https://doi.org/10.1016/j.trc.2021.103293

Learning ride-sourcing drivers' customer-searching behavior: A dynamic discrete choice approach

Junji Urata[1], Zhengtian Xu*[2], Jintao Ke[3], Yafeng Yin[4], Guojun Wu[5], Hai Yang[6], and Jieping Ye[7]

[1]*Department of Civil Engineering, University of Tokyo, Tokyo, Japan*
[2]*Department of Civil and Environmental Engineering, The George Washington University, Washington DC, Unites States*
[3]*Department of Logistic and Maritime Studies, The Hong Kong Polytechnic University, Hong Kong, China*
[4]*Department of Civil and Environmental Engineering, University of Michigan, Ann Arbor, United States*
[5]*Data Science Program, Worcester Polytechnic Institute, Worcester, United States*
[6]*Department of Civil and Environmental Engineering, The Hong Kong University of Science and Technology, Hong Kong, China*
[7]*AI Labs, Didi Chuxing, Beijing, China*

May 26, 2022

Abstract

Ride-sourcing drivers spend a significant portion of their service time being idle, during which they can move freely to search for the next customer. Such customer-searching movements, while not being directly controlled by ride-sourcing platforms, impose great impacts on the service efficiency of ride-sourcing systems and thus need to be better understood. To this purpose, we design a dynamic discrete choice framework by modeling drivers' customer search as absorbing Markov decision processes. The model enables us to differentiate three latent search movements of idle drivers, as they either remain motionless, cruise around without a target area, or reposition towards specific destinations. Our calibration takes advantage of large-scale empirical datasets from Didi Chuxing, including the transaction information of five million passenger requests and the trajectories of 32,000 affiliated drivers. The calibration results uncover the variations of drivers' attitudes in customer search across time and space. In general, ride-sourcing drivers do respond actively and positively to the repetitive market variations when idle. They are comparatively more mobile at high-demand hotspots while preferring to stay motionless in areas with long time of waiting being expected. Our results also suggest that drivers' search movements are not confined to local considerations. Instead, idle drivers show a clear tendency of repositioning towards the faraway hotspots, especially during the evening when the demand cools down in the suburb. The discrepancies between full-time and part-time drivers' search behavior are also examined quantitatively.

Keywords: Ride-sourcing service, customer search, driver behavior, dynamic discrete choice

*Corresponding author. E-mail address: zhengtian@gwu.edu (Z. Xu).

1 Introduction

The maturity of mobile internet technology catalyzes on-demand ride-sourcing services provided by companies like Uber, Lyft, and DiDi Chuxing. Compared to traditional street-hailing taxi services, these emerging ride-sourcing services significantly reduce the meeting frictions between riders and drivers, and thus become unprecedentedly popular among urban travelers in recent years (Conway et al., 2018). The great success of ride-sourcing services has attracted a lot of interests on the analysis and management of such on-demand ride-hailing systems (see Wang and Yang, 2019 for a recent review). However, less attention has been paid to the behaviors of ride-sourcing drivers, partially due to the lack of access to service data. To better serve ride-sourcing drivers and facilitate a cohesive platform environment, it is crucial for system managers and policy-makers to understand drivers' behaviors and concerns in service provision (see, e.g., Sun et al., 2019; Xu et al., 2020). By virtue of comprehensive empirical data from Didi Chuxing, this paper thus aims to comprehend the behavior of ride-sourcing drivers in customer search, which constitutes a significant portion of drivers' service time and strongly associates with their profitability.

Many empirical studies have been carried out to investigate the customer-searching behavior of taxi drivers, who shares a significant similarity with drivers in ride-sourcing markets. A group of researchers from Hong Kong first applied multinomial logit (MNL) models to capture strategic zonal choices of Hong Kong taxi drivers' customer search (see, e.g., Wong et al., 2014b, 2015). They proposed a cell-based logit-opportunity model to tackle the local customer-searching behavior of taxis by considering the opportunities along search paths. Recently, Tang et al. (2019) argued that between different destination choices of vacant taxis, there are substantial overlaps in paths, which invalidate the use of MNL models. Instead, they proposed a mixed path size logit-based customer-searching model and tested its effectiveness in predicting routing choices over the trajectory of 36,000 taxis in Beijing. Although these static search models substantially facilitate empirical calibrations, they fall short in capturing the dynamic choice behavior of drivers under the highly varying market conditions. Zheng et al. (2018) modeled vacant taxi drivers' anticipatory behavior by using a time-dependent framework. However, their study focused on the one-shot decision choices of taxi drivers between urban areas and condensed-demand areas, such as airports and railway stations, and is unsuitable for behavioral calibration. One of the major difficulties in calibrating the search behavior is that drivers' trajectories do not fully reflect their real preferences. It is common for drivers to get matched to passengers before reaching the actual cruising destinations, especially in app-based ride-hailing markets. Sometimes, drivers do not even have specific search destinations in mind. Therefore, we are in need of a behavioral model that can cope with the intense market variations and identify drivers' latent search patterns with modeling differentiation.

An alternative way of modeling drivers' customer-searching movement is to formulate it as Markov decision process (MDP). Oftentimes, an MDP framework is coupled with learning approaches to seek for the optimal searching policy for idle drivers. Recently, Liu et al. (2013), Verma et al. (2017), Gao et al. (2018), and Lin et al. (2018) employed Q-learning to investigate the optimal dynamic routing strategy. This approach does not explicitly characterize the intervening opportunity, but instead implicitly incorporates it into the action rewards through learning. Besides, Qu et al. (2014), Rong et al. (2016), Yu et al. (2020), and Shou et al. (2020) specified structured reward functions based on various zonal features and then solved the problem using dynamic programming. However, the weights of different features were exogenously given to concretize certain search objectives. In general, these learning-based approaches aim at deriving the optimal policy of repositioning drivers for centralized control, and lack the behavioral implications of individual drivers.

To fill this research gap, a dynamic discrete choice model is developed in this paper to investigate the customer-searching behavior of ride-sourcing drivers. The model translates market conditions, including both supply and demand information, into a spatiotemporal continuum of opportunity values to idle drivers. Then, by adopting an absorbing MDP framework, we formulate and evaluate drivers' manifold decisions underlying customer search. To foster a comprehensive understanding, heterogeneous behavior among different driver types and time of day is further explored and compared. It is worth noting that our model differentiates three modes of search movements, respectively being staying motionless, cruising around without a target area, or repositioning towards a specific destination. These movements, although corresponding to completely different mentality of drivers, are challenging to separate through trajectory data. Indeed, to the best of our knowledge, none of the previous studies has investigated or calibrated the driver behavior in this regard. Leveraging large-scale empirical datasets (with uninterrupted trajectories of 32,000 drivers and transaction information of five million trip requests), this study is dedicated to deepening our understanding of ride-sourcing drivers' customer-searching behavior and provide policy insights for the platform's labor supply management.

The remainder of this paper is organized as follows. Section 2 details the dynamic discrete choice model designed for learning ride-sourcing drivers' customer-searching behavior, while Section 3 illustrates the data preparation for model calibration. Section 4 presents the results of parametric estimation and then interprets the behavioral implications of drivers in zonal search. At the end, Section 5 concludes the paper and points out an future research avenue.

2 Dynamic Discrete Choice Model

This section first introduces the dynamic discrete choice model that we formulate in line with ride-sourcing drivers' customer-searching movements. We first dissect the process of how idle drivers search for customers, and then formulate a mathematical model to delineate drivers' sequential choice-making. The method for parametric estimation is also discussed.

2.1 Drivers' customer-searching movements

When ride-sourcing drivers are idle, they enjoy full freedom of deciding where and how to search for the next customer based on the market condition they perceive. They may either remain motionless awaiting the next match, cruise around the neighborhood to actively search for customers, or reposition themselves towards a target hotspot. To cope with these different ways of customer search, we materialize idle drivers' movements as cyclic decision-making by segmenting them into a series of steps. Within each step, a driver either stays in the current zone, or chooses her next destination from the finite set of adjacent zones and then moves forward. A series of choice decisions are made sequentially by a driver along a chain of connecting zones, until she successfully gets matched with a customer (see Figure 1 for a graphical explanation).

For the convenience of behavioral analysis, we further treat the sequential movements as Markov decision processes, assuming that the decision-making of drivers at each step is independent of her previous choices. Notwithstanding, drivers do plan several steps ahead when making each zonal choice decision. Note that such a setting adheres well with that of the dy-

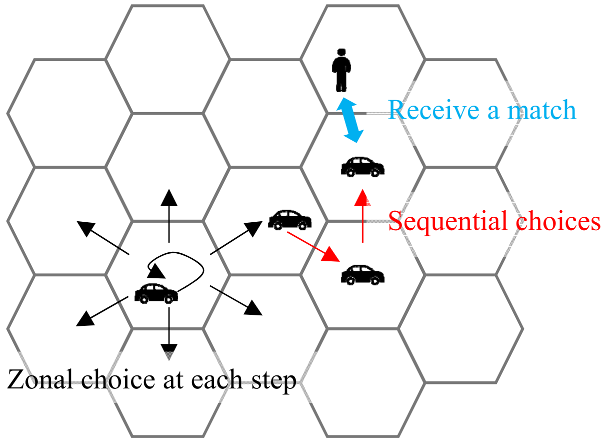

Figure 1: Sequential movements of idle drivers in customer search

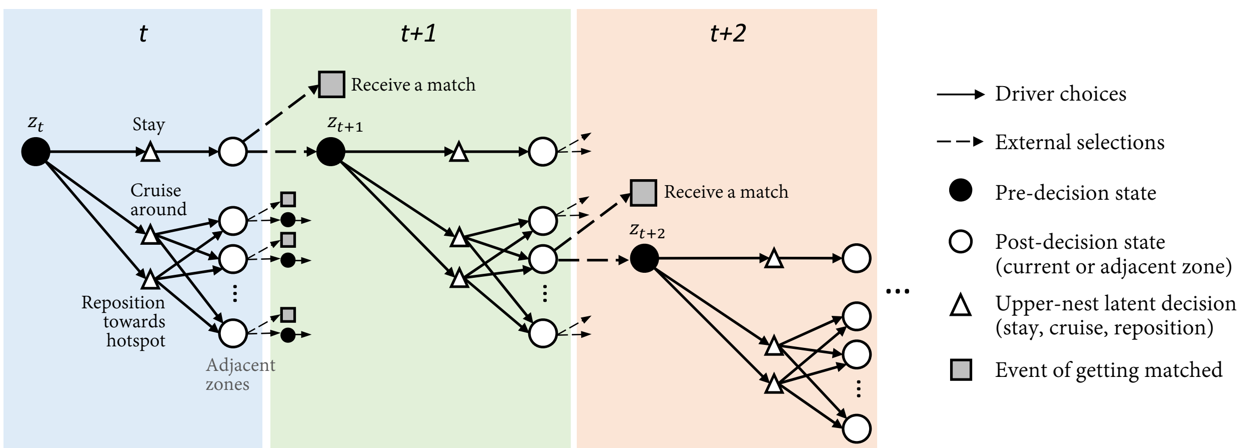

Figure 2: Decision-making processes of ride-sourcing drivers

namic discrete choice theory (Rust, 1987). We thus introduce absorbing Markov chains in this paper to describe the dynamic discrete choices of drivers under finite evaluation horizons.

