Transportation Research Part B 150 (2021) 587–605

Contents lists available at ScienceDirect

Transportation Research Part B

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

A generalized fluid model of ride-hailing systems

Zhengtian Xu[a], Yafeng Yin[b,c,*], Xiuli Chao[c], Hongtu Zhu[d], Jieping Ye[d]

[a] Department of Civil and Environmental Engineering, The George Washington University, Washington DC, United States
[b] Department of Civil and Environmental Engineering, University of Michigan, Ann Arbor, United States
[c] Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, United States
[d] DiDi AI Labs, Didi Chuxing, Beijing, China

ARTICLE INFO

Article history:
Received 3 January 2021
Revised 14 May 2021
Accepted 23 May 2021

Keywords:
Ride-hailing
Fluid model
Spatiotemporal characteristics
Optimal pricing

ABSTRACT

This paper proposes a macroscopic fluid modeling framework to assist with strategic decision making of a platform for operating a large-scale on-demand app-based ride-hailing system. The framework captures the spatiotemporal characteristics of a ride-hailing system, and is flexible in representing control policies that a platform is implementing. It thus enables the analysis of a large-scale ride-hailing system with observable market responses and facilitates the optimization of the control polices. As a demonstration of the proposed framework, we customize it for a ride-hailing system operated by Didi Chuxing in a large city of China, and conduct an empirical study of the system. In doing so, we calibrate the model with a high degree of granularity using empirical data, and then investigate optimal pricing strategies to achieve various managerial objectives of the platform. To our best knowledge, the proposed model is the first of its kind that offers a tractable way to support the analysis and optimization of large-scale ride-hailing systems. The empirical study also provides one of the first testbeds for examining various operations strategies for ride-hailing services.

1. Introduction

The proliferation of smartphones has catalyzed rapid growth of app-based ride-hailing services offered by companies like Uber, Lyft and Didi Chuxing. When offering the service, these platforms have unprecedented controls over their ride-hailing systems, but finding optimal control policies is very challenging due to the complexities underlying the system dynamics. The operations and management of ride-hailing services has attracted extensive research interests covering various aspects of the services. See Wang and Yang (2019) for a recent, comprehensive review. This paper contributes to this ongoing quest by proposing a macroscopic fluid model to help with strategic decision making of a ride-hailing platform. The model describes the system dynamics of a spatial ride-hailing market under control polices by a platform, and can thus evaluate the impacts of these policies. It can also be applied to seek for better policies to further improve the system performance and achieve the managerial objectives of the platform.

The literature has relied excessively on canonical models that assume steady states in a ride-hailing system to derive various regulation, control, or policy to manage the system. A number of models are constructed based on highly simplified market/system settings, without or with stylized consideration of spatiotemporal variations (see e.g., Zha et al., 2016; Xu et al., 2017; Zha et al., 2018b; Yang et al., 2020b; Zhang et al., 2019). These models aim to offer profound theoretical or

* Corresponding author.
E-mail address: yafeng@umich.edu (Y. Yin).

https://doi.org/10.1016/j.trb.2021.05.014
0191-2615/

Nomenclature

Notation	Description
Sets	
$\mathcal{N}$	Set of regions
$\mathcal{T}$	Set of time period
$\mathcal{W}$	Set of OD pairs of ride-sourcing customer demands
$\mathcal{W}^c$	Set of bordered pairs of regions
Variables	
$a_{ij}^c(t)$	Cumulative number of type-(i, j) riders arriving at region i by time t
$m_{ij}(t)$	Cumulative number of type-(i, j) riders being matched at region i by time t
$\pi_{ij}^c m_{ij}(t, s)$	Number of type-(i, j) riders who enter at interval t while matched by time s
$q_{ij}^c(t)$	Number of type-(i, j) riders queueing at region i at time t
$w_{ij}^q(t)$	Time needed to match the number of type-(i, j) riders queueing at time t
$d_{ij}(t)$	Cumulative number of type-(i, j) riders delivered at destination j by time t
$w_{ij}^s(t)$	Service time for type-(i, j) riders who receive matches at time t
$a_i^d(t)$	Total arrival of idle drivers at region i by time t
$c_{ij}(t)$	Cumulative number of idle drivers that cruise from region i to k
$q_i^d(t)$	Number of idle drivers at region i at time t
$b_i(t)$	Cumulative number of drivers abandoning the system at region i by time t
$w_{ij}^m(t)$	Matching time for type-(i, j) riders who arrive at time t
$w_{ij}^p(t)$	Pickup time for type-(i, j) riders matched at time t
$\mathcal{P}_{ij}(t)$	Pricing strategy on type-(i, j) riders at time t
$\boldsymbol{\mathcal{A}}_i(t)$	Matching mechanisms at region i at time t
$\boldsymbol{\mathcal{R}}_i(t)$	Service interventions on idle drivers at region i at time t
Parameters	
$e_i(t)$	Cumulative external arrivals of idle drivers at region i by time t
$\alpha_i(t), \gamma_i(t)$	Coefficients of the pickup time function at region i in period t
$\eta_{ij}^s(t)$	Discounting ratio of ride-pooling for type-i, j riders in period t
$B_{ij}(t)$	Frequency of type-(i, j) users accessing the application in period t
$\boldsymbol{\beta}_{ij}(t)$	Type-(i, j) users' sensitivities on costs in period i
q_i^0	Threshold on the number of idle drivers in matching at region i
$\Lambda_{ij}^c(t)$	Probability of a driver repositioning from region i to j in period t
$\Lambda_i^b(t)$	Probability of a driver abandoning service at region i in period t
Functions	
$D_{ij}(t)$	Demand function of type-(i, j) riders at time t
$M_{ij}(\cdot)$	Aggregate matching function of type-(i, j) riders
$T_{ij}(t)$	Function on type-(i, j) riders' pickup time at time t
$C_{ij}(t)$	Function on drivers' rate of transition from region i to j at time t
$B_i(t)$	Function on drivers' rate of service abandonment at region i time t

