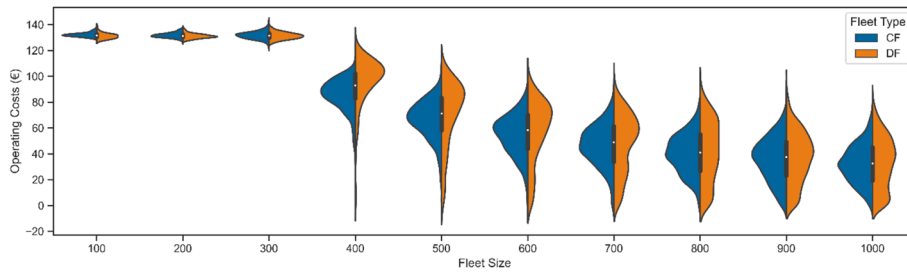


Fig. 5. Drivers’ operating costs in the centralised and decentralised platforms.

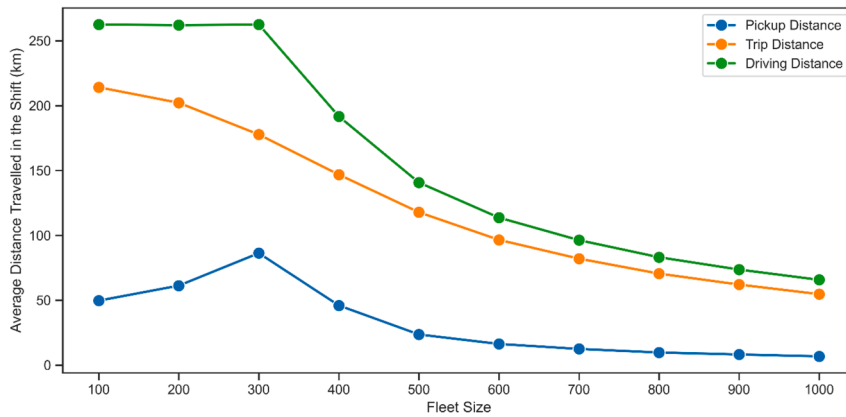


Fig. 6. Average distance travelled per driver.

number of drivers is between 100–500. As mentioned above, this is due to marginally higher revenue in the undersupplied state and/or lower operating costs in the balanced state of the CF during the shift in these scenarios.

The platform revenue during the simulation time does not significantly change as soon as all requests are served. Nonetheless, driver’s acceptance behaviour may result in a higher average waiting time for passengers in the balanced and oversupplied state. This may decrease the platform market share which may potentially in the long term lead to lower income for a platform with a DF.

| Fleet Size | Acceptance Rate | | | | |
|------------|-----------------|---------|--------|--------|--------|
| | 60 or less | 60-70 | 70-80 | 80-90 | 90-100 |
| 100 | NA | 1381.00 | 767.55 | 510.66 | 218.83 |
| 200 | NA | 1119.50 | 631.45 | 369.64 | 139.31 |
| 300 | NA | 445.43 | 319.94 | 191.27 | 76.97 |
| 400 | 422.67 | 156.64 | 115.33 | 71.94 | 32.65 |
| 500 | 228.16 | 116.00 | 93.77 | 61.37 | 23.79 |
| 600 | 197.31 | 106.03 | 82.21 | 51.41 | 15.57 |
| 700 | 176.52 | 97.50 | 73.26 | 47.85 | 15.29 |
| 800 | 151.69 | 88.76 | 68.12 | 44.06 | 13.64 |
| 900 | 161.38 | 81.41 | 62.69 | 36.87 | 11.63 |
| 1000 | 122.94 | 78.36 | 60.52 | 36.17 | 12.47 |

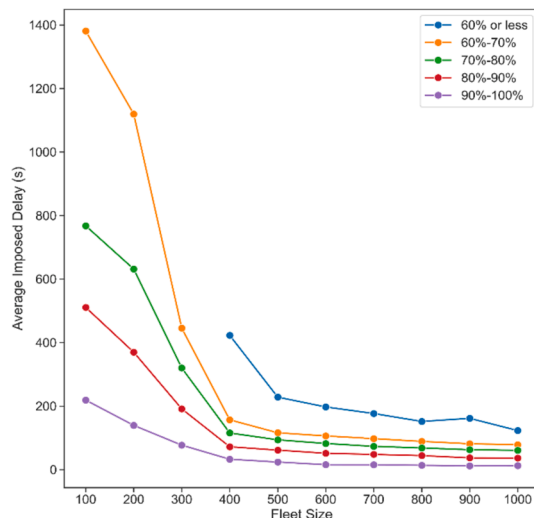


Fig. 7. Average imposed delay per driver.

2.5. Ride acceptance rate implications

In this section, we observe how the drivers' acceptance behaviour affects their performance. To this end, we track the drivers' behaviour across the simulation and classify them based on their acceptance rate, which is obtained from the ratio between the number of accepted requests and the total number of requests received during the shift. The drivers are then grouped into five segments with an acceptance rate of 60 % or less, 60 %-70 %, 70 %-80 %, 80 %-90 %, and more than 90 %. The performance of drivers is analysed for each of these groups.

First, we analyse the delay imposed on passengers by drivers rejecting rides. Fig. 7 depicts the imposed delay in the shift for different fleet sizes based on the drivers' ride acceptance rate. In the undersupplied state, none of the drivers has an acceptance rate of less than 60 % given that the number of requests received during the shift is high enough. The delay imposed by a driver is the time that the corresponding passenger whose request has been rejected needs to wait to be matched with another available driver or cancel the requests because of exceeding the maximum allowable waiting time. Expectedly, the imposed delay decreases when more drivers are available. In the undersupplied state, the delay imposed by drivers with a low acceptance rate is much higher given that the rejected passengers should wait longer owing to the low number of drivers available to serve them.