Figure 2 illustrates the cross-nested structure of drivers' sequential searching choices. During the process, a driver at each stage can select to either rest for a while in the current zone or move to one of the adjacent zones, based on the expected utility of each choice. The process continues as the cumulative probability of being matched increases along the trajectory. Within each stage, the upper nest indicates three latent scenarios for drivers' zonal choices: to either remain motionless, cruise nearby without a specific destination, or reposition toward a hotspot area. Correspondingly, the lower nest expands the potential zonal choices under each scenario. Built on such a cross-nested structure, our model differentiates the latent searching scenarios and describes drivers' movements more precisely.

2.2 Model formulation

We formulate the model based on the multinomial-type dynamic discrete choice model proposed by Rust (1987). Let z_t be the location of a driver at time t. For brevity, we use z_t as an abbreviation for vector (z_t, t), which marks the state of an idle driver being at zone z_t at time t. To facilitate understanding, we first present a base model for non-nested choice cases and then extend it to more complicated contexts of nested and cross-nested choices. The driver-specific indicators are omitted in this subsection for clarity.

Base model

With a non-nested choice structure, the expected utility of an idle driver $V^T(z_t)$ at state z_t under a time horizon of T adheres to the following relationship,

$$V^T(z_t) = \mathbb{E}\left[\max_{z_{t+1}\in C(z_t)} \Big(v(z_{t+1}|z_t;\theta) + \rho_{z_t}\beta V^T(z_{t+1}) + \epsilon(z_{t+1})\Big)\right]$$

where $C(z_t)$ denotes the choice set of drivers at z_t; ρ_{z_t} is the probability of drivers remaining unmatched coming through state z_t; β denotes a time-discounting factor ($0 \leq \beta \leq 1$). Specifically, on the right side, $v(z_{t+1}|z_t;\theta)$ represents the observed value of an idle driver transitioning from z_t to z_{t+1}, with θ being a vector of parameters; $\rho_{z_t}\beta V^T(z_{t+1})$ denotes the discounted value of attaining the state z_{t+1}; and $\epsilon(z_{t+1})$ entails the error associated with unobserved factors post state z_{t+1}. Note that ρ_{z_t} in the equation acts similarly as the survival probability in a general dynamic programming (Rust, 2016). In operations, drivers' customer-searching movements are forced to termination once they receive matches from the platform. The probability ρ_{z_t} thus captures the chance that drivers' idleness continues for at least one more periods following z_t. It is assumed that drivers' customer-searching movements follow utility maximization in each step, and the values of a driver being matched during the search or reaching the end of a horizon are both set to zero in the model.

Assuming that the error term ϵ follows a Gumbel distribution that has a scaling factor μ ($\mu \geq 1$) indicating the degree of independence for the unobserved confounders, yields the following choice probability (with parameter θ omitted for clarity, same below),

$$P^T(z_{t+1}|z_t) = \frac{\exp\left(\mu\left(v(z_{t+1}|z_t) + \rho_{z_t}\beta V^T(z_{t+1})\right)\right)}{\sum_{z\in C(z_t)}\exp\left(\mu\left(v(z|z_t) + \rho_{z_t}\beta V^T(z)\right)\right)}, \qquad \forall\, z_{t+1} \in C(z_t)$$

where the expected value $V^T(\cdot)$ of each choice can be derived as

$$V^T(z_t) = \frac{1}{\mu}\cdot\ln\sum_{z_{t+1}\in C(z_t)}\exp\Big(\mu\Big(v(z_{t+1}|z_t) + \rho_{z_t}\beta V^T(z_{t+1})\Big)\Big)$$

The choice probability $P^T(z_{t+1}|z_t)$ is then utilized as a surrogate of the state transition probability from z_t to z_{t+1}, following the same treatment adopted by the recursive logit model (Fosgerau et al., 2013). It is worth noting that the time horizon T is incorporated to avoid the state explosion of a time-expanded network amid the varying market conditions. Behaviorally, it could be seen as the upper bound sensed by drivers for potential search duration. With the horizon T specified externally, the value of $V^T(z_t)$ can be calculated via backward induction detailed in Section 2.3.

Nested model

As shown by Figure 2, drivers' customer-searching decisions in each step are decomposed as two levels of nests. They first make an upper-nest decision among stay, cruising, or repositioning, and then select an adjacent zone out of the lower nest to move one-step forward. Accordingly, we extend the above base model to accommodate such a context of a nested structure (with no cross alternatives). The value function under the nested choices is reformulated as follows,

$$V^T(z_t) = \mathbb{E}\left[\max_{l \in C_u(z_t)} \left(\epsilon(l) + \max_{z_{t+1} \in C(z_t,l)} \left(v(z_{t+1}|z_t) + \rho_{z_t}\beta V^T(z_{t+1}) + \epsilon(z_{t+1}, l)\right)\right)\right]$$

where l is the choice in the upper nest $C_u(\cdot)$. Let $\epsilon(z_{t+1}, l) = \frac{1}{\mu_l}\varepsilon(z_{t+1})$ and the error term $\varepsilon(z_{t+1})$ follow a standard Gumbel distribution. The value function and inclusive (log-sum) utility v_u are given by

$$V^T(z_t) = \mathbb{E}\left[\max_{l \in C_u(z_t)} \left(\epsilon(l) + v_u(l|z_t) + \epsilon_u(l)\right)\right]$$

$$v_u(l|z_t) = \frac{1}{\mu_l} \ln \sum_{z_{t+1} \in C(z_t,l)} \exp\left(\mu_l\left(v(z_{t+1}|z_t) + \rho_{z_t}\beta V^T(z_{t+1})\right)\right), \quad \forall\, l \in C_u(z_t)$$

Again, assuming $\epsilon(l) + \epsilon_u(l) = \frac{1}{\mu_n}\varepsilon(l)$ with $\varepsilon(l)$ following a standard Gumbel distribution yields

$$V^T(z_t) = \frac{1}{\mu_n} \ln \sum_{l \in C_u(z_t)} \exp\left(\mu_n v_u(l|z_t)\right), \quad (\mu_l \geq \mu_n \geq 1) \tag{1}$$

and the choice probability exhibits the following multiplicative form,

$$\begin{aligned} P^T(z_{t+1}|z_t) &= P^T(z_{t+1}|l) \cdot P^T(l|z_t) \\ &= \frac{\exp\left(\mu_l\left(v(z_{t+1}|z_t) + \rho_{z_t}\beta V^T(z_{t+1})\right)\right)}{\sum_{z \in C(z_t,l)} \exp\left(\mu_l\left(v(z|z_t) + \rho_{z_t}\beta V^T(z)\right)\right)} \cdot \frac{\exp\left(\mu_n v_u(l|z_t)\right)}{\sum_{k \in C_u(z_t)} \exp\left(\mu_n v_u(k|z_t)\right)} \end{aligned}$$

Cross-nested model

Particularly, in each step, the zonal choices in the lower nest is cross shared by the two upper-nest intentions, i.e. cruising and repositioning. Further extensions for the cross-nested choices are obtained through generating functions. In the cases of nested logit (NL) and cross-nested logit (CNL), the generating functions are respectively given as follows (Train, 2009),

$$\text{NL}: G(\boldsymbol{Y}) = \sum_{l=1}^{K} \left(\sum_{z \in B_l} Y_z^{\mu_l}\right)^{\frac{1}{\mu_l}}$$

$$\text{CNL}: G(\boldsymbol{Y}) = \sum_{l=1}^{K} \left(\sum_{z \in B_l} (\alpha_{z,l} Y_z)^{\mu_l}\right)^{\frac{1}{\mu_l}}$$

where $\alpha_{z,l}$ is an allocation parameter that reflects the likelihood of alternative z being a member of nest B_l with $\alpha_{zl} \geq 0$ and $\sum_l \alpha_{zl} = 1$; the function Y_z characterizes $\exp(V(z))$ in both models, where $V(z)$ stands for the value of state z. In comparison, we have $V^T(z_t)$ continue to hold as Eq. (1) and the log-sum (inclusive) utility v_u under the cross-nested logit be respecified as

$$v_u(l|z_t) = \frac{1}{\mu_l} \ln \sum_{z_{t+1} \in C(z_t,l)} \left(\alpha_{z_{t+1},l} \exp\left(v(z_{t+1}|z_t) + \rho_{z_t}\beta V^T(z_{t+1})\right)\right)^{\mu_l}, \quad \forall\, l \in C_u(z_t) \tag{2}$$

and the choice probability now sums over all the probability multiplications,

$$P^T(z_{t+1}|z_t) = \sum_{l \in C_u(z_t)} P^T(z_{t+1}|l) \cdot P^T(l|z_t) \tag{3}$$

$$= \sum_{l \in C_u(z_t)} \frac{\left(\alpha_{z_{t+1},l} \exp\left(v(z_{t+1}|z_t) + \rho_{z_t}\beta V^T(z_{t+1})\right)\right)^{\mu_l}}{\sum_{z \in C(z_t,l)} \left(\alpha_{z,l} \exp\left(v(k|z_t) + \rho_{z_t}\beta V^T(k)\right)\right)^{\mu_l}} \cdot \frac{\exp\left(\mu_n v_u(l|z_t)\right)}{\sum_{k \in C_u(z_t)} \exp\left(\mu_n v_u(k|z_t)\right)}$$

with $\mu_l \geq \mu_n \geq 1$.