policy insights, but fall short of supporting strategic decision making for operating a real ride-hailing system, especially those with drastic spatiotemporal heterogeneity. As demonstrated by Xu et al. (2020), the heterogeneity can cause severe operational inefficiencies. Since Yang and Wong (1998), a series of studies have been carried out to better describe the spatial structure of a street-hailing taxi market (see Salanova et al., 2011 for a comprehensive survey). Recent studies have also extended the analysis to app-based ride-hailing systems (see e.g., Ban et al., 2019; Xu et al., 2019 as well as the citations therein). Another school of work apply a queueing-network-based approach to better capture market frictions and examine how various operations strategies impact system performance (e.g., Braverman et al., 2019). A common assumption made by all these studies is that the market condition is stationary with no variations or correlations across periods. However, a real ride-hailing system is subject to time-varying supply and demand that interact over transient system states. Therefore, to better cope with the operations of real ride-hailing systems, we need models or approaches that accommodate the system dynamics and spatial heterogeneity. Notably, among a few exceptions in the literature, Ramezani and Nourinejad (2018) and Buchholz (2019) proposed dynamic network models that capture the spatial search and matching between passengers and street-hailing taxis. Buchholz (2019) focused on search frictions and investigated pricing policies to decrease the supply misallocation and increase the customer welfare, while Ramezani and Nourinejad (2018) stressed the inter-dependencies

between taxi and traffic systems and explored centralized dispatching of vacant taxis to improve the service efficiency. Comparatively, our model generalizes to the prevailing app-based ride-hailing services, where riders after receiving matches need to spend extra time waiting for drivers' pickup. Both online matching and physical pickup phases are methodically modulized in this paper, and then integrated endogenously with passenger demand to enable the analysis of operation policies with more targeted rider differentiation. In addition, Zha et al. (2018a), Chen et al. (2020), and Nourinejad and Ramezani (2020) introduced dynamic models that track the time-varying number of riders, vacant and occupied vehicles in a ride-hailing system. However, their model is constructed for an aggregate market without spatial heterogeneity.

Another approach that becomes increasingly popular is reinforcement learning, which has been adopted to assist ride-hailing platforms with online matching and dispatching policies (see e.g., Ke et al., 2019; Shou et al., 2020). These studies disentangle the system dynamics as Markov decision processes, and then estimate policy values via designated data-driven approaches. However, the underlying learning process is a "black box" to the model builders, which makes the resultant policies ineffective for generalization. This paper takes another route by proposing a generalized fluid modeling framework of ride-hailing systems constructed closely following its spatiotemporal dynamics. Managerial strategies of a platform are embedded mathematically into the framework and thus can be optimized systematically with observable market responses. The transparency of our framework's physical structure makes it highly transferable for customization and update. We demonstrate this by conducting an empirical study and customizing a model for a ride-hailing system in a large city in China. We calibrate the model with a high degree of granularity using a large-scale dataset from Didi Chuxing, the largest ride-hailing platform in China. It is demonstrated that the model is capable of describing the market conditions and predicting the consequences of control policies; moreover, it can be incorporated into an optimization model to prescribe optimal operation strategies. To our best knowledge, the proposed modeling framework is one of the first to offer a tractable way to support the analysis and optimization of real, large-scale ride-hailing systems. The empirical study also provides one of the first testbeds for examining various operations strategies for ride-sourcing services.

For the remainder, Section 2 first presents the generalized fluid-based modeling framework for ride-hailing systems by assembling two groups of relationships, respectively resulting from the conservation law and statistical characteristics in fluid dynamics. Section 3 introduces a model customized and calibrated for a real ride-hailing system in a large city in China. In Section 4, we demonstrate how the model can be applied to prescribe optimal control policies and discuss a solution procedure to solve the resulting optimization model. Section 5 then presents numerical results that show the potentials of price leverages for the ride-hailing platform in service management. Lastly, Section 6 concludes the paper.

2. Fluid-based system model

This section first introduces various conservation relationships held in a ride-hailing system, followed by discussions on its system dynamics. We then formulate a generalized fluid-based model for a ride-hailing system. We consider a city with $\mathcal{N}$ regions, where a ride-hailing platform manages a group of affiliated drivers to serve customers, each characterizing a ride from one region $i(\in \mathcal{N})$ to another $j(\in \mathcal{N})$. The dynamics of such a ride-hailing system can be largely captured by 10-tuple processes. Below details each process over a period of time $\mathcal{T}$, with a complete notational glossary provided in the nomenclature section for reference.

2.1. Conservation in system dynamics

Denote $\mathcal{W}$ as the set of origin-destination (OD) pairs for ride requests. For brevity, we refer type-(i, j) riders to those whose requests are from i to j with $(i, j) \in \mathcal{W}$. Then, the processes $\{a^c_{ij}(t), t \in \mathcal{T}\}$ and $\{m_{ij}(t), t \in \mathcal{T}\}$ are both $\mathcal{W}$-dimensional, where $a^c_{ij}(t)$ and $m_{ij}(t)$ respectively denote the cumulative number of type-(i, j) riders having arrived and been matched at region i by time t. The process in $\{q^c_{ij}(t), t \in \mathcal{T}\}$ represents the number of type-(i, j) riders at region i waiting to be matched at time t. Naturally, we have the following relationship among these three processes,

$$a^c_{ij}(t) = m_{ij}(t) + q^c_{ij}(t), \quad \forall (i, j) \in \mathcal{W} \tag{1}$$

In contrast to $\{q^c_{ij}(t)\}$, the processes $\{w^q_{ij}(t), t \in \mathcal{T}\}$ represent the frictions in serving ride requests. We thus define:

$$a^c_{ij}(t) = m_{ij}\big(t + w^q_{ij}(t)\big), \quad \forall (i, j) \in \mathcal{W} \tag{2}$$

We note that $w^q_{ij}(t)$ reflects the average time that type-(i, j) riders need to wait for matching around time t under a general matching algorithm.