Fig. 8 presents drivers' average profit benchmarked against their acceptance rate in each fleet size. In the undersupplied state, the acceptance rate does not have a significant effect on the average profit per shift given the high number of incoming requests despite a relatively high rejection rate. Once the system is no longer in undersupply, the average profit significantly varies based on the acceptance rate. From this point onward, drivers with an acceptance rate of up to 60 % have the lowest income as the idle time is higher. With 500 or more drivers, the highest income is earned by the drivers whose acceptance rate is between 80 % and 90 %.

To gain a better understanding of the ride acceptance rate implications on drivers' profits, Fig. 9 shows the distribution and quartiles of the profit based on drivers' acceptance rate for different fleet sizes. Interestingly, the profit of drivers with an acceptance rate of 80 %-90 % is either more than or equal to the other categories in each quartile once the number of drivers is 500 or higher. In the oversupplied situation, a high-profit variation with an acceptance rate of 90 % is seen. In other words, this group of drivers can make a profit as high as or more than the drivers with an acceptance rate of 80 %-90 % or as low as or less than drivers with a 60 %-70 % acceptance rate. To explain this, we need to investigate the performance of each driver. A scatterplot of the number of rides received by each driver against the acceptance rate is shown in Fig. 9.

2.6. Surge pricing

At the operational level, surge pricing is one of the platform's most prevalent pricing strategies. This surge in price, applied as a multiplier or an additive value to the trip fare, is paid by riders and results in higher revenue for drivers and the platform. Some passengers may revoke their request due to their sensitivity to trip fare. Moreover, surge pricing motivates drivers to accept more rides given that they can earn more money through surge pricing in which each monetary unit is valued higher than the trip fare (Ashkrof et al., 2022b). To implement surge pricing, the case study area of Amsterdam is divided into 55 hexagonal zones, and the demand-supply ratio is dynamically calculated in each zone. If a zone is undersupplied, the surge function (as explained in section 2.1.3) is used to assign a surge multiplier ranging from 1.0x to 5.0x and estimate the corresponding surge fee. Fig. 10 illustrates the granular hyperlocal zones in Amsterdam and the calculated maximum demand-supply ratio in each zone using a colour palette ranging from dark red (high surge) to light red (low surge).

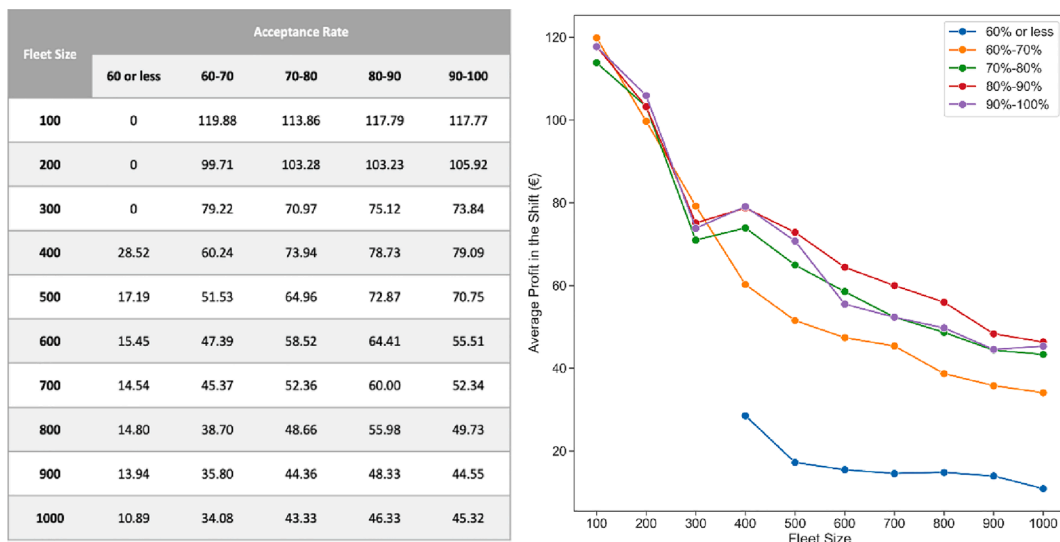


Fig. 8. Average profit in the shift based on the acceptance rate.