2.3 Model estimation

All parameters in the model $(\theta, \mu, \beta, \alpha)$, collectively referred to as Θ, can be estimated by maximizing the following log-likelihood (LL), i.e.,

$$\text{LL}(\Theta) = \sum_{i,t} \delta_{i,t}(z_{t+1}|z_t) \cdot \log P^T(z_{t+1}|z_t;\Theta)$$

where $\delta_{i,t}$ indicates the choice made by individual driver i at state z_t. The indicator equals to 1 if the driver chose z_{t+1} and 0 otherwise. Note that given all the parameters Θ, the choice probabilities $\boldsymbol{P^T}$ can be calculated through Eq. (3) using the expected utilities $\boldsymbol{V^T}$, which is treated as fixed via backward induction. As per the recursive Eqs. (1) and (2), the utilities $\boldsymbol{V^T}$ of preceding states at time t can be computed based on the succeeding states' at $t+1$. Since the terminating state value at the end of each horizon is prespecified as 0, we can thus recursively calculate the utility of all the states. The estimation of Θ can be carried out in a similar fashion as that for the network Generalized Extreme Value models (Daly and Bierlaire, 2006). The Student's t-tests are conducted to examine the significance of parametric estimates in the model, while Watson and Westin pooling tests are applied for model comparisons (Watson and Westin, 1975).

3 Data Description

Two datasets from Didi Chuxing are used for model calibration. One records the trajectory information of drivers while the other contains the complete transaction information of trip requests from one medium-sized city in China, spanning over the 10 weekdays in August 7-18, 2017.

The trajectory data comprises the spatiotemporal records of 32,000 DiDi drivers. Each trajectory characterizes a series of status points of a particular driver recorded every 3 seconds throughout a nonstop customer-searching segment; and each status point consists of a timestamp, longitude and latitude coordinates of a driver, as well as her service state at that moment, either being idle (waiting to be matched), deadheading (picking up a customer), or occupied (delivering a customer to the destinations). The entire city is partitioned into regular hexagonal lattices, each being side-connected with six adjacent ones. The major advantage for hexagonal partition is that each unit has six symmetrically equivalent and unambiguous side-adjacent neighbors, while square partition results in two different types of neighboring, respectively being side-connected or corner-connected. We map the latitude/longitude sequences to the 660-meter-side-length hexagonal lattices (Sahr et al., 2003) to produce the base sample for movement identification. Overall, the searching distances of idle drivers were not overwhelming in this

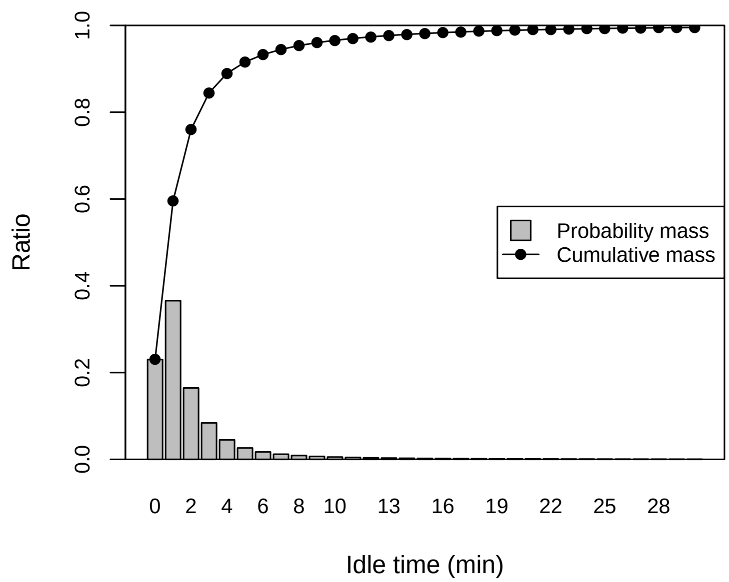

Figure 3: Histogram of drivers' idle time staying in one zone

city. Above 80% of the trajectories in the dataset cover no more than 3 zones. The case when an idle driver remains in one specific zone for a duration exceeding a threshold κ is regarded as a "stay", while the case that a driver repositions to a neighboring zone within κ is taken as a "move". Those searching trajectories with a total duration less than κ are removed from our sample. Figure 3 plots the distribution of drivers' consecutive idle time of staying at one single zone. We finalize the selection of κ to be four minutes to buffer the time needed by divers to wait for traffic signals and pass through zones.

The transaction dataset contains the information of approximately five million request records. Each record comprises the timestamp when the trip gets requested, the origin/destination of the trip, the matching/pickup/delivery time of the passenger, the driver in service, and the trip fare. The transaction information is aggregated by hexagonal zones and time of day to produce various covariates associated with the market conditions. It is worth noting that to overcome the sparsity of data under the high-granularity partition, this study does not differentiate between days. All the records fall in the same time interval across different days are merged together to create the within-day explanatory variables. Consequently, the coefficient estimation results likely reflect drivers' behavior in response to the regular market variations from day to day rather than the transient and random fluctuations.

As an overview of the datasets, Figures 4a and 4b visualize the spatial distributions of drivers' movements and the number of orders being requested, respectively. In Figure 4a, the blue arrow represents the number of samples moving between two adjacent zones ("move" action), while the green circle refers to the number of samples staying in the current zone ("stay" action). Figure 4b shows the heat map of the number of orders averaged over all the five-minute intervals for each hexagonal zone. It can be observed that the variations of supply and demand are highly consistent in space. Figure 5 displays a boxplot for ratio ρ_{z_t} that idle drivers remain unmatched in each 12-min interval of a day. As clearly shown by the figure, the matching probability of drivers changes drastically throughout a day. Following the searching processes detailed in Section 2, this essentially implies the time-varying range and composition of drivers' zonal choices, thereby

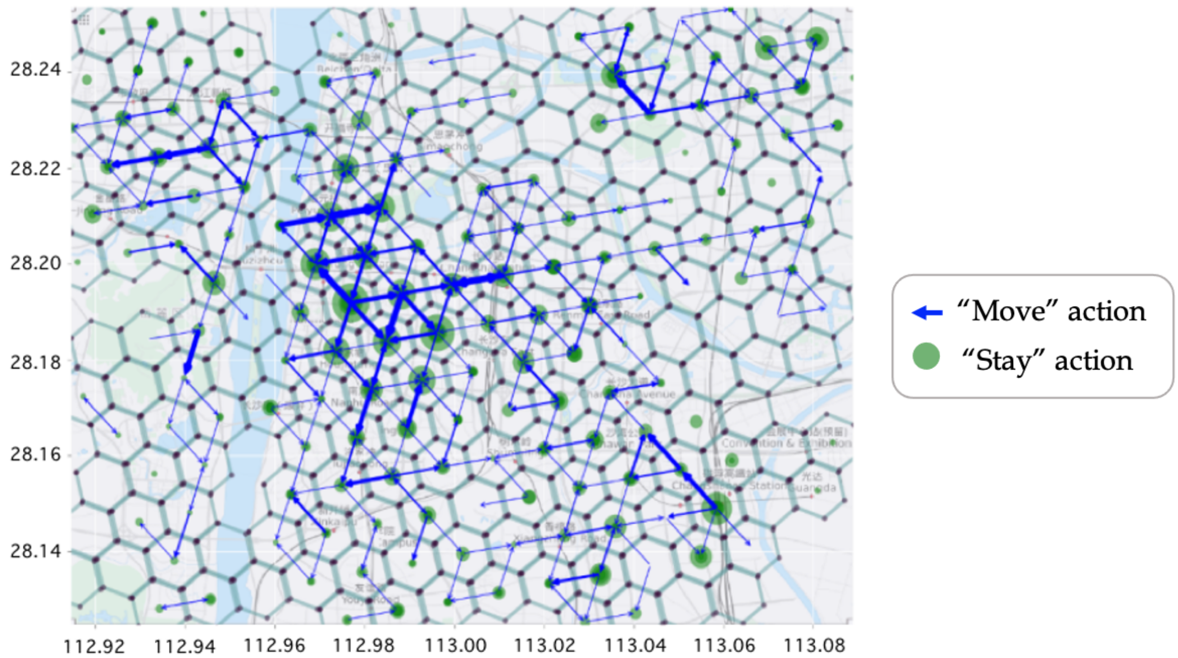

(a) Drivers' movements

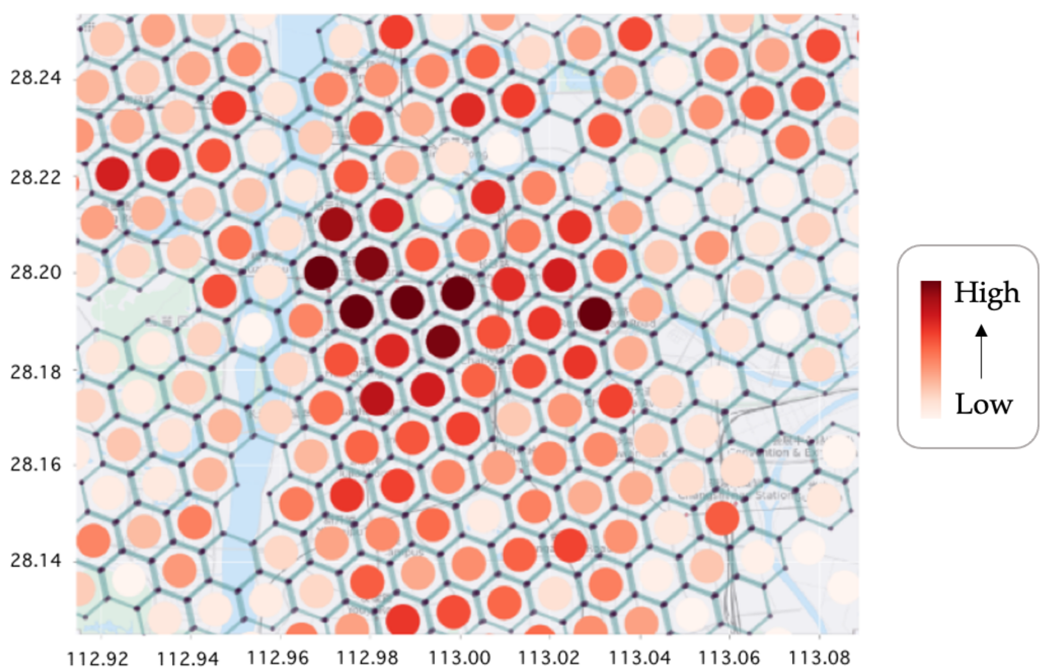

(b) Order requests

Figure 4: Spatial distributions of a) drivers' search trajectories and b) orders requested. The blue arrows and the green circles in the top figure represent the number of "move" and "stay" actions, with thicker arrows and larger circles standing for higher counts, respectively. The heat map at the bottom illustrates the contrasts of ride requests over the space, with deeper colors representing higher service demand.