Let each process in $\{d_{ij}(t), t \in \mathcal{T}\}$ be the cumulative number of type-(i, j) riders delivered to their destination j by time t. Define $w^s_{ij}(t)$ as the service time (pickup plus in-vehicle time) for type-(i, j) riders matched at time t. Accordingly, the cumulative numbers of matches and deliveries are connected as follows,

$$d_{ij}(t) = \sum_{\tau \in \mathcal{T}} \mathbb{1}\big(\tau + w^s_{ij}(\tau) \leqslant t\big) \cdot \delta m_{ij}(\tau), \quad \forall (i, j) \in \mathcal{W} \tag{3}$$

where $\delta m_{ij}(\tau)$ represents the number of type-(i, j) riders newly matched at time τ.

At each region, idle drivers are generated from three sources: external arrivals, dropping off riders at the region, or cruising from neighboring areas. Define the processes $\{e_i(t), t \in \mathcal{T}\}$ and $\{a_i^d(t), t \in \mathcal{T}\}$ respectively as the external and total arrivals of idle drivers at region $i \in \mathcal{N}$ by time t. Let $\mathcal{W}^c$ be the full set of pairs for neighbouring or adjacent regions, and define the processes $\{c_{kl}(t), t \in \mathcal{T}\}$ as the cumulative number of idle drivers that successfully transit from region k to l with $(k, l) \in \mathcal{W}^c$ and contribute to the supply at l by time t. Then, we obtain the following equation regarding the entry of idle drivers:

$$a_i^d(t) = \sum_{j:(j,i)\in\mathcal{W}} d_{ki}(t) + \sum_{k:(k,i)\in\mathcal{W}^c} c_{ki}(t) + e_i(t), \quad \forall i \in N \tag{4}$$

On the other hand, for drivers already in a region, they could either stay idle, be matched to riders, get offline, or transit to another region. Define the process in $\{q_i^d(t), t \in \mathcal{T}\}$ as the number of idle drivers at region i at time t, and each in $\{b_i(t), t \in \mathcal{T}\}$ to be the cumulative number of drivers getting offline at region i by time t. We thus have another conservation for drivers as follows:

$$a_i^d(t) = q_i^d(t) + \sum_{j:(i,j)\in\mathcal{W}} m_{ij}(t) + b_i(t) + \sum_{k:(i,k)\in\mathcal{W}^c} c_{ik}(t), \quad \forall i \in N \tag{5}$$

The above five conservation equations hold in general for non-pooling ride-hailing systems. They reveal the intrinsic interactions and dynamics of rider and driver flows that are isolated from market turbulence and system realizations. We will leverage these relationships to trace causal transmissions throughout the spatiotemporal network.

2.2. Statistical characteristics of system dynamics

Governed by the aforementioned conservation conditions, a ride-hailing system evolves with stochasticity and randomness, under the control policies by the platform. Thanks to the law of large numbers, below we capture statistical characteristics of the system dynamics.

At the demand side, the number of rides realized in the system is largely influenced by the service costs that riders experience at the stages of matching, pickup and delivery. The impacts can be summarized as the following latent function D_{ij} for each pair of $(i, j) \in \mathcal{W}$ during small interval δt around time t,

$$\mathbb{E}[\delta a_{ij}^c(t)]/\delta t = D_{ij}\Big(\mathbb{D}^t[w_{ij}^m(t)],\ \mathbb{D}^t[w_{ij}^p(t + w_{ij}^m(t))],\ \mathcal{P}_{ij}(t),\ t\Big) \tag{6}$$

where δa_{ij}^c represents the increment of type-(i, j) rider arrivals in δt, while $w_{ij}^m(\cdot)$ and $w_{ij}^p(\cdot)$ respectively denote arriving riders' matching and pickup time. The symbol $\mathbb{D}^t[\cdot]$ denotes the probabilistic distribution over time realized by the variable inside. The last term $\mathcal{P}_{ij}(t)$ denotes service fare of type-(i, j) riders at time t, which can be adjusted dynamically in accordance with the platform's pricing policy (Zha et al., 2018a; Bimpikis et al., 2019).

The online matching process performed by the ride-hailing platform can be generalized as the following many-to-one function $M_{ij}(\cdot)$ for each pair of $(i, j) \in \mathcal{W}$ and $t \in \mathcal{T}$, i.e.,

$$\mathbb{E}[\delta m_{ij}(t)]/\delta t = M_{ij}(\mathbb{D}^s[\boldsymbol{m_i}(\leqslant t)],\ \mathbb{D}^s[\boldsymbol{q_i^c}(\leqslant t)],\ \mathbb{D}^s[q_i^d(\leqslant t)],\ \boldsymbol{\mathcal{A}_i}(t)) \tag{7}$$

where $\boldsymbol{m_i}(\leqslant t)$, $\boldsymbol{q_i^c}(\leqslant t)$, and $q_i^d(\leqslant t)$ respectively indicate the historical realizations of processes $\{m_{ij}(t)\}$ and $\{q_{ij}^c(t)\}$ with $(i, j) \in \mathcal{W}$ and $\{q_i^d(t)\}$ at each region $i \in \mathcal{N}$ up to the time t. The action term $\boldsymbol{\mathcal{A}_i}(t)$ denotes the matching policy implemented by the platform at region i at time t. The uncertainty implied by the expression $\mathbb{D}^s[\cdot]$ results from the spatial distributions of riders and drivers, as the platform considers not only their arriving sequence/time waited but also the costs of meeting physically (Yang et al., 2020a).