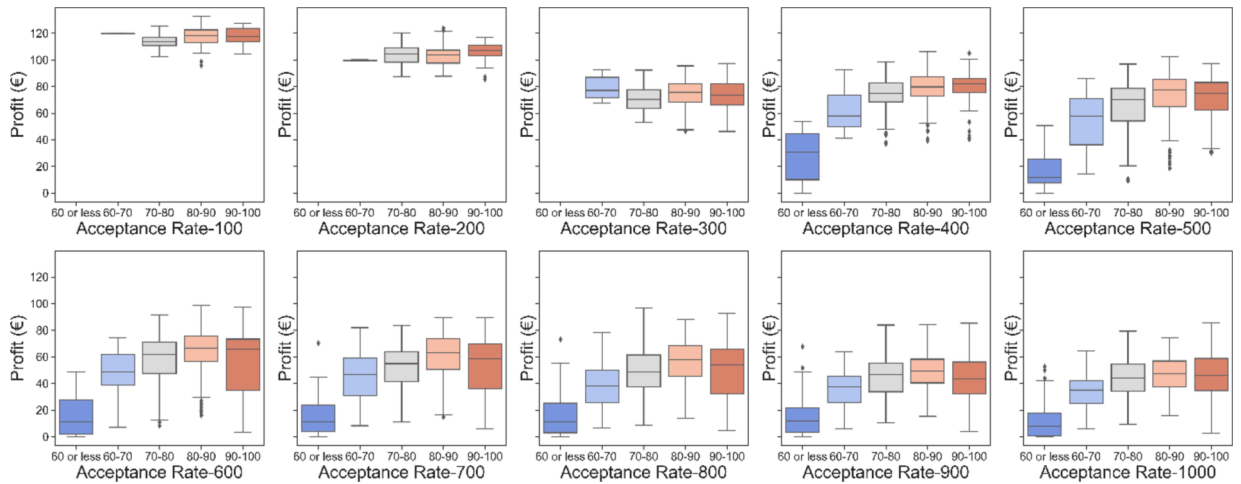


Fig. 9. Driver's profit variation for different acceptance rates.

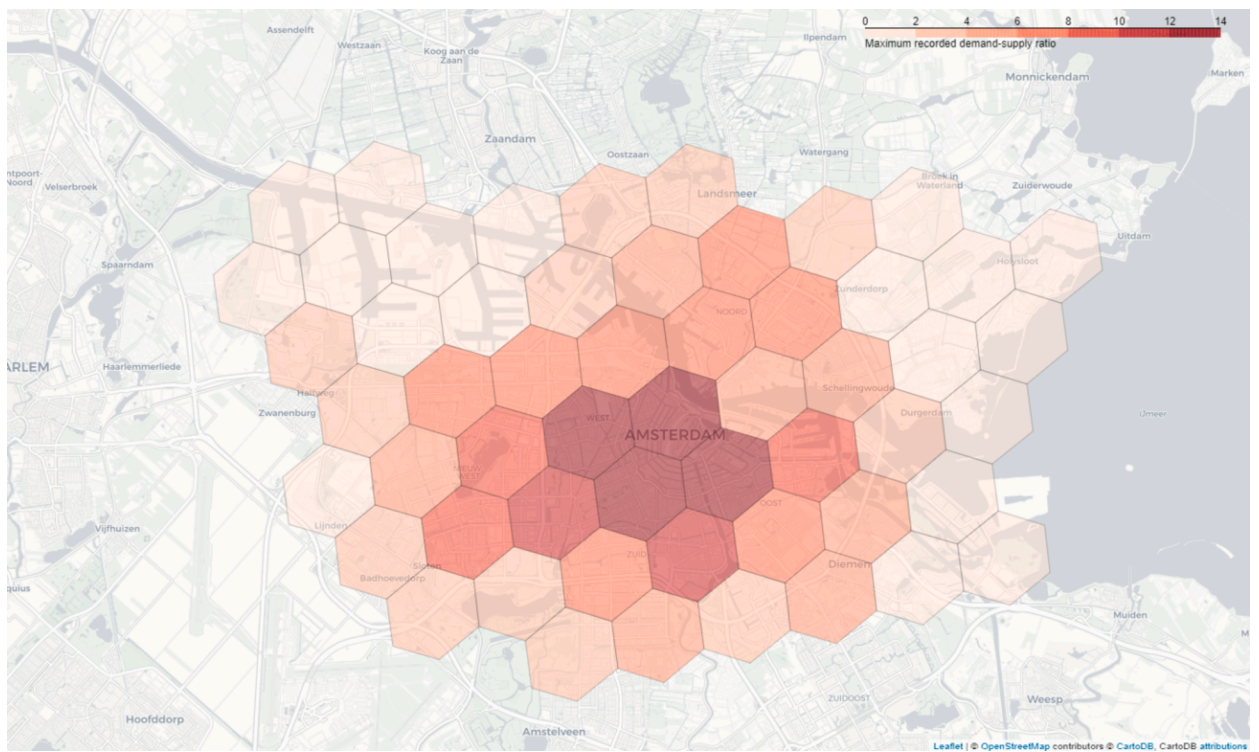


Fig. 10. Maximum demand-supply ratio in the granular hyperlocal zones in Amsterdam.

Table 1
Summary statistics of surge multiplier and drivers' surge income.

| Fleet size | Surge Multiplier | | Drivers' Surge Income in the Shift | |
|------------|------------------|--------------------|------------------------------------|--------------------|
| | Mean | Standard Deviation | Mean (Euros) | Standard Deviation |
| 100 | 1.51x | 0.08 | 119.83 | 21.71 |
| 200 | 1.34x | 0.08 | 74.77 | 18.27 |
| 300 | 1.15x | 0.06 | 29.34 | 12.33 |

Table 1 presents the mean and the standard deviation of the surge multiplier and drivers' surge income (i.e., the revenue derived from surge pricing) in various fleet sizes. Simulation results indicate that surge pricing is activated only for fleet sizes of up to 300, in which the system is undersupplied. With 69 % of the rides being surge priced in the extremely undersupplied state (i.e., 100 drivers), drivers have approximately 120 euros (on average) surge income during a working shift with an average surge multiplier of 1.51x. With 200 drivers, surge pricing is applied to 49 % of the rides and an average driver can earn 75 euros from surge pricing with a multiplier of 1.34x. In the moderately undersupplied situation where the fleet size is 300, only 21 % of the rides receive surge pricing with a multiplier of 1.15x resulting in an average surge income of 30 euros for drivers. Remarkably, the average surge multiplier and the number of rides with dynamic pricing in the moderately undersupplied fleet size scenario (300 drivers) perfectly match the corresponding indicators in real-world Uber system operations. Moreover, surge pricing decreases the demand by no more than 4 % in all the fleet size scenarios as a consequence of the relatively inelastic passenger demand towards surge pricing reported by Cohen et al. (2016).