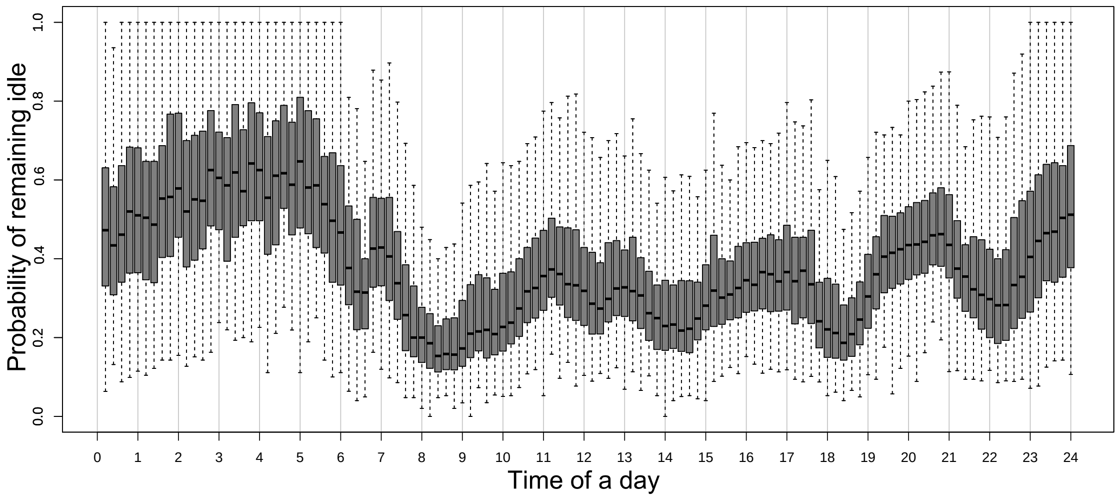

Figure 5: Probability of idle drivers remaining unmatched in each 12-min interval of a day. The bold line traces the median ratios across all the zones within each interval, while the dark gray boxes and the dotted segments display the [25th, 75th] and the [5th, 95th] percentile ranges, respectively.

questing the need for a dynamic model to delineate the choice-making.

In addition, depending on their service patterns, drivers considered in this study are divided into two classes, respectively as full-time and part-time drivers. Full-time drivers are those with the most extended service hours but also the highest earners and most profitable ones. They are mostly daytime workers and display very little within-group heterogeneity in terms of profitability. In contrast, part-time drivers are active for less amount of time and show significant within-group disparities in driving and working experiences, as well as profitability and earnings.

4 Model Estimation and Behavioral Interpretation

A set of explanatory variables are generated from the above datasets and then fed into the dynamic discrete choice model for parametric estimation. In this section, we base on the coefficient estimates to interpret drivers' factorial focuses underlying the customer-searching movements.

4.1 Parametric setup

The following explanatory variables are generated to represent the variable market condition:

- Trip fare TF_z^t: The average trip fare of orders originated in zone z during time interval t;
- Number of requests NR_z^t: The total number of requests sent out in zone z during time interval t;
- Pickup time PT_z^t: The average pickup time of passengers requesting for trips in zone z during time interval t;

- Matching probability MP_z^t: The ratio of idle drivers receiving matches in zone z during time interval t.

To be consistent with the movement threshold κ, the time granularity for the four market variables are set as 4 minutes. The first three variables are further normalized to distributions with zero mean and unit standard deviation. It is also verified that the partial correlation among the explanatory variables is fairly low, and collinearity should not be a concern here. All these variables are then invited to construct the observed utility $v\left(z_{t+1}|z_t\right)$ in Eq. (2) as follows[1],

$$v\left(z_{t+1}|z_t,\theta\right) = \theta_{TF}TF_{z_{t+1}}^t + \theta_{NR}NR_{z_{t+1}}^t + \theta_{PT}PT_{z_{t+1}}^t + \theta_{MP}MP_{z_{t+1}}^t + \theta_{SZ}\mathbb{1}(z_{t+1} = z_t)$$

where coefficient $\boldsymbol{\theta}$ represents drivers' sensitivity to different factors. Specifically, θ_{SZ} is a movement constant that indicates drivers' preference for remaining motionless when idle; $\mathbb{1}(\cdot)$ characterizes an indication function that values 1 when the condition within the bracket holds and 0 otherwise.

To better capture the strategic movements of idle drivers, we introduce another variable —distance to the hotspot DH_z^h —that calculates the Euclidean distance from the centroid of zone z, where the driver stays, to that of the hotspot h. The hotspot areas are carefully selected in each of the five periods of a day, i.e. morning (6AM-10AM), daytime (10AM-4PM), evening (4PM-8PM), night (8PM-11PM), and midnight (11PM-6AM). For each zone, we sum up the number of requests by 20-min intervals and then calculate the 80th-percentile and the maximum counts in each separate period. A zone is then identified as a hotspot in the period if the 80th-percentile count falls below the maximum by less than 10 percent. The selected hotspots are mostly downtown areas, commercial areas in the suburb, and railway stations, which are further categorized into the downtown-hotspot area (DTH) and the outer-hotspot area (OUH). Our cross-nested logit model treats these two hotspot categories separately in the upper nest, and the variable DH_z^h calculates the distance to the closest hotspot in each category. The allocation parameter $\alpha_{z_t,l}$ in Eq. (3) are specified as follows:

$$\alpha_{z_t,h} = \frac{\exp\left(\gamma_{DH}DH_{z_t}^h\right)}{S}, \quad \forall\, h \in \{\text{DTH, OUH}\}$$

$$\alpha_{z_t,c} = \frac{\exp\left(\gamma_{NR}NR_{z_t}^t\right)}{S}$$

$$S = \exp\left(\gamma_{DH}DH_{z_t}^{\text{DTH}}\right) + \exp\left(\gamma_{DH}DH_{z_t}^{\text{OUH}}\right) + \exp\left(\gamma_{NR}NR_{z_t}^t\right)$$

where $\alpha_{z_t,h}$ and $\alpha_{z_t,h}$ denote the weights for the target-specific repositioning and aimless-cruising movements, respectively; the coefficients γ represent a set of parameters that indicate drivers' sensitivity to the factors in different searching categories. We note that the factor specifications for allocation parameters above are refined through stepwise selections from a list of variables.

The choice set of drivers $C(z_t)$ at each state is extracted from the observations, and the scaling factor μ_n is set to 1. Meanwhile, the sample is split evenly into two for the purpose of learning and validation. We employ the Nelder–Mead method, one of the best-known algorithms for derivative-free optimization of unconstrained problems, to estimate the parameters.

[1]Note that many other factors, such as drivers' home location, parking availability, and traffic congestion etc., may dictate drivers' customer-search movements but are omitted in this study due to our data limitation. Once available, they are advised to be incorporated to reduce the potential omitted variable bias in result interpretation.

Table 1: Model comparisons for different evaluation horizons T

Horizon T (step)	0 (static)	1	2	3	4	5
LL	-112,311	-112,235	-112,221	-112,220	-112,220	-112,220
LL (validation)	-109,803	-109,743	-109,726	-109,724	-109,724	-109,724

Table 2: Model comparisons for juxtaposed nest structures and granularities of ρ_{z_t}

Nest structure	Cross-nested			Nested	Plain
ρ_{z_t}'s updating frequency	4 min	2 h	Never	4 min	4 min
LL	-112,221	-112,266	-112,261	-112,253	-112,875
LL (validation)	-109,726	-109,758	-109,755	-109,761	-110,440

4.2 Model refinement

With the value functions specified, we then refine the selection of critical parameters and structures in the model, i.e., the horizon of evaluation T, the updating frequency of the probability ρ_{z_t}, and the nest structure.

Table 1 compares the models under a set of different evaluation horizons T, based on the sample of full-time drivers. As shown in the table, the LLs in validation follow the same trend as that from model estimations, relieving the concerns for overfitting. The LL for the model when $T = 0$ (a static model) appears the lowest, implying the relative superiority of our proposed dynamic choice model. Besides, the LL increases constantly as T grows from 0 to 2, and stays constant afterwards. Considering the computational efficiency, we adopt $T = 2$ in the final implementation to account for the ahead-planning behavior of drivers in customer search.

Table 2 compares the models with different updating frequencies of ρ_{z_t} and juxtaposed nest structures. First, for the cross-nested case, we present three models where ρ_{z_t} gets updated every 4 minutes, 2 hours, and never (by fixing ρ_{z_t} to 1), respectively. According to the LLs, the model with the most frequently updated ρ_{z_t} works the best. Then, we zoom out to compare the modeling effectiveness of different nest structures. The nested logit model keeps only the stay/leave options in the upper nest (therefore, with no crossings in the zonal alternatives), while the plain structure practices the most basic logit model. The utility specifications of both models include the hotspot variables DH_z^h by linear combinations. Not surprisingly, the cross-nested model outperforms the other two and is thus selected for our later analyses.

4.3 Full-time versus part-time drivers

Based on the cross-nested model, we apply the pooling test by Watson and Westin (1975) to examine whether full-time and part-time drivers in general behave differently in customer search. The restricted model applies a unified set of coefficients for both types of drivers, while the alternative model specifies two separate sets of coefficients. The likelihood-ratio (LR) test statistic is significant at the 1% level, which confirms the heterogeneous searching behavior among different types of drivers. The specific parametric estimation as well as the significance of each estimate are presented in Table 3. The left two sets of columns respectively summarize the estimates of Θ for full-time and part-time drivers, while the right two columns present the two classes of drivers' comparative differences $\Delta\Theta$ in response to the various factors.