As aforementioned, the service time of riders $w_{ij}^s(t)$ sums up the time in pickup and delivery, i.e.,

$$w_{ij}^s(t) = w_{ij}^p(t) + w_{ij}^l\Big(t + w_{ij}^p(t)\Big), \quad \forall (i, j) \in \mathcal{W} \tag{8}$$

where $w_{ij}^l(\tau)$ denotes the delivery time of the type-(i, j) rider who gets picked up at time τ. Since each region covers an area of space, pickup time $w_{ij}^p(t)$ and delivery time $w_{ij}^l(t)$ are both random. As an immediate result of matching, the statistics on pickup time $w_{ij}^p(t)$ are likely to rely on the same parametric setting as Eq. (7), i.e.,

$$\mathbb{E}\Big[w_{ij}^p(t)\Big] = T_{ij}\Big(\mathbb{D}^s[\boldsymbol{m_i}(\leqslant t)],\ \mathbb{D}^s[\boldsymbol{q_i^c}(\leqslant t)],\ \mathbb{D}^s\Big[q_i^d(\leqslant t)\Big],\ \mathcal{A}_i(t),\ t\Big), \quad \forall (i, j) \in \mathcal{W} \tag{9}$$

where $T_{ij}(t)$ characterizes the time-dependent function for type-(i, j) riders' expected pickup time that is variable to the real-time traffic condition in reality.

The states of a ride-hailing system also depend on drivers' decision-makings. For idle drivers, they can either wait in the current region or cruise to another one in search of the next rider. The aggregate movements of drivers can be measured by the following statistics,

$$\mathbb{E}\big[\delta c_{ij}(t)\big]/\delta t = C_{ij}\Big(q_i^d(t),\ \mathcal{R}_i(t),\ t\Big), \quad \forall (i, j) \in \mathcal{W}^c \tag{10}$$

$$\mathbb{E}[\delta b_i(t)]/\delta t = B_i\left(q_i^d(t),\ \mathcal{R}_i(t),\ t\right),\quad \forall i \in \mathcal{N} \tag{11}$$

where $C_{ij}(t)$ and $B_i(t)$ are two functions respectively related to drivers' expected rates of zonal transitions and service abandonment. The additional term $\boldsymbol{\mathcal{R}}_i(t)$ in the above functions stands for the intervention actions taken by the platform on drivers at region $i \in \mathcal{N}$ and time $t \in \mathcal{T}$.

2.3. Generalized fluid-based system model

The statistical relationships identified in the previous section reveal key physics or mechanisms of a ride-haling system. However, the involvement of random variables and statistical distributions makes it hard to connect with practical implementation. As a remedy, we apply fluid approximation. Fluid approximation first treats discrete riders and drivers as continuum and then takes the deterministic analog of system dynamics by omitting the uncertainty (Newell, 1971). Specifically for the processes previously discussed, the fluid approximation requires the removal of system uncertainty both in time and space (mathematically, by substituting the distributions of temporal and spatial processes with the corresponding mean values). Although the fluid approximation undermines to a certain level the accuracy of system state estimates, it yields a continuous and differentiable mathematical framework that provides mathematical tractability for system analysis and optimization. In this paper, we will demonstrate the descriptive capability of the resulting fluid model through an empirical study, but leave its qualitative properties for future investigation.

Applying the two approximation procedures to Eqs. 1–(11) yields the following generalized fluid-based model for ride-hailing systems:

$$a_{ij}^c(t) = m_{ij}(t) + q_{ij}^c(t),\quad \forall (i,j) \in \mathcal{W} \tag{12a}$$

$$a_{ij}^c(t) = m_{ij}\left(t + w_{ij}^q(t)\right),\quad \forall (i,j) \in \mathcal{W} \tag{12b}$$

$$a_i^d(t) = \sum_{j:(j,i)\in\mathcal{W}} d_{ki}(t) + \sum_{k:(k,i)\in\mathcal{W}^c} c_{ki}(t) + e_i(t),\quad \forall i \in \mathcal{N} \tag{12c}$$

$$a_i^d(t) = q_i^d(t) + \sum_{j:(i,j)\in\mathcal{W}} m_{ij}(t) + b_i(t) + \sum_{k:(i,k)\in\mathcal{W}^c} c_{ik}(t),\quad \forall i \in \mathcal{N} \tag{12d}$$

$$d_{ij}(t) = \int \mathbb{1}\left(\tau + w_{ij}^s(\tau) \leqslant t\right)\cdot \mathrm{d}m_{ij}(\tau),\quad \forall (i,j) \in \mathcal{W} \tag{12e}$$

$$w_{ij}^s(t) = w_{ij}^p(t) + w_{ij}^l\left(t + w_{ij}^p(t)\right),\quad \forall (i,j) \in \mathcal{W} \tag{12f}$$

$$w_{ij}^p(t) = T_{ij}\left(\boldsymbol{m}_i(\leqslant t),\ \boldsymbol{q}_i^c(\leqslant t),\ q_i^d(\leqslant t),\ \mathcal{A}_i(t),\ t\right),\quad \forall (i,j) \in \mathcal{W} \tag{12g}$$

$$\frac{\mathrm{d}a_{ij}^c(t)}{\mathrm{d}t} = D_{ij}\left(w_{ij}^q(t),\ w_{ij}^p\left(t + w_{ij}^q(t)\right),\ \mathcal{P}_{ij}(t),\ t\right),\quad \forall (i,j) \in \mathcal{W} \tag{12h}$$

$$\frac{\mathrm{d}m_{ij}(t)}{\mathrm{d}t} = M_{ij}\left(\boldsymbol{m}_i(\leqslant t),\ \boldsymbol{q}_i^c(\leqslant t),\ q_i^d(\leqslant t),\ \mathcal{A}_i(t)\right),\quad \forall (i,j) \in \mathcal{W} \tag{12i}$$

$$\frac{\mathrm{d}c_{ij}(t)}{\mathrm{d}t} = C_{ij}\left(q_i^d(t),\ \mathcal{R}_i(t),\ t\right),\quad \forall (i,j) \in \mathcal{W}^c \tag{12j}$$

$$\frac{\mathrm{d}b_i(t)}{\mathrm{d}t} = B_i\left(q_i^d(t),\ \mathcal{R}_i(t),\ t\right),\quad \forall i \in \mathcal{N} \tag{12k}$$