The analysis of drivers' income shows that the average profit in the shift increases with surge pricing by about 100 %, 70 %, and 30 % with 100, 200, and 300 drivers, respectively. Given that no significant change in the operating cost is observed, this rise in profit stems from the higher revenue obtained from drivers' surge income. In addition, a higher variation in driver's profit is observed, implying that surge pricing reinforces income inequality.

Given that the platform charges a 25 % commission per ride, the platform's revenue is higher when surge pricing is introduced. The platform can earn up to 50 % higher revenue in a shift with surge pricing when the fleet size is 100. With 200 and 300 drivers in the surge pricing scenario, the platform revenue is 30 % and 10 % higher, respectively. Interestingly, platform's revenue reaches a peak with 300 drivers when surge pricing is applied. This means that the platform can maximise its revenue by applying surge pricing in the moderately undersupplied state even though this results in leaving about 10 % of the requests unserved.

Furthermore, the platform can take advantage of having greater control over the demand–supply ratio. Platforms wish to minimise the number of unanswered ride requests and ensure riders can at least receive an offer. It appears that surge pricing can be used to serve this purpose. Given that surge pricing is paid out of riders' pockets, some riders may reject the offer due to their lower willingness to pay. As shown in Table 2, the number of ride requests with no offer decreases when surge pricing is applied. This is because the ride acceptance rate is slightly higher thanks to surge pricing and some riders reject the offer, making drivers available for the next rides. Notably, in both the balanced and oversupplied states, riders do not reject any rides. Rejections are rather triggered by surge pricing, which is activated only in the undersupplied state.

Our simulation model results suggest that surge pricing does not have a significant impact on the average passenger's waiting time. With 100 and 200 drivers, the waiting time slightly decreases, but with 300 drivers, a minor increase is observed in the surge pricing scenario. Although surge pricing primarily occurs when demand is considerably higher than supply and then a higher waiting time is expected, the drop in demand due to high surge fees can nearly offset the extra waiting time. This is in line with the findings of Cohen et al. (2016), in which they concluded that no correlation exists between passenger waiting time and surge pricing.

Fig. 11 presents the variation in travellers' waiting times (upper row) and drivers' income (lower row) across granular hyperlocal zones in Amsterdam, for fleet sizes of 100 (first column), 200 (second column), and 300 (third column). Note that the travellers' origin and the drivers' initial position are used to assign them to a specific zone. The average waiting time per traveller exhibits an increasing trend as the distance from the city center increases, a pattern that is more noticeable with the fleet size of 300. The average revenue per driver shows minimal variation based on their initial position. Note that the white suburban zones denote areas that were initially unoccupied by drivers at the beginning of the simulation.

The linear function used for estimating the surge multiplier (Eq. (8)) can be replaced with alternative functional forms. Fig. 12 illustrates the maximum recorded surge multipliers for the convex and the concave quadratic functions for the fleet size of 100, compared to the linear function employed in this study. We expect the concave quadratic function to result in higher trip fares and increased number of rejected requests by travellers, and the convex quadratic function to lead to lower trip fares and decreased number of rejected requests by travellers. Furthermore, drivers' income highly depends on the trade-off between the trip fares and the number of rejected requests by travellers.

3. Discussions and implications

Based on the ratio between demand and supply in a given area, the system state can be divided into three categories: undersupplied (more requests than available drivers), balanced, and oversupplied (more available drivers than the number of requests). The implications of such distinction are highlighted in a study by de Ruijter et al. (2022b). They establish how the ratio between supply and demand governs the balance between the matching time (i.e., request match time and driver idle time) and the match quality (i.e., average pickup distance). Furthermore, they show that the system is more efficient in an asymmetrical two-sided market (either

Table 2
Number of rides rejected by riders or receiving no offer.

| Fleet Size | Dynamic Pricing Rejected by Rider | No Offer | Regular Pricing No Offer |
|------------|-----------------------------------|----------|--------------------------|
| 100 | 841 | 5492 | 6218 |
| 200 | 1093 | 2004 | 2908 |
| 300 | 638 | 501 | 787 |