Table 3: Coefficient estimates of a cross-nested model with two types of drivers

Param. Est.	Full-time $\hat{\Theta}$		Part-time $\hat{\Theta}$		Difference $\Delta\hat{\Theta}$	
	Coef.	*t*-Stat.	Coef.	*t*-Stat.	Coef.	*t*-Stat.
Utility parameter						
θ_{TF}	0.02	2.40†	0.05	3.83$*$	0.03	1.77
θ_{NR}	0.08	37.02$*$	0.06	21.40$*$	-0.01	-3.77$*$
θ_{PT}	-0.02	-3.11$*$	0.08	7.91$*$	0.10	8.27$*$
θ_{MP}	1.00	18.29$*$	0.60	8.00$*$	-0.40	-4.30$*$
θ_{SZ}	1.70	199.62$*$	1.78	135.86$*$	0.09	5.45$*$
Allocation parameter						
γ_{DH}	-3.68	-8.00$*$	-3.17	-6.25$*$	0.51	0.74
γ_{NR}	0.52	8.46$*$	0.45	6.13$*$	-0.07	-0.78
Discount β	0.31	23.08$*$	0.17	7.46$*$	-0.14	-5.30$*$
Scaling μ_l	1.74	68.90$*$	1.91	39.22$*$	0.17	3.16$*$
Observations		166,171		83,778		249,188
LL(0)		-184,763		-93,047		-277,810
Final LL		-112,221		-53,042		-165,264
LL (validation)		-109,726		-52,217		-161,942
Adjusted ρ^2		0.39		0.43		0.41

Notes: The left two sets of columns show the coefficient estimates respectively for full-time and part-time drivers, while the rightmost two columns present the differences in drivers' response between the two classes of drivers (i.e. the part-time drivers' minus the full-time drivers' counterparts). The former results are estimated using separate datasets corresponding to each driver type, while the latter uses the full set of data.

† - Significance to the 5% level;

$*$ - Significance to the 1% level.

As shown by Table 3, most of the coefficient estimates $\hat{\Theta}$ are significant at the 1% level for both full-time and part-time drivers, with intuition-consistent signs. This implies an encouraging fact that in general ride-sourcing drivers do respond actively and positively to the repetitive market variations. The only counter-intuitive result is that $\hat{\theta}_{PT}$ is positive for part-time drivers. It means that part-time drivers tend to drive to the area with longer pickup time for passengers. In fact, different from the rest explanatory variables, the metric of passengers' pickup time PT_z^t is obscure to drivers, and the estimate $\hat{\theta}_{PT}$ may embody drivers' response to other market factors in relevant. For instance, the negative $\hat{\theta}_{PT}$ of full-time drivers may partially reflect their aversion to the congested areas in zonal choice, while the counter-intuitively positive $\hat{\theta}_{PT}$ for part-time drivers could be that they prefer the areas where they can match to customers in a wider space (higher matching opportunity but longer pickup time), or the areas with more condensed customer demand but also more congestion in usual. Besides, as revealed specifically by our model, drivers' searching movements are not confined to the local considerations. Instead, they show a clear tendency of repositioning towards the faraway hotspots, which strengthens significantly as they move closer to those areas. It is highly probable that drivers take the distant hotspots as back-up options given the certainty of receiving quick matches therein, similar to the inclination of taxi drivers for taxi stands outside the city center (Szeto et al., 2019). But as a result, all the supply of idle drivers at the neighborhood of a hotspot might be drained up, causing deceptive supply shortage in local regions.

While drivers of different groups hold consistent preferences for the various factors in customer search, there is a significant disparity between full-time and part-time drivers in response sensitivity. In contrast to full-time drivers, drivers work in part-time are much less sensitive to the number of requests and the matching probability in zonal search. This is consistent with our intuition that full-time drivers are more experienced in service provision and can thus respond scrupulously under different circumstances. Interestingly, part-time drivers characterize a significantly lower discount factor β compared to full-time drivers. In accordance with the role of β as a time-discounting factor, this implies that full-time drivers are more far-sighted by planning ahead, while part-time drivers focus more on the near-future opportunities. Meanwhile, the scaling parameter μ_l for full-time drivers is significantly lower, meaning that their

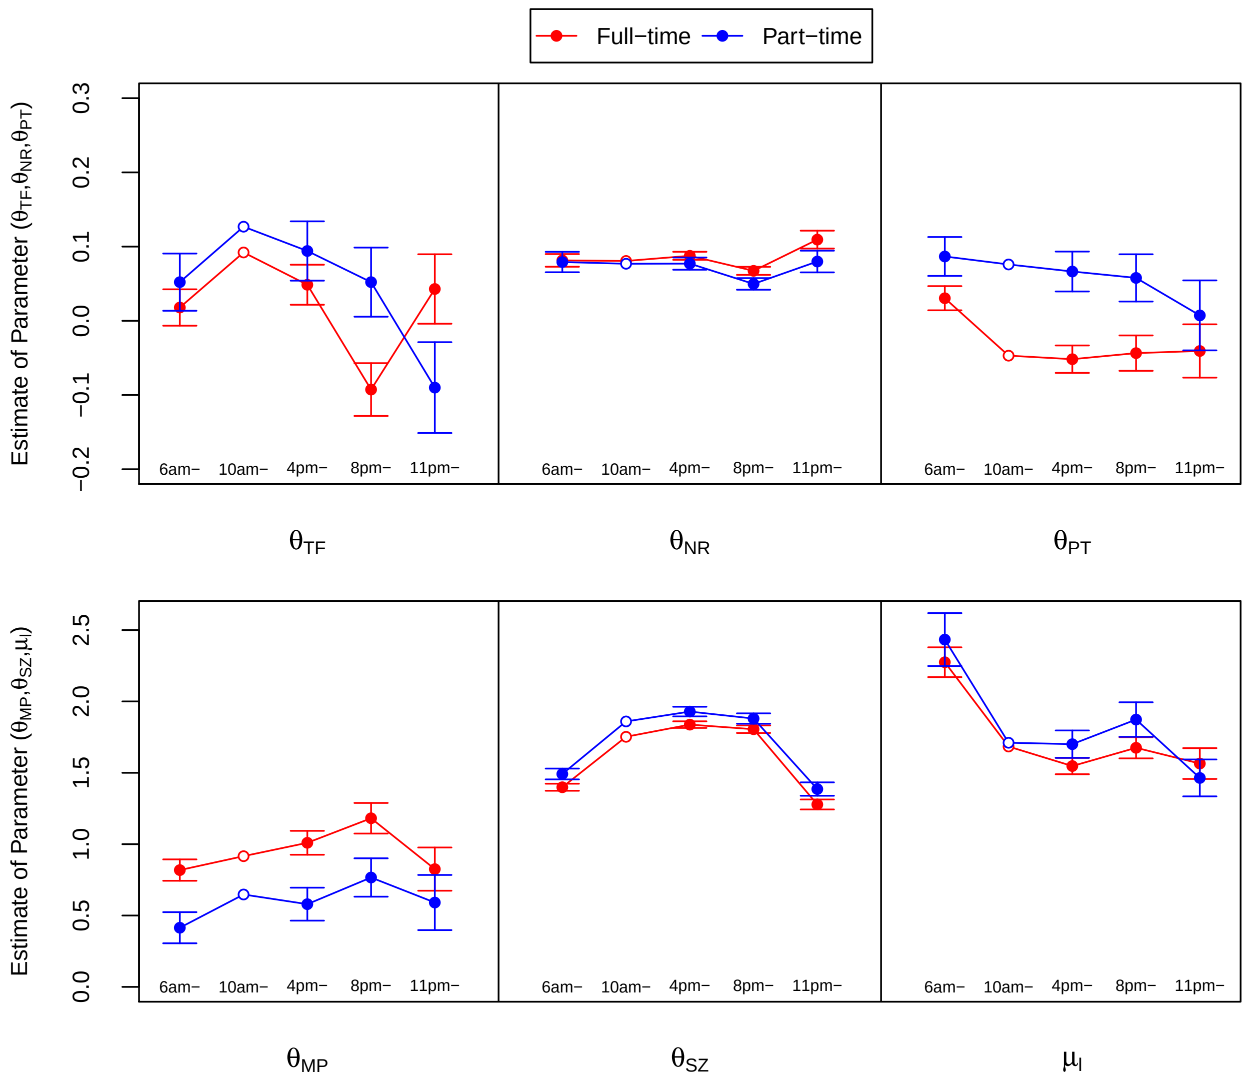

Figure 6: Comparisons of drivers' behavioral responses across the time of a day. On each parameter, the estimate $\hat{\Theta}$ corresponding to the daytime period (10:00 AM-4:00 PM) is set as the baseline, while the relative deviations $\Delta\hat{\Theta}$ of the other time periods are tested. The dots present the absolute value of estimates $\hat{\Theta}$ across different periods, while the error bars around indicate the standard deviations for $\Delta\hat{\Theta}$.

customer-searching movements are dictated more strongly by unobserved confounders.

We then proceed to examine whether drivers' searching behavior varies in time by allowing the coefficients to change by periods of time. To ensure the validity of calibration over the segmented datasets, we fix the discounting factor β and allocation factor γ in each sub-model to the value estimated previously using the full sample. Figure 6 displays the parametric estimates of full-time and part-time drivers across the time of a day, respectively. According to the figure, the two classes of drivers again exhibit very similar temporal patterns in response to the various factors. They both show higher preference on the trip fare during the daytime, while focusing more on the matching probability in the evening. Such a behavioral transition adheres with the nature of ride-sourcing markets, as the travel demand becomes much more sparse and heterogeneous spatially after the evening peak, when drivers need to switch searching goals to first secure the chances of getting matched. Meanwhile, the estimates of the movement coefficient θ_{SZ} indicate that drivers prefer to stay motionless in the afternoon and evening (specifically, from 10:00 AM to 11:00 PM), while moving more actively at late night and early morning. This contrast partially results from the fact that many drivers start their shifts before the morning peak and end at late night. During those periods, the searching behavior of drivers can be vastly influenced by their inclination to either reposition towards ideal service areas or move back home.