By specifying the functional forms for Eqs. (12 g-k) and the platform's policies in pricing $\boldsymbol{\mathcal{P}}$, matching $\boldsymbol{\mathcal{A}}$, and dispatching $\boldsymbol{\mathcal{R}}$, the dynamics Ω of a ride-hailing system, represented by the following 10-tuple processes, solves the system (12):

$$\Omega_{\mathcal{P},\mathcal{A},\mathcal{R}} \triangleq \left\{ \begin{matrix} a_{ij}^c(t),\ m_{ij}(t),\ q_{ij}^c(t),\ w_{ij}^q(t),\ d_{ij}(t) \\ e_i(t),\ a_i^d(t),\ q_i^d(t),\ c_{ik}(t),\ b_i(t) \end{matrix} \middle| \begin{matrix} \text{Subject to Eqs. (12) under } \mathcal{P},\mathcal{A},\mathcal{R} \\ \forall i \in \mathcal{N},\ (i,j) \in \mathcal{W},\ t \in \mathcal{T},\ (i,k) \in \mathcal{W}^c \end{matrix} \right\}$$

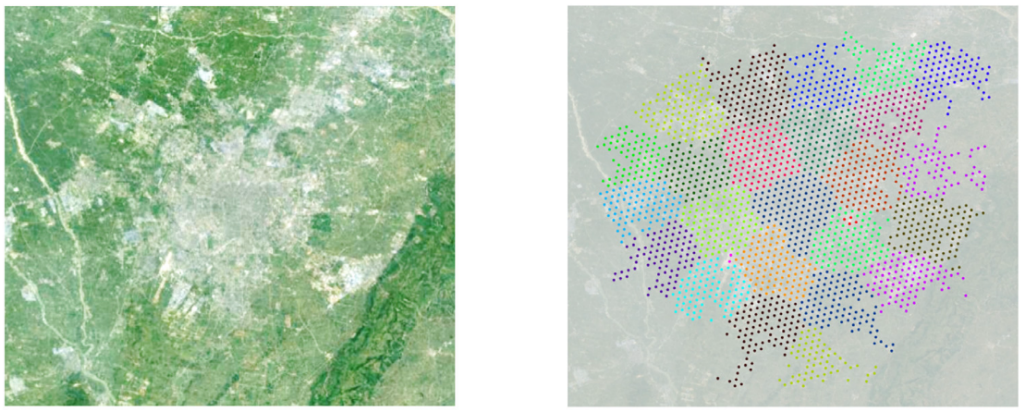

(a) (b)

Fig. 1. Maps of (a) the city and (b) its 27 clustered regions, including 24 spatial service areas and 3 service spots (2 railway stations and 1 airport).

The above fluid modeling framework allows us to quickly identify the system states and capture the interaction of time-varying demand and supply across the spatiotemporal network. It is very general and can be customized to represent various types of strategies of pricing, matching, or dispatching, by virtue of the controllers embedded in the fluid model, and then evaluate their impacts on system performance. Moreover, to further improve the performance, it is trivial to add another layer of optimization to prescribe optimal operation strategies, although it often requires non-trivial efforts to develop efficient solution algorithms. In the following sections, we exemplify how the proposed modeling framework can be customized and calibrated for a realistic system, and how the resulting model can be incorporated into an optimization model to prescribe optimal pricing policies for the system.

3. Parameterization, calibration, and validation of an empirical model

This section customizes the proposed generalized framework to model and analyze a real-world ride-hailing system. We consider the system operated by Didi Chuxing in a large city in China. The city is partitioned into 27 regions, consisting of 24 spatial areas (neighborhoods) and 3 service spots (two railway stations and one airport). See Fig. 1 for the spatial partitioning. Meanwhile, to facilitate the model calibration, we transform the above framework into a discrete-time counterpart and then parameterize the characteristic functions based on both empirical observations and physical insights. The empirical data we have spans over an eight-week period in Summer 2019 (from June 24, to August 11). A granularity of five-minute interval is chosen for model calibration given the highly variable market condition, especially during the peak periods.

Assume all the state transitions of drivers and riders happen only at the beginning of each time interval and the system conditions remain stationary within the interval. Catering to such a discrete-time context, we re-specify the label t to refer the time interval in $\mathcal{T}$ rather than a timestamp. The state variables (e.g. $q^c_{ij}(t)$ and $q^d_i(t)$) now denote the corresponding system condition at the beginning of interval t, while the prefix Δ in combination with the cumulative state variables (e.g., $a^c_{ij}(t)$ and $a^d_i(t)$) represent the respective increments during interval t. Below we then specify, calibrate, and validate different components of an empirical model with the contextual transitions.

3.1. Flow propagation and conservation

First of all, the flow conservation relationship can naturally be adapted as follows,

$$q^c_{ij}(t+1) = q^c_{ij}(t) + \Delta a^c_{ij}(t) - \Delta m^c_{ij}(t), \quad \forall (i,j) \in \mathcal{W} \tag{13a}$$

$$q^d_i(t+1) = q^d_i(t) + \Delta a^d_i(t) - \sum_{j:(i,j)\in\mathcal{W}} \Delta m^d_{ij}(t) - \Delta b_i(t) - \sum_{j:(i,j)\in\mathcal{W}^c} \Delta c_{ij}(t), \quad \forall i \in \mathcal{N} \tag{13b}$$

$$\Delta a^d_i(t) = \sum_{j:(j,i)\in\mathcal{W}} \Delta d_{ji}(t) + \sum_{j:(j,i)\in\mathcal{W}^c} \Delta c_{ji}(t) + \Delta e_i(t), \quad \forall i \in \mathcal{N} \tag{13c}$$

$$\Delta d_{ij}(t) = \sum_{\tau\in\mathcal{T}} \mathbb{1}\left(\tau + w^p_{ij}(\tau) + w^l_{ij}\left(\tau + w^p_{ij}(\tau)\right) = t\right) \cdot \Delta m^d_{ij}(\tau), \quad \forall (i,j) \in \mathcal{W} \tag{13d}$$

$$\Delta m^d_{ii}(t) = \eta^s_{ii}(t) \cdot \Delta m^c_{ii}(t), \quad \forall (i,j) \in \mathcal{W} \tag{13e}$$