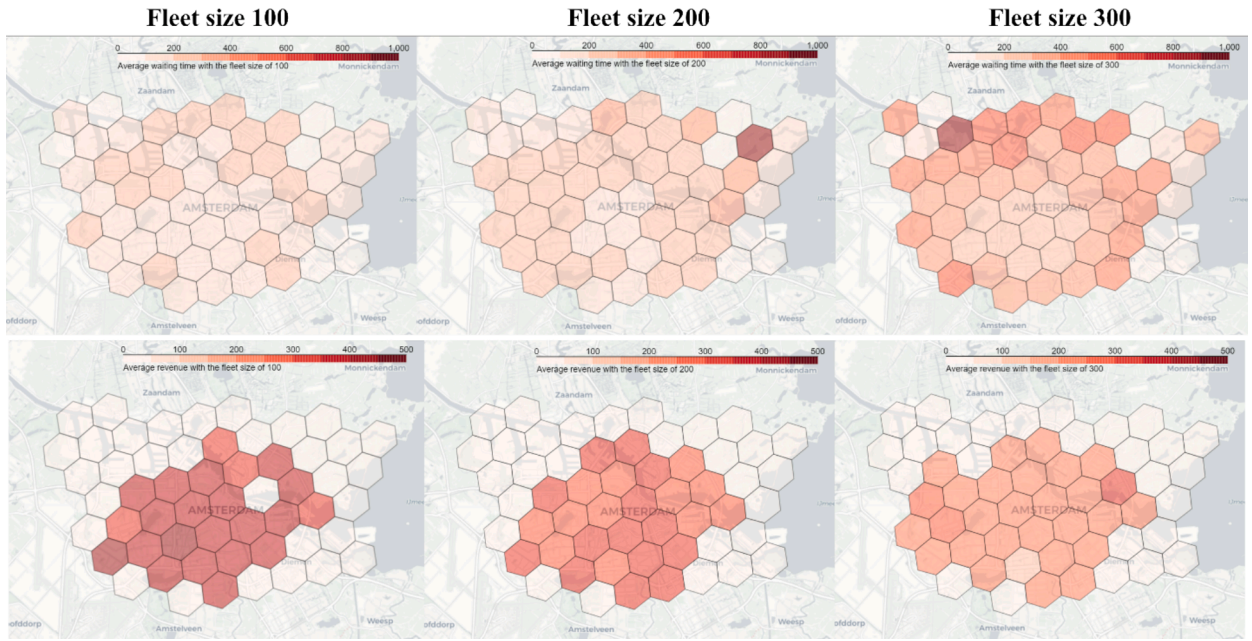


Fig. 11. Average waiting time per traveller (upper row) and average revenue per driver (lower row) in granular hyperlocal zones within Amsterdam.

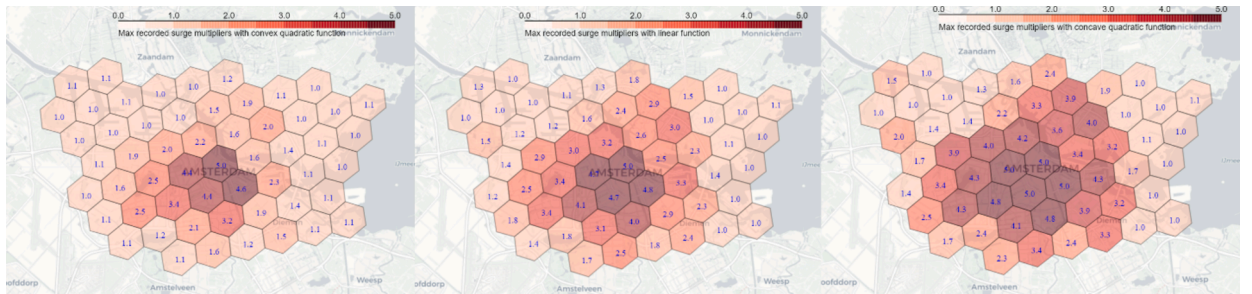


Fig. 12. Maximum surge multipliers for convex (left), linear (middle), and concave (right) surge functions in the granular hyperlocal zones in Amsterdam.

undersupplied or oversupplied) in which one side benefits from higher match quality at the expense of higher matching time for the other side.

In our research, a distinctive pattern is observed in each state, especially in the case of undersupplied conditions. For instance, passengers’ average waiting time decreases with the fleet size (for a fixed demand) except for the undersupplied condition where a larger fleet size leads to a higher match rate and a higher pickup distance resulting in longer passengers’ waiting time. This suggests that the system is less reliable in the undersupplied situation. At the same time, drivers’ average profit decreases when more drivers join the system regardless of the ratio between supply and demand, but the profit variation significantly rises once the system moves beyond the undersupplied state given the greater competition between drivers. The within-day platform’s revenue increases with a larger fleet size in the undersupplied state and does not significantly change in the balanced and oversupplied ranges.

The extent of imbalance between supply and demand is found to be critical. In an extremely undersupplied state where the number of incoming requests in a particular zone is excessively higher than the number of drivers, drivers earn significantly higher income at the expense of a low match rate and long waiting time for passengers. Conversely, an extremely oversupplied system benefits passengers, as it yields a high acceptance rate and short waiting time, but creates intense competition between drivers resulting in longer idle times and, consequently, a lower income. The platform may adopt measures to mitigate the emergence of such extreme situations. When the system is moderately misbalanced, both sides can asymmetrically benefit from the system. Under a moderately undersupplied state passengers experience long waiting times caused by a higher match rate. In a near-balanced or slightly oversupplied situation, the competition between drivers is not as fierce as in the highly oversupplied situation resulting in a higher income/satisfaction, and at the same time, the number of drivers is sufficient to handle all the requests with a relatively short waiting time for passengers. The platform can also benefit directly from the revenue obtained from serving all the requests and indirectly from having a

potentially higher market share, given the higher level of service provisioned.

Comparing the performance of a centralised and a decentralised fleet operation, we argue that the type of fleet associated with drivers' ride acceptance behaviour plays a crucial role in system operations. In the balanced and oversupplied condition, passengers' average waiting time is higher in a DF given that each rejection imposes a delay to the system. This increase in waiting time can reduce the system's capacity and efficiency. This is in line with the findings of Nahmias-Biran et al. (2019) that suggest that a centralised fleet operated by automated vehicles, which dutifully follow the platform instructions on repositioning, performs more efficiently than a decentralised human-driven fleet. They found that such higher efficiency increases the platform market share by four times once an automated fleet takes over the operations. Regarding the relation between drivers' profit and fleet type, we find that human interventions as manifested in ride acceptance decisions cause higher income inequality leading to drivers of a DF having higher income variation, especially if the system is not undersupplied. Such operational differences in CF and DF suggest that ignoring drivers' behaviour when analysing the within-day operations of real-world two-sided ride-sourcing system potentially leads to misrepresentation of the system performance.