4.4 Space-time-dependent preference of searching movements

Applying the calibrated model above, we then derive the choice probabilities of full-time drivers at different locations and periods to further investigate how their latent behavior adapts to the variable market conditions. We note that full-time drivers constitute the majority of labor supply in the ride-sourcing service, characterizing a group comparable to the traditional taxi drivers. To yield comparable choice probabilities across time, we drop the outer-hotspot choices from the upper nest across the periods and then estimate the coefficients using models with a consistent upper nest.

Figure 7 and 8 visualize the zonal choice probabilities of full-time drivers at different areas in two typical periods to highlight the difference of drivers' latent searching movements between daytime and nighttime. Figure 7 first presents the probabilities for drivers in each zone to reposition to the downtown hotspot when being idle, with larger and darker pies marking the higher probabilities (same for the figures presented later in this section). As can be seen clearly, idle drivers, except for those at the few zones on the outskirt of the city, prefer less the choice of repositioning to the downtown hotspot during the daytime. In contrast, as the suburban market cools down significantly during the evening, idle drivers show much stronger willingness to reposition and escape the potentially long time of wait therein. Such insights are also suggested by Figure 8, which details the contrasts by visualizing the stepwise choice probabilities over the space. Each arrow denotes the choice of moving from the origin zone towards a neighboring zone, and again darker color indicates that the corresponding movement is chosen with a higher probability among all the choices available at the origin. Connecting all these preferences pictures the movement tendency of idle drivers within the spatial market. It can be easily observed that compared to the daytime (Figure 8a,b), idle drivers during the nighttime show much higher preferences for moving rather than staying motionless (Figure 8c,d). Meanwhile, the upward(downward) arrows in the south(north) side carry darker colors, which essentially implies the inclination of idle drivers moving to the central area to receive matches more easily. The pattern of drivers gathering from the suburban areas to the downtown is greatly strengthened at night. Similar behaviors were also observed and reported by Wong et al. (2014a, 2015) for taxi drivers in Hong

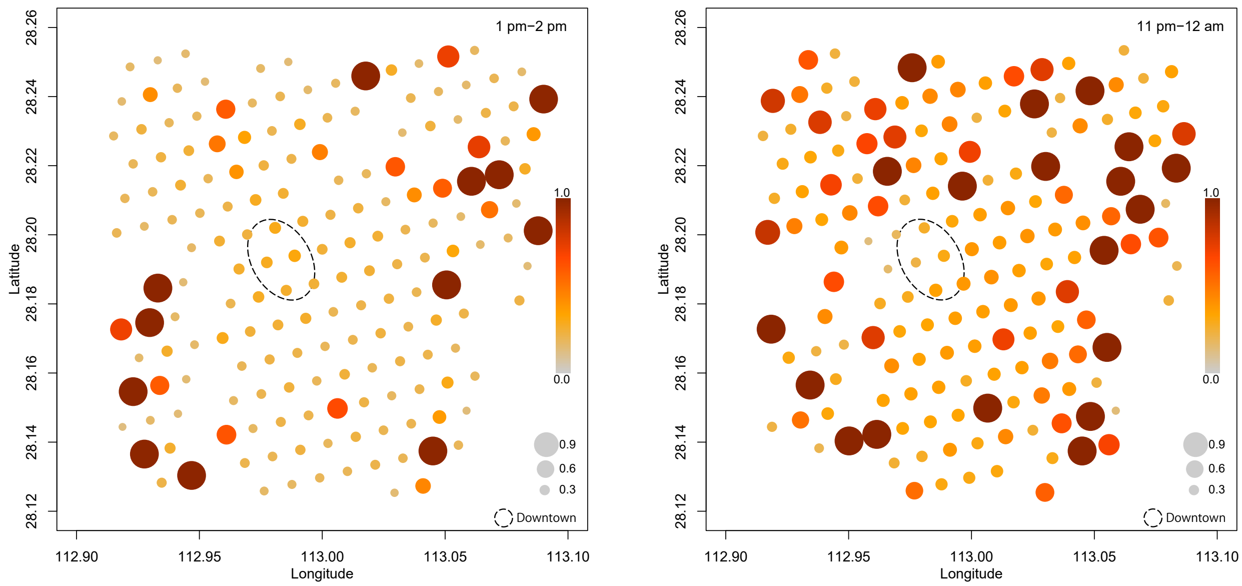

(a) Daytime

(b) Nighttime

Figure 7: Choice probability of idle drivers repositioning towards the downtown in each zone, with larger and darker pies representing the higher probabilities. The dotted circles in the center mark the downtown area of the city.

Kong.

Figure 9 and 10 display the choice probabilities of full-time drivers for staying motionless and cruising around the neighborhood, respectively, at four representative periods of a day. Interestingly, we find that drivers are consistently more mobile in the central areas compared to the suburbs. At those regions with rarer travel demand, ride-sourcing drivers prefer to stay motionless rather than moving and searching for customers, which somehow differs from that of taxi drivers. As per Wong et al. (2014a), taxi drivers do not show a clear preference for traveling towards taxi stands and waiting there for customers at the low-demand areas. In fact, we suspect that such an attitudinal difference between ride-sourcing and taxi drivers might be due to the nature of search frictions under the two ride-hailing modes. Taxi drivers mainly serve customers waving on the curbside and can thus improve their service efficiency substantially through local hunting (Zhang et al., 2014). In contrast, the app-based e-hailing services eliminate the physical barriers between drivers and passengers in matching, under which zonal search of drivers does not necessarily increase the chances of being matched but pushes up the operational costs, especially in places with sparse trip demand.

The temporal variations of choice probabilities are also intriguing. For example, the cruising behavior of drivers appears weakly in a wide range of areas during the morning peak (Figure 10a), and then almost disappears afterward (Figure 10b). Later, starting from the evening peak, while the cruising effect remains weak in the suburbs, it rebounds in the downtown as well as around the railway station (at the southeast corner) and intensifies in the evening up until midnight. Such a trend agrees well with those of taxi drivers, who are reported to be more willing to circulate within local regions during the morning peak but prefer to wait motionlessly for customers at the evening peak (Wong et al., 2015). Again, such time-varying behavior of drivers could also be a result of their strategic reaction to the changeable contrasts between

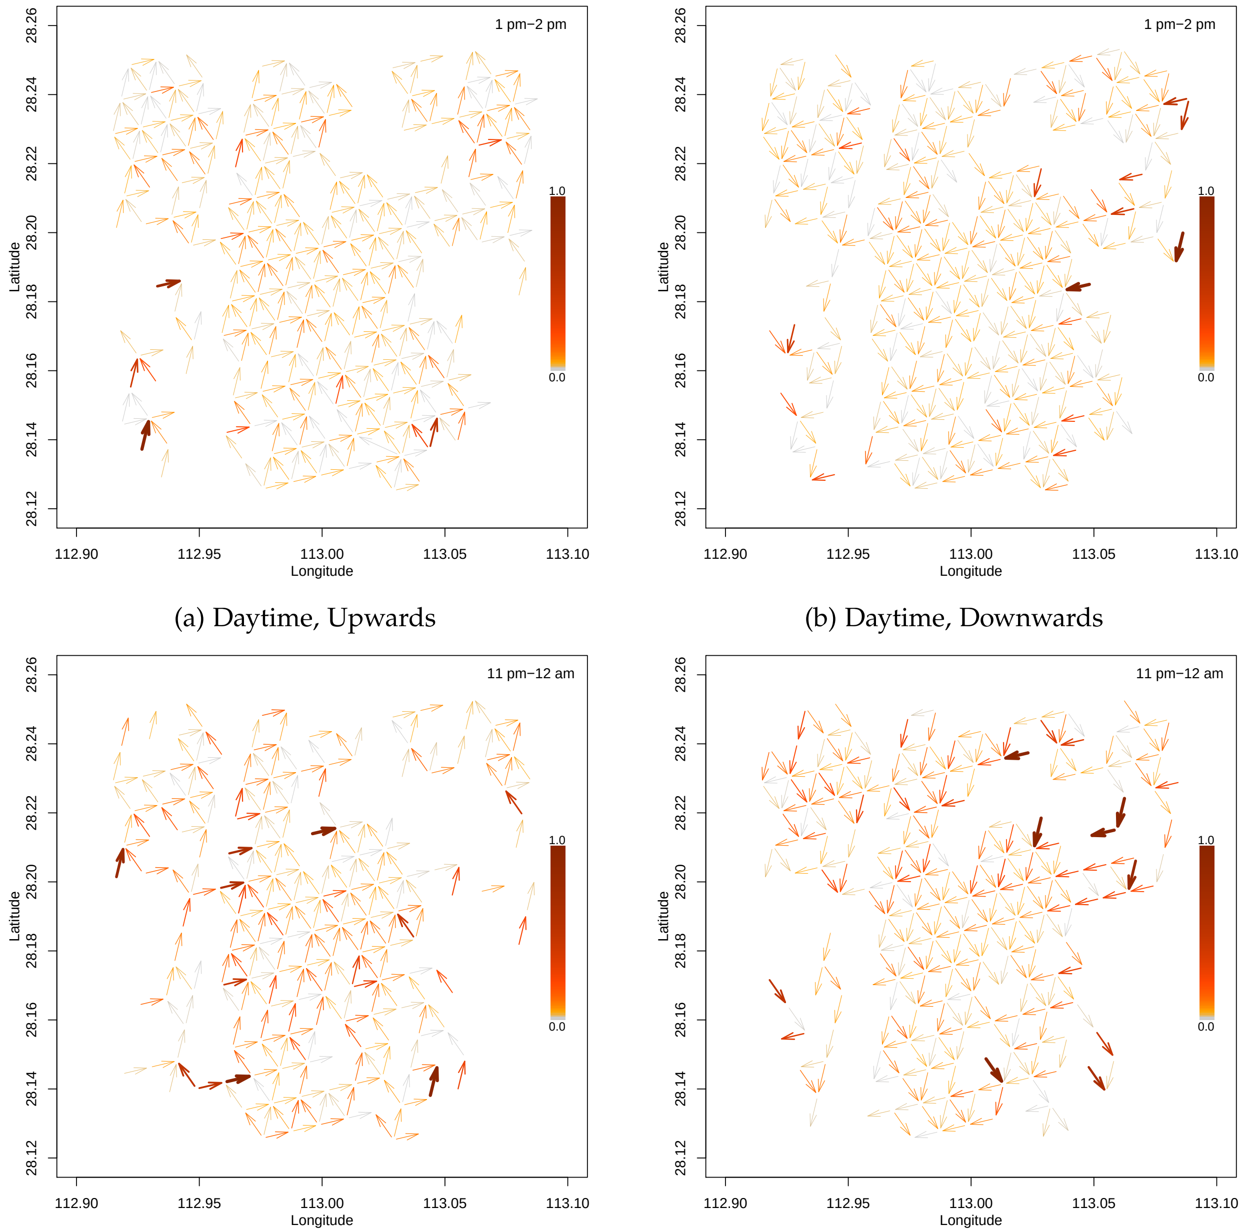

(a) Daytime, Upwards

(b) Daytime, Downwards

(c) Nighttime, Upwards

(d) Nighttime, Downwards

Figure 8: Choice probability of idle drivers on stepwise movements in zonal search during the a, b) daytime and c, d) nighttime. Compared to the daytime, idle drivers during the nighttime show a much higher tendency to move rather than stay motionless. Further, the pattern of drivers gathering from the suburban areas to the downtown is greatly strengthened at night. The upward(downward) arrows in the south(north) side carry darker colors compared to the north(south), which indicates drivers' inclination to move towards the central area.