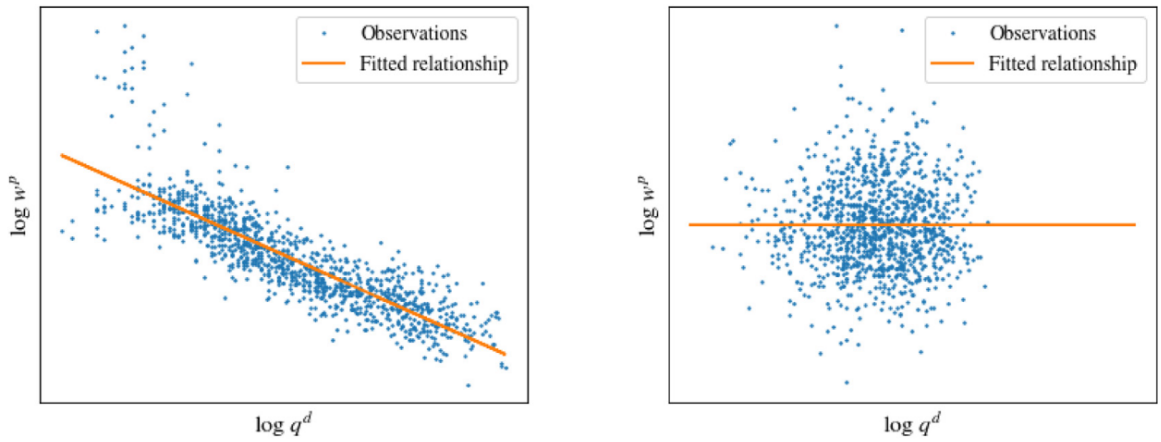

Fig. 2. Relationships of the average riders' pickup time versus the number of idle drivers at (a) the urban district and (b) the airport between 13:00 to 16:00 of weekdays. Logarithms have been taken for both indices for the purpose of visualization.

where Eqs. (13 a,b) respectively describe the updating of riders and drivers' queues; Eqs. (13 c,d) respectively inherit Eqs. (12 c,e) for each time slice, where the external arrivals of idle drivers $\Delta e_i(t)$ and the delivery time $w^l_{ij}(t)$ are both exogenous parameters directly retrieved/estimated from the empirical data; the last relationship (13 e) is newly added, with the matching tuple $\{m_{ij}\}$ now split by drivers $\{m^d_{ij}\}$ and riders $\{m^c_{ij}\}$ to accommodate the unbalanced matching quantities resulting from the ride-pooling service, which shares the same pool of drivers together with the solo service during operations. The parameter $\eta^s_{ij}(t)(\leq 1)$ characterizes an OD-specific discounting ratio for the service demand given the availability of ride-pooling, where multiple customers can be served by a single driver simultaneously. In our empirical analysis, each bundle of pooled trips that a driver uninterruptedly fulfilled is considered an "augmented" solo trip, which originates at the earliest boarding and terminates at the latest drop-off. The trip fare sums up over all the pooled trips, while the trip time comprises both the delivery and detours made for servicing pooled riders. The discounting ratio then equals the number of solo trips divided by the original number of requests.

3.2. Parameterization and calibration of input functions

To fully capture the system dynamics, we then parametrize and calibrate the five input functions (12 g-k) in the fluid model.

Pickup time function

The pickup time of riders who get matched in interval t is specified as a function of the instantaneous number of idle drivers within the region, i.e.,

$$w^p_i(t) = \alpha_i(t) \cdot \left[q^d_i(t)\right]^{\gamma_i(t)}, \quad \forall i \in \mathcal{N}, t \in \mathcal{T}$$

where parameters α_i and γ_i are time-dependent constants to be calibrated. Such a power-function form has been rationalized in theory by many previous studies (see e.g., Daganzo, 1978; Arnott, 1996 for pioneering derivations). Fig. 2 visualizes the relationships between the average riders' pickup time and the number of idle drivers in two example regions. Specifically, our calibration differentiates two types of service areas, neighborhood or spot. In neighborhoods, drivers are scattered spatially and matched by the platform to faraway riders who then require longer-distance pickups. Thus, the power index γ_i in this case is negative (see Fig. 2a), implying the physics that the average pickup distances of riders stay shorter under higher density of idle drivers. In contrast, for service spots like airport and railway station, which usually accumulate a large number of drivers waiting at designated areas, the pickup time of riders can be fairly constant (see Fig. 2b). In these areas, the power γ_i is set to be zero. We note that such a specification aligns reasonably well with empirical data.