The surge pricing model introduced in this study is based on the ratio between supply and demand which is dependent on spatial-temporal factors. Our observations show that surge pricing is more likely to occur in the city centre during peak hours where demand concentrations are higher. Analysing the impacts of surge pricing on system operations suggests that surge pricing can be beneficial for drivers and the platform while having mixed effects on the demand side. On the one hand, an inflated price decreases the utility of a ride for passengers, leading to some offers being rejected by riders. On the other hand, this reduction in the total number of requests leaves the remaining passengers better off as it creates more balance between supply and demand, enhancing the chance of those passengers to receive a ride. This can improve the matching efficiency by allocating the rides to passengers with a high willingness to pay when drivers are scarce (Castillo, 2018). According to Uber, surge pricing is meant to help passengers find a reliable ride by restoring the balance between supply and demand (Uber Marketplace Surge pricing, 2022). In fact, Uber strives to minimise the number of requests that receive no response. Using surge pricing, the platform nudges riders with a lower willingness to pay (e.g. due to differences in urgency, income or quality of alternatives) to opt out. Regardless of the final decision, this choice can boost riders' surplus. Moreover, the platform benefits from higher revenue due to higher trip fares.

From the suppliers' perspective, drivers, on average, earn significantly higher average profits during surge pricing, the extent of which depends on the ratio between demand and supply. Nevertheless, surge pricing reinforces income inequality between drivers. Overall, it seems that surge pricing can offer benefits to all parties at the aggregate level, albeit their advantages are asymmetric. Nevertheless, the surge pricing mechanism is reportedly not transparent, leading to complaints and mistrust which cities may want to avoid if they are focusing on equity and transparency in the mobility system (Ashkrof et al., 2020; Castillo, 2018).

Building upon the critical insights gained in this study, the subsequent policy and practice implications are derived:

- Policies such as enforcing supply caps or regulating the number of drivers during surplus periods should be devised to prevent extreme oversupply, which benefits passengers but results in lower income for drivers due to intensified competition.
- Ride-sourcing platforms need to invest in clear and unbiased communication regarding system operations to build trust and prevent potential misunderstandings among riders and drivers.
- Introducing regulatory measures for ride-sourcing platforms to ensure transparent implementation of surge pricing (e.g., specifying clear criteria that detail the specific circumstances and timing of surge pricing occurrences) and the effective dissemination of information and instructions can be highly beneficial for both riders and drivers.
- Policies may target the reduction of income inequality among drivers caused by ride acceptance behaviour, especially when transitioning from an undersupplied state. Initiatives for rewarding drivers maintaining a fair acceptance rate as opposed to a low acceptance rate even during high-demand or oversupplied conditions could prove advantageous.

4. Conclusions

This study sheds light on the within-day operations of a two-sided ride-sourcing platform in which the interactions between the platform, individual riders and drivers are explicitly taken into account. To this end, we adapt and adopt a discrete-event agent-based simulation for the case of Amsterdam to investigate system dynamics. To the best of our knowledge, this is the first study that models the ride acceptance decisions of ride-sourcing drivers while interacting with riders and the platform in a simulation framework.

Our investigation extends to various perspectives, contrasting centralised fleet operations (mandatory acceptance on each ride request) with decentralised fleet dynamics (ride acceptance decisions made by individual drivers). We explore the implications of supply-demand intensity, ride acceptance rates, and surge pricing, tracking and analysing key performance indicators. This approach allows us to capture potential effects stemming from agents' dynamic behaviour on system operations, providing insights into certain (ir-)regularities.

We have shown that the decisions made by ride-sourcing drivers regarding ride acceptance significantly impact the overall performance of the system, influencing factors such as passengers' waiting time, drivers' revenue, operating costs, and profit. While a higher acceptance rate among drivers may result in more rides, it does not necessarily translate to increased profits for them. Furthermore, surge pricing, despite its asymmetric impact on the demand side through higher trip fares, is found to favour parties involved.

In this study, driver agents are considered homogeneous. However, key individual characteristics such as employment status, experience level with the platform, and working shifts can significantly influence their perceived values of pickup time, thereby affecting their ride acceptance decisions. Future studies may incorporate heterogeneous drivers to account for the implications of these

characteristics on ride acceptance behavior, however, one must collect enough statistically significant information on these variables in order to generate these attributes.

Furthermore, we assumed constant speed for the vehicles in experiments without accounting for traffic congestion on the road network. In cities with a large share of ride-sourcing trips, it is by now well-acknowledged that congestion can lead to prolonged pick-up and travel times, which not only deteriorates driver income but also increases passenger waiting time (Beojone and Geroliminis, 2021; Li et al., 2016; Dandl et al., 2017). Despite low modal share of ride-sourcing platforms in Amsterdam, it might still be insightful to investigate the impact of traffic congestion.

While the focus of this study centres on within-day operations with assumed fixed demand and supply, future research avenues may broaden the scope. Integrating within-day and day-to-day decisions for both demand and supply sides in the ride-sourcing market holds promise for a more comprehensive analysis at the operational level.