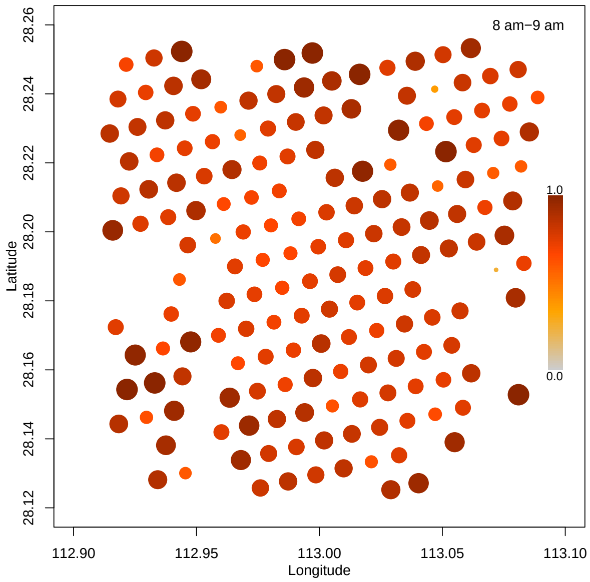

(a) Morning peak

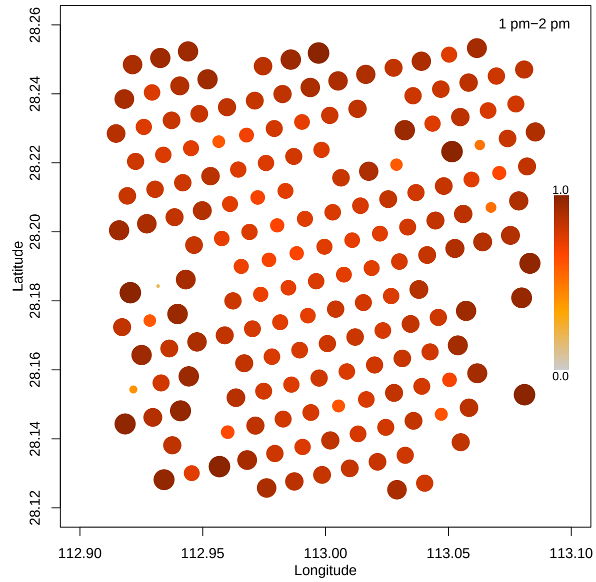

(b) Midday

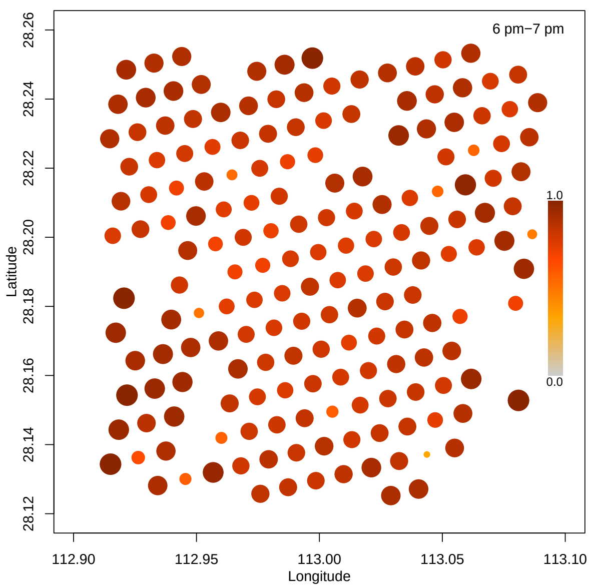

(c) Evening peak

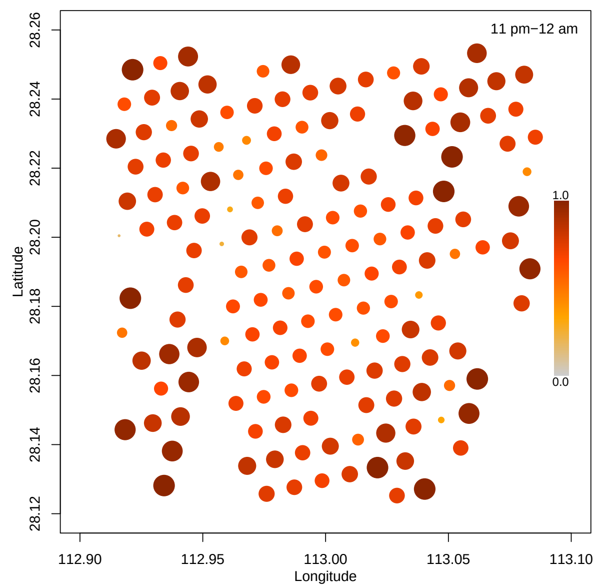

(d) Midnight

Figure 9: Choice probability of idle drivers staying motionless in each zone

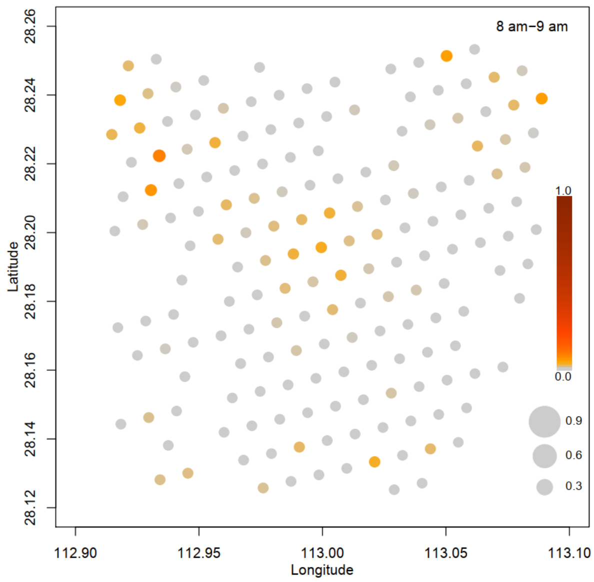

(a) Morning peak

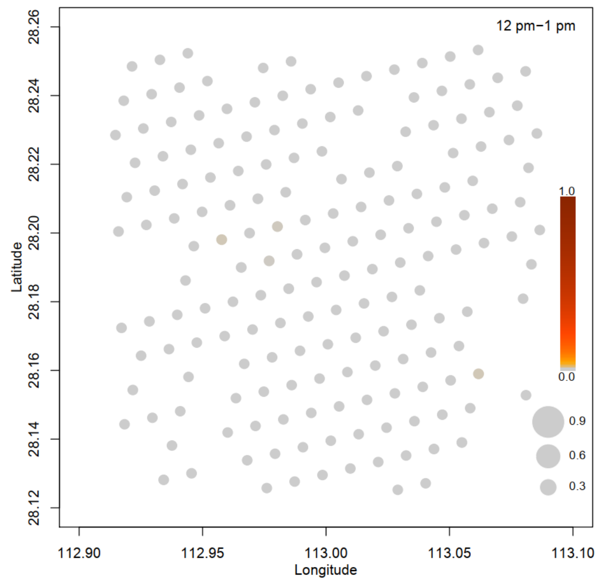

(b) Midday

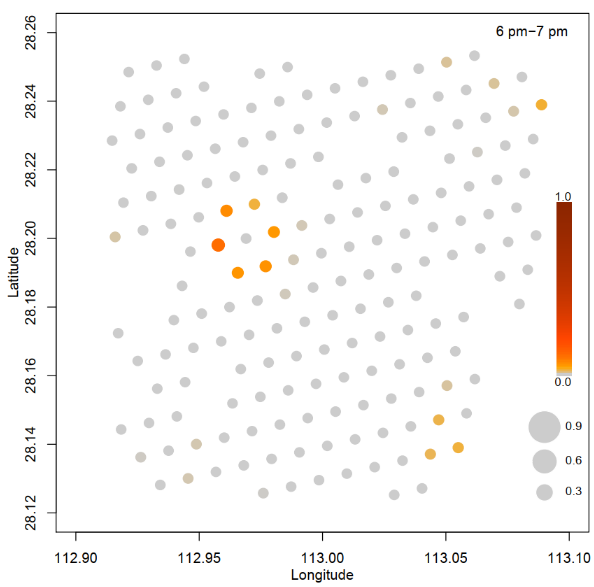

(c) Evening peak

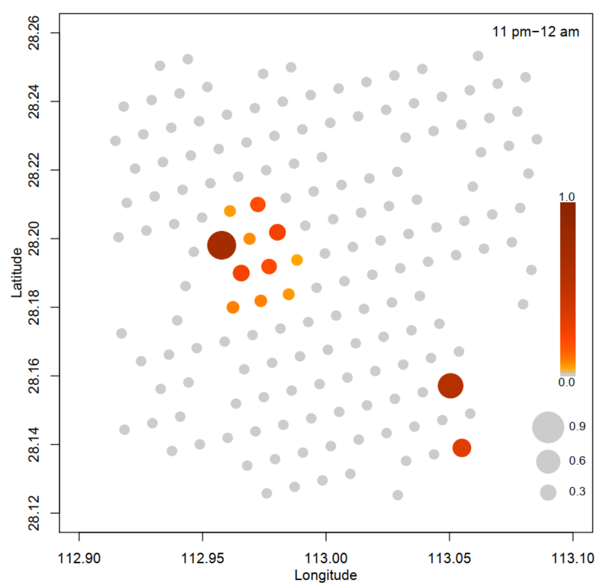

(d) Midnight

Figure 10: Choice probability of idle drivers cruising nearby in each zone. Note that the three pies in the right corner carry the sizes corresponding to the probability levels of 0.3, 0.6, and 0.9 for reference.

supply and demand at different time and locations in the market. For the idle drivers at hotspots, zigzagging in space may substantially lessen their matching time, given the relatively denser fleet supply and thus more fierce competition there (see Figure 4a). But the same strategy could lead to marginal improvements for ride-sourcing drivers in the cases with both demand and supply being sparse in space. It would be worthwhile to empirically verify these hypotheses in the future, to help reinforce the supply management of app-based ride-hailing services and yield more desirable guidance for idle drivers.