Demand function

The demand of rides $\Delta a^c_{ij}(t)$ in each interval t is calibrated via the following specification,

$$\Delta a^c_{ij}(t) = B_{ij}(t) \cdot \exp\left(\beta^q_{ij}(t)\overline{w}^q_{ij}(t) + \beta^p_{ij}(t)\overline{w}^p_{ij}(t) + \beta^f_{ij}(t)\mathcal{P}_{ij}(t)\right), \quad \forall (i,j) \in \mathcal{W}, t \in \mathcal{T} \tag{14}$$

where coefficient $B_{ij}(t)$ denotes the potential demand for the ride-hailing service during interval t and the vector $\boldsymbol{\beta}$ represents users' sensitivity to costs in different service stages, i.e., online matching ($\boldsymbol{\beta^q}$), pickup ($\boldsymbol{\beta^p}$), and delivery ($\boldsymbol{\beta^f}$). Conceptually, each of the latter exponents represents the successful conversion rate of users after they sense the costs in waiting

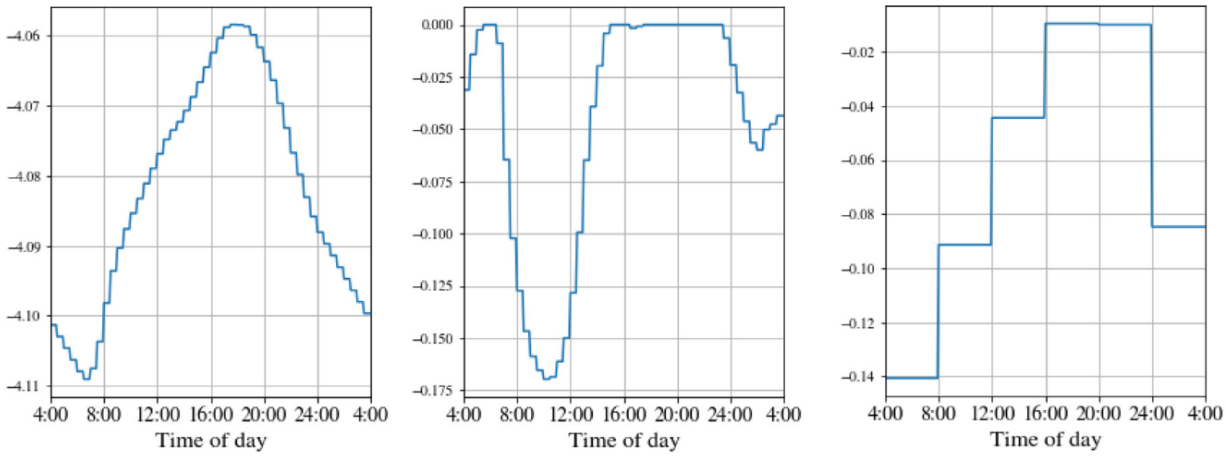

(a) β^q (b) β^p (c) β^f

Fig. 3. The sensitivities of demand in the central district on the matching time, pickup time, and fare for a typical weekday. The time sensitivities are calibrated with a 30-min granular, while the price sensitivity updates every four hours. The matching and pickup time is processed in the unit of hour, while the price is in Chinese Yuan (¥).

and price. In our calibration, the average trip distance and trip time are also incorporated to control for their possible impacts on the realized demand. The time and price sensitivities are calibrated altogether across different period of a day but with different granularity. An optimization is constructed to perform the least square estimation with side-constraints restricting the estimates to stay in the negative regime. Fig. 3 presents a set of calibrated sensitivities for the demand at the central district in a typical weekday. Interestingly, the calibration suggests a huge contrast between the morning and evening peaks. The demand is much more sensitive to additional waiting or price surge during the morning rush hours, but becomes neutral to the pickup time and price during the evening peak.

Note that $\bar{w}_{ij}^q(t)$ and $\bar{w}_{ij}^p(t)$ in the original demand function (14) are the average matching time and pickup time experienced by the riders who send out request at time t. In the discrete-time context, the two items are correspondingly calculated as

$$\overline{w}_{ij}^q(t) = \frac{\sum_{\{s\in\mathcal{T}|s\geq t\}} \pi_{ij}^c(t,s)\cdot(s-t)}{\sum_{s\in\mathcal{T}} \pi_{ij}^c(t,s)}, \quad \forall (i,j)\in\mathcal{W}, t\in\mathcal{T} \tag{15a}$$

$$\overline{w}_{ij}^p(t) = \frac{\sum_{\{s\in\mathcal{T}|s\geq t\}} \pi_{ij}^c(t,s)\cdot w_i^p(s)}{\sum_{s\in\mathcal{T}} \pi_{ij}^c(t,s)}, \quad \forall (i,j)\in\mathcal{W}, t\in\mathcal{T} \tag{15b}$$

where $\pi_{ij}^c(t,s)$ denotes the number of type-(i,j) riders who enter the system at interval t while getting matched at interval $s(\geq t)$. The determination of $\pi_{ij}^c(t,s)$ will depend on the specific matching disciplines adopted by the platform, but is always guarded by the following two relationships,

$$\begin{aligned} \Delta m_{ij}^c(s) &= \textstyle\sum_{t\in\mathcal{T}} \pi_{ij}^c(t,s), \qquad \forall (i,j)\in\mathcal{W}, s\in\mathcal{T} \\ \Delta a_{ij}^c(t) &= \textstyle\sum_{s\in\mathcal{T}} \pi_{ij}^c(t,s), \qquad \forall (i,j)\in\mathcal{W}, t\in\mathcal{T} \end{aligned}$$

In the later empirical analysis, we assume that riders of each OD type are matched with the first-come-first-serve (FCFS) discipline to pin down $\{\pi_{ij}^c(t,s)\}$.

Further, as implied by Eq. (15), both cost terms of riders are indeed realized endogenously. In other words, the market is dynamically equilibrated over time, where the demand rate of riders is realized based on the future supply condition.