CRedit authorship contribution statement

Peyman Ashkrof: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft, visualization. **Farnoud Ghasemi:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing – original draft. **Rafał Kucharski:** Conceptualization, Software, Project administration, Supervision, Writing – review & editing. **Gonçalo Homem de Almeida Correia:** Conceptualization, Project administration, Supervision, Writing – review & editing. **Oded Cats:** Funding acquisition, Conceptualization, Project administration, Resources, Supervision, Writing – review & editing. **Bart van Arem:** Conceptualization, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This research was supported by the CriticalMaaS project (grant number 804469), which is financed by the European Research Council and the Amsterdam Institute for Advanced Metropolitan Solutions. F.G. and R.K. are supported by the National Science Centre in Poland program OPUS 19 (Grant Number 2020/37/B/HS4/01847).

References

- “Standard Mileage Rates for 2021 | Internal Revenue Service”, 2022. *IRS*. Accessed October 31. <https://www.irs.gov/newsroom/irs-issues-standard-mileage-rates-for-2021>. doi:10.1007/s11116-019-10070-2.
- Arentze, T.A., Timmermans, H.J.P., 2004. A Learning-Based Transportation Oriented Simulation System. *Transportation Research Part b: Methodological* 38 (7), 613–633. <https://doi.org/10.1016/j.trb.2002.10.001>.
- Ashkrof, P., Homem, G., de Almeida Correia, Oded Cats, Bart van Arem, 2020. Understanding Ride-Sourcing Drivers' Behaviour and Preferences: Insights from Focus Groups Analysis. *Research in Transportation Business and Management* 37 (March). Elsevier, 100516. <https://doi.org/10.1016/j.rtbm.2020.100516>.
- Ashkrof, P., Homem, G., de Almeida Correia, Oded Cats, Bart van Arem, 2022. “Ride Acceptance Behaviour of Ride-Sourcing Drivers.” *Transportation Research Part C: Emerging Technologies* 142 (June). Elsevier Ltd 103783. <https://doi.org/10.1016/j.trc.2022.103783>.
- Ashkrof, P., Homem, G., de Almeida Correia, Oded Cats, Bart van Arem, 2023a. On the Relocation Behaviour of Ride-Sourcing Drivers. *Transportation Letters* March, 1–8. <https://doi.org/10.1080/19427867.2023.2192581>.
- Ashkrof, P., Homem, G., de Almeida Correia, Oded Cats, Bart van Arem, 2023b. On the Relocation Behavior of Ride-Sourcing Drivers. *Transportation Letters* March, 1–8. <https://doi.org/10.1080/19427867.2023.2192581>.
- Beojone, C.V., Geroliminis, N., 2021. On the inefficiency of ride-sourcing services towards urban congestion. *Transportation Research Part c: Emerging Technologies* 124, 102890. <https://doi.org/10.1016/j.trc.2020.102890>.
- Bokányi, E., Hannák, A., 2020. Understanding Inequalities in Ride-Hailing Services Through Simulations. *Scientific Reports* 10 (1), 1–11. <https://doi.org/10.1038/s41598-020-63171-9>.
- Cachon, G.P., Daniels, K.M., Lobel, R., 2017. The Role of Surge Pricing on a Service Platform with Self-Scheduling Capacity. *Manufacturing and Service Operations Management* 19 (3), 368–384. <https://doi.org/10.1287/msom.2017.0618>.
- Castillo, J.C., 2018. Who Benefits from Surge Pricing? *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3245533>.
- Chen, Le, Alan Mislove, and Christo Wilson. 2015. “Peeking beneath the Hood of Uber.” *Proceedings of the ACM SIGCOMM Internet Measurement Conference, IMC 2015*—October (October). Association for Computing Machinery: 495–508. doi:10.1145/2815675.2815681.
- Cohen, P., Hahn, R., Hall, J., Levitt, S., Metcalfe, R., 2016. Using Big Data To Estimate Consumer Surplus. *Nber Working Paper Series* 42.
- Dandl, F., Bracher, B., Bogenberger, K., 2017. Microsimulation of an autonomous taxi-system in Munich. In: *5th IEEE International Conference on Models and Technologies for Intelligent Transportation Systems*. <https://doi.org/10.1109/MTITS.2017.8005628>.
- de Ruijter, A., Cats, O., Kucharski, R., van Lint, H., 2022a. Evolution of Labour Supply in Ridesourcing. *Transportmetrica B* 10 (1), 599–626. <https://doi.org/10.1080/21680566.2021.2024917>.
- de Ruijter, A., Cats, O., van Lint, H., 2022b. Emerging Dynamics in Ridesourcing Markets. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4258151>.
- GemeenteAmsterdam. 2019. “Agenda Taxi 2020–2025. Technical Report. Verkeer En Openbare Ruimte, Gemeente Ams- Terdam.”.
- Guda, H., Subramaniana, U., 2019. Your Uber Is Arriving: Managing on-Demand Workers through Surge Pricing, Forecast Communication, and Worker Incentives. *Management Science* 65 (5), 1995–2014. <https://doi.org/10.1287/mnsc.2018.3050>.
- Kucharski, R., Cats, O., 2022. Simulating Two-Sided Mobility Platforms with MaaSsim. *PLoS ONE* 17 (6 June), 1–20. <https://doi.org/10.1371/journal.pone.0269682>.
- Li, Manzi, Gege Jiang, and Hong K. Lo. 2022. “Pricing Strategy of Ride-Sourcing Services under Travel Time Variability.” *Transportation Research Part E: Logistics and Transportation Review* 159 (March 2020). Elsevier Ltd: 102631. doi:10.1016/j.tre.2022.102631.
- Li, Z., Hong, Y., Zhang, Z., 2016. An Empirical Analysis of On-Demand Ride Sharing and Traffic Congestion. *Proc. International Conference on Information Systems (ICIS '16)*, Available at SSRN.