5 Conclusion and discussion

To the best of our knowledge, this paper is among the first attempts to investigate ride-sourcing drivers' customer-searching behavior. A dynamic discrete choice model has been proposed to rationalize the time-dependent search movements of idle drivers within the spatial market. The proposed model enables us to evaluate the impacts of spatiotemporal market conditions and understand the searching behavior of different classes of drivers. In particular, our model considers the unobservable intentions behind drivers' searching movements. Based on two large-scale datasets from real-world operations, we calibrate drivers' context-aware sensitivity to various factors in their decision-making when idle. Statistical testing results confirm that there exists a significant disparity between full-time and part-time drivers, and drivers' preferences in customer search vary across time and space. The supply management of ride-sourcing platforms could be further enhanced by accounting for these differentiated preferences of drivers:

- In general, ride-sourcing drivers respond actively to the repetitive market variations, with full-time drivers being more sensitive and far-sighted compared to part-time drivers. Platforms could thus customize searching guidance by accenting opportunities in nearby and broader spaces for part-time and full-time drivers, respectively.

- Catering better to drivers' time-dependent appetites, ride-hailing platforms need to vary the strategies for supply management. Drivers in idle prefer to stay motionless during the daytime but become significantly more mobile late at night actively seeking matching opportunities. Correspondingly, monetary incentives can be essential in stimulating idle drivers to reposition favorably in the day, while sharing information that helps reduce their idle time may be more welcomed when the market cools down at night.

- Drivers' aversion to the moving cost gives rise to their profound propensity to stay motionless when idle, especially in the suburbs where matching opportunities are scarce. It is thus difficult and costly to reposition idle drivers out of those less demanded areas. Once ended up there, drivers may be trapped with a long time of idleness. Such weak "self-adjustments" of idle drivers stress the importance of demand rationing for supply management. As customer trips deeply shape the supply availability in space, strategic pricing and matching that account for riders' destinations are critical for efficient circulation of supply resources.

- Customer-searching movements of drivers are not confined to local considerations. Instead, they show a clear tendency of repositioning towards faraway targets, and such inclination rises significantly as they stay closer to hotspot areas. As a result, compared to coldspots and hotspots, the supply at the middle ground can be relatively unsustainable. It may be drained up by hotspots, causing deceptive supply shortages. Hence, the platform may need to pay more attention to the areas surrounding hotspots to prevent overwhelmed supply rebalancing.

- On the whole, ride-sourcing drivers in the city do not cruise vigorously in local, as online matching overcomes the physical obstacles in customer search. Cruising behavior only gets intense isolatedly in the evening near the downtown and the railway station, where drivers strive to win over other competitors by moving inches closer to potential riders. The intense cruising in those few circumstances essentially signifies the mismatch of overall supply and demand therein. Appropriately, platforms should discourage drivers from dwelling in those oversupplied areas or adopt more transparent matching mechanisms to ease the fruitless competition.

The outcomes of this study suggest that multifaceted concerns/attitudes of drivers, other than regular market factors, can significantly dictate their customer-searching behavior. All these complexities and uncertainties of drivers in idle/searching movements pose challenges to the system operations. Therefore, to improve the supply management in the market, some ride-sourcing platforms have started recruiting contracted drivers, who are required to follow the platform's matching and repositioning instructions and paid with fixed income (Dong et al., 2020). One of the promising future topics is thus to investigate how to effectively utilize such a group of contractors and turn them into system actuators/controllers. Differentiated matching and repositioning of contractors could be effective in addressing the spatial imbalance of supply and demand, and substantially improve the efficiency of a ride-hailing system (Yang et al., 2020).

Acknowledgements

We would like to thank Dr. Pinghua Gong's team at Didi Chuxing for their professional and invaluable assistance in data access for this empirical research. The work described in this paper was partly supported by research grants from the US National Science Foundation (CMMI-1854684; CMMI-1904575), the Hong Kong Research Grants Council (No. HKUST16208619), the NSFC/RGC Joint Research Scheme under project N_HKUST627/18 (NSFC-RGC 71861167001), and Didi Chuxing.

References

Conway, M. W., Salon, D., and King, D. A. (2018). Trends in taxi use and the advent of ridehailing, 1995–2017: Evidence from the us national household travel survey. *Urban Science*, 2(3):79.

Daly, A. and Bierlaire, M. (2006). A general and operational representation of generalised extreme value models. *Transportation Research Part B: Methodological*, 40(4):285–305.

Dong, T., Xu, Z., Luo, Q., Yin, Y., Wang, J., and Ye, J. (2020). Optimal contract design for ride-sourcing services under dual sourcing. Available at SSRN: `https://ssrn.com/abstract=3658335`. Accessed on August 9, 2020.

Fosgerau, M., Frejinger, E., and Karlstrom, A. (2013). A link based network route choice model with unrestricted choice set. *Transportation Research Part B: Methodological*, 56:70–80.

Gao, Y., Jiang, D., and Xu, Y. (2018). Optimize taxi driving strategies based on reinforcement learning. *International Journal of Geographical Information Science*, 32(8):1677–1696.

Lin, K., Zhao, R., Xu, Z., and Zhou, J. (2018). Efficient large-scale fleet management via multi-agent deep reinforcement learning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1774–1783.

Liu, S., Araujo, M., Brunskill, E., Rossetti, R., Barros, J., and Krishnan, R. (2013). Understanding sequential decisions via inverse reinforcement learning. In *2013 IEEE 14th International Conference on Mobile Data Management*, volume 1, pages 177–186. IEEE.

Qu, M., Zhu, H., Liu, J., Liu, G., and Xiong, H. (2014). A cost-effective recommender system for taxi drivers. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 45–54.

Rong, H., Zhou, X., Yang, C., Shafiq, Z., and Liu, A. (2016). The rich and the poor: A markov decision process approach to optimizing taxi driver revenue efficiency. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 2329–2334.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033.

Rust, J. (2016). *Dynamic Programming*, pages 1–26. Palgrave Macmillan UK, London.

Sahr, K., White, D., and Kimerling, A. J. (2003). Geodesic discrete global grid systems. *Cartography and Geographic Information Science*, 30(2):121–134.

Shou, Z., Di, X., Ye, J., Zhu, H., Zhang, H., and Hampshire, R. (2020). Optimal passenger-seeking policies on e-hailing platforms using markov decision process and imitation learning. *Transportation Research Part C: Emerging Technologies*, 111:91–113.

Sun, H., Wang, H., and Wan, Z. (2019). Model and analysis of labor supply for ride-sharing platforms in the presence of sample self-selection and endogeneity. *Transportation Research Part B: Methodological*, 125:76–93.

Szeto, W., Wong, R., and Yang, W. (2019). Guiding vacant taxi drivers to demand locations by taxi-calling signals: A sequential binary logistic regression modeling approach and policy implications. *Transport Policy*, 76:100–110.

Tang, J., Wang, Y., Hao, W., Liu, F., Huang, H., and Wang, Y. (2019). A mixed path size logit-based taxi customer-search model considering spatio-temporal factors in route choice. *IEEE Transactions on Intelligent Transportation Systems*.

Train, K. E. (2009). *Discrete choice methods with simulation*. Cambridge university press.

Verma, T., Varakantham, P., Kraus, S., and Lau, H. C. (2017). Augmenting decisions of taxi drivers through reinforcement learning for improving revenues. In *Twenty-Seventh International Conference on Automated Planning and Scheduling*.

Wang, H. and Yang, H. (2019). Ridesourcing systems: A framework and review. *Transportation Research Part B: Methodological*, 129:122–155.

Watson, P. L. and Westin, R. B. (1975). Transferability of disaggregate mode choice models. *Regional Science and Urban Economics*, 5(2):227–249.

Wong, R., Szeto, W., and Wong, S. (2014a). Bi-level decisions of vacant taxi drivers traveling towards taxi stands in customer-search: Modeling methodology and policy implications. *Transport Policy*, 33:73–81.

Wong, R., Szeto, W., and Wong, S. (2014b). A cell-based logit-opportunity taxi customer-search model. *Transportation Research Part C: Emerging Technologies*, 48:84–96.

Wong, R. C., Szeto, W., and Wong, S. (2015). Sequential logit approach to modeling the customer-search decisions of taxi drivers. *Asian Transport Studies*, 3(4):398–415.

Xu, Z., AMC Vignon, D., Yin, Y., and Ye, J. (2020). An empirical study of the labor supply of ride-sourcing drivers. *Transportation Letters*, pages 1–4.

Yang, K., Tsao, M. W., Xu, X., and Pavone, M. (2020). Planning and operations of mixed fleets in mobility-on-demand systems. Available at arXiv: `https://arxiv.org/pdf/2008.08131.pdf`. Accessed on December 10, 2020.

Yu, J. J., Tang, C. S., Max Shen, Z.-J., and Chen, X. M. (2020). A balancing act of regulating on-demand ride services. *Management Science*, 66(7):2975–2992.

Zhang, D., Sun, L., Li, B., Chen, C., Pan, G., Li, S., and Wu, Z. (2014). Understanding taxi service strategies from taxi gps traces. *IEEE Transactions on Intelligent Transportation Systems*, 16(1):123–135.

Zheng, Z., Rasouli, S., and Timmermans, H. (2018). Modeling taxi driver anticipatory behavior. *Computers, Environment and Urban Systems*, 69:133–141.