- Marquet, Oriol. 2020. "Spatial distribution of ride-hailing trip demand and its association with walkability and neighborhood characteristics". *Cities* Elsevier: 102926. doi:10.1016/j.cities.2020.102926.
- Martínez, L. Miguel, Gonçalo Homem de Almeida Correia, Filipe Moura, and Mafalda Mendes Lopes. 2017. "Insights into Carsharing Demand Dynamics: Outputs of an Agent-Based Model Application to Lisbon, Portugal." *International Journal of Sustainable Transportation* 11 (2). Taylor & Francis: 148–159. doi:10.1080/15568318.2016.1226997.
- Nahmias-Biran, B.H., Oke, J.B., Kumar, N., Basak, K., Araldo, A., Seshadri, R., Akkinepally, A., Azevedo, C.L., Ben-Akiva, M., 2019. From Traditional to Automated Mobility on Demand: A Comprehensive Framework for Modeling On-Demand Services in SimMobility. *Transportation Research Record* 2673 (12), 15–29. <https://doi.org/10.1177/0361198119853553>.
- Peng, C., Jin, X., Wong, K.-C., Shi, M., Liò, P., 2012. Collective Human Mobility Pattern from Taxi Trips in Urban Area. *PLOS ONE*. <https://doi.org/10.1371/annotation/f0d48839-ed4b-4cb2-822a-d449a6b4fa5d>.
- Qian, X., Ukkusuri, S.V., 2015. Spatial variation of the urban taxi ridership using GPS data. *Applied Geography*. 31–42. <https://doi.org/10.1016/j.apgeog.2015.02.011>.
- Ramezani, Mohsen, Yue Yang, Jacob Elmasry, and Porsiem Tang. 2022. "An Empirical Study on Characteristics of Supply in E-Hailing Markets: A Clustering Approach." <https://doi.org/10.1080/19427867.2022.2079869>. Taylor & Francis. doi:10.1080/19427867.2022.2079869.
- Rapport, R.O.A., 2023. *Gebruikers Uber app, gemaakte trips, verdiensten en tevredenheid*. Research Centre for Education and the Labour Market at Maastricht University.
- Ruch, Claudio, H Sebastian, and Emilio Frazzoli. 2018. "AMoDeus , a Simulation-Based Testbed for Autonomous Mobility-on-Demand Systems," 3639–3644.
- Santi, P., Resta, G., Szell, M., Ratti, C., 2014. Quantifying the benefits of vehicle pooling with shareability networks. *Computer Sciences*. 13290–13294. <https://doi.org/10.1073/pnas.1403657111>.
- Son, Dong Hoon. 2023. "On-Demand Ride-Sourcing Markets with Cryptocurrency-Based Fare-Reward Scheme." *Transportation Research Part E: Logistics and Transportation Review* 171 (March 2022). Elsevier Ltd: 103027. doi:10.1016/j.tre.2023.103027.
- Uber Marketplace Surge pricing. 2022. "Uber Marketplace Surge Pricing." <https://www.uber.com/us/en/marketplace/pricing/surge-pricing/>.
- Uber, 2022c. "Privacy protection for rides | Uber" <https://help.uber.com/riders/article/privacy-protection-for-riders?nodeId=74ecf382-9922-40d9-b4f6-6e5b6f30bea4>.
- Uber. 2022a. "H3: Uber's Hexagonal Hierarchical Spatial Index | Uber Blog." <https://www.uber.com/en-NL/blog/h3/>.
- Uber. 2022b. "Uber Marketplace Matching | Uber." <https://www.uber.com/us/en/marketplace/matching/>.
- Wang, Z., Kun, Fu., Ye, J., 2018. Learning to Estimate the Travel Time. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1, 858–866. <https://doi.org/10.1145/3219819.3219900>.
- Wang, D.i., Miwa, T., Morikawa, T., 2021. Comparative Analysis of Spatial–Temporal Distribution between Traditional Taxi Service and Emerging Ride-Hailing. *ISPRS International Journal of Geo-Information* 10 (10), 690. <https://doi.org/10.3390/ijgi10100690>.
- Xu, K.e., Sun, L., Liu, J., Wang, H., 2018. An Empirical Investigation of Taxi Driver Response Behavior to Ride-Hailing Requests: A Spatio-Temporal Perspective. *PLoS ONE* 13 (6), 1–17. <https://doi.org/10.1371/journal.pone.0198605>.
- Zha, L., Yin, Y., Yang, H., 2016. "Economic Analysis of Ride-Sourcing Markets." *Transportation Research Part C: Emerging Technologies* 71. Elsevier Ltd 249–266. <https://doi.org/10.1016/j.trc.2016.07.010>.
- Zha, L., Yin, Y., Yuchuan, Du., 2018a. Surge Pricing and Labor Supply in the Ride-Sourcing Market. *Transportation Research Part b: Methodological* 117. Elsevier Ltd 708–722. <https://doi.org/10.1016/j.trb.2017.09.010>.
- Zha, L., Yin, Y., Zhengtian, Xu., 2018b. Geometric Matching and Spatial Pricing in Ride-Sourcing Markets. *Transportation Research Part c: Emerging Technologies* 92 (April). Elsevier, 58–75. <https://doi.org/10.1016/j.trc.2018.04.015>.