Aggregate matching function

As evidenced by Fig. 4, the numbers of riders and drivers waiting for matching in a neighborhood largely adhere to the relationship below,

$$\left(q_i^d(t) - q_i^0\right)\cdot \sum_{j:(i,j)\in\mathcal{W}} q_{ij}^c(t) = 0, \quad \forall i\in\mathcal{N}, t\in\mathcal{T}$$

where q_i^0 denotes a threshold on the number of idle drivers. The equation essentially summarizes the following matching regime: with a sufficient number of drivers (above the threshold that corresponds to the intersection of two dashed lines in Fig. 4), incoming riders receive immediate matches. Otherwise, to avoid devastating long pickup, the platform intends to maintain the number of drivers to be above the threshold. In this case, riders have to wait. The waiting riders with

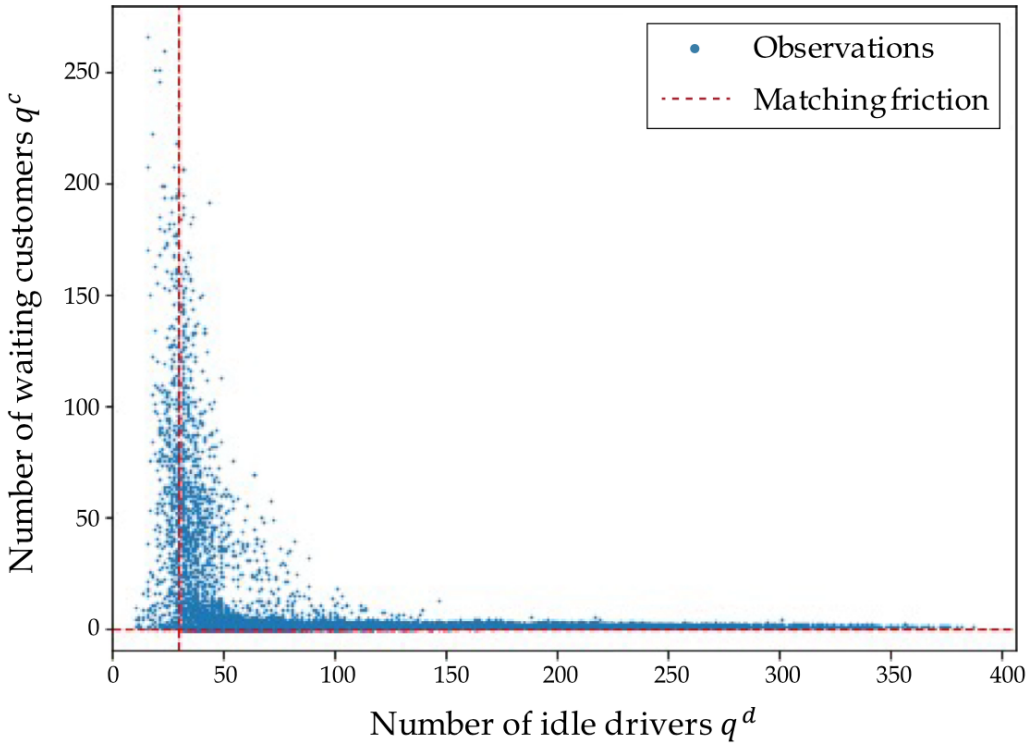

Fig. 4. Relationship between the numbers of riders and drivers waiting to be matched at an urban district. The auxiliary dashed lines rationalize the use of a threshold-based matching function to model matching frictions.

different destinations are then matched proportionally. It is worth noting that the materialized matching function considers each neighborhood separately, ignoring the possible inter-dependencies across the spatial market under batch matching. The consequences of such a simplification are marginalized in this empirical instance as we adopt relatively large zonal partitions, under which most matches take place within separate neighborhoods. However, if a higher spatial granularity is adopted, with cross-neighborhood matching becoming significant, an inter-zonal matching function needs to be calibrated as a replacement. See Xu et al. (2019) for such an attempt.

Enlighten by the aforementioned regime, the aggregate matching function under the discrete-time context is specified as the equation below,

$$\Delta m_{ij}^d(t) = \min\left(1,\ \frac{n_i^d(t) - q_i^0}{\sum_{k:(i,k)\in\mathcal{W}} \eta_{ik}^s(t) n_{ik}^c(t)}\right) \cdot \eta_{ij}^s(t) n_{ij}^c(t) \cdot \Big(1 + \mathcal{A}_{ij}(t)\Big), \quad \forall (i,j) \in \mathcal{W}, t \in \mathcal{T}$$

with the action $\boldsymbol{\mathcal{A}}(t)$ subject to the following two additional constraints,

$$\begin{aligned} &\textstyle\sum_{j:(i,j)\in\mathcal{W}} \eta_{ij}^s(t) n_{ij}^c(t) \cdot \mathcal{A}_{ij}(t) = 0 \\ &\Delta m_{ij}^d(t) \in \left[0,\ \eta_{ij}^s(t) n_{ij}^c(t)\right] \end{aligned}$$

where $n_i^d(t)$ and $n_{ij}^c(t)$ respectively denote the numbers of idle drivers and type-(i,j) riders participating in matching in region i at interval t. As per the equation, the platform in each interval either matches up all riders given sufficient supply or pairs them proportionately to maintain the supply level at the threshold. Specifically, the platform's action $\mathcal{A}_{ij}(t)$ is materialized as the weights/priorities in matching riders based on their destinations. In particular, if the platform does not prioritize riders according to their destinations (i.e., $\mathcal{A}_{ij}(t) = 0$), then all the riders waiting in the same area experience the same matching probability.

For our empirical analysis, the matching threshold $q_i^0(t)$ is determined to be the third percentile of the idle driver accumulation distribution for each region in each four-hour time interval of a day. This yields six thresholds for each region, each corresponding to one period of the day. Besides, the parameters $\boldsymbol{\mathcal{A}}(t)$ are preset to zero, assuming there was no matching differentiation under the status quo.

State transition rate

The last two functions are idle drivers' rate of regional transition and service abandonment, specified as the multiplicative form below for calibrations,

$$\begin{aligned} \Delta c_{ij}(t) &= q_i^d(t) \cdot \Lambda_{ij}^c(t) \cdot \Big(1 + \mathcal{R}_{ij}(t)\Big), \quad && \forall (i,j) \in \mathcal{W}^c \\ \Delta b_i(t) &= q_i^d(t) \cdot \Lambda_i^b(t), && \forall i \in \mathcal{N} \end{aligned}$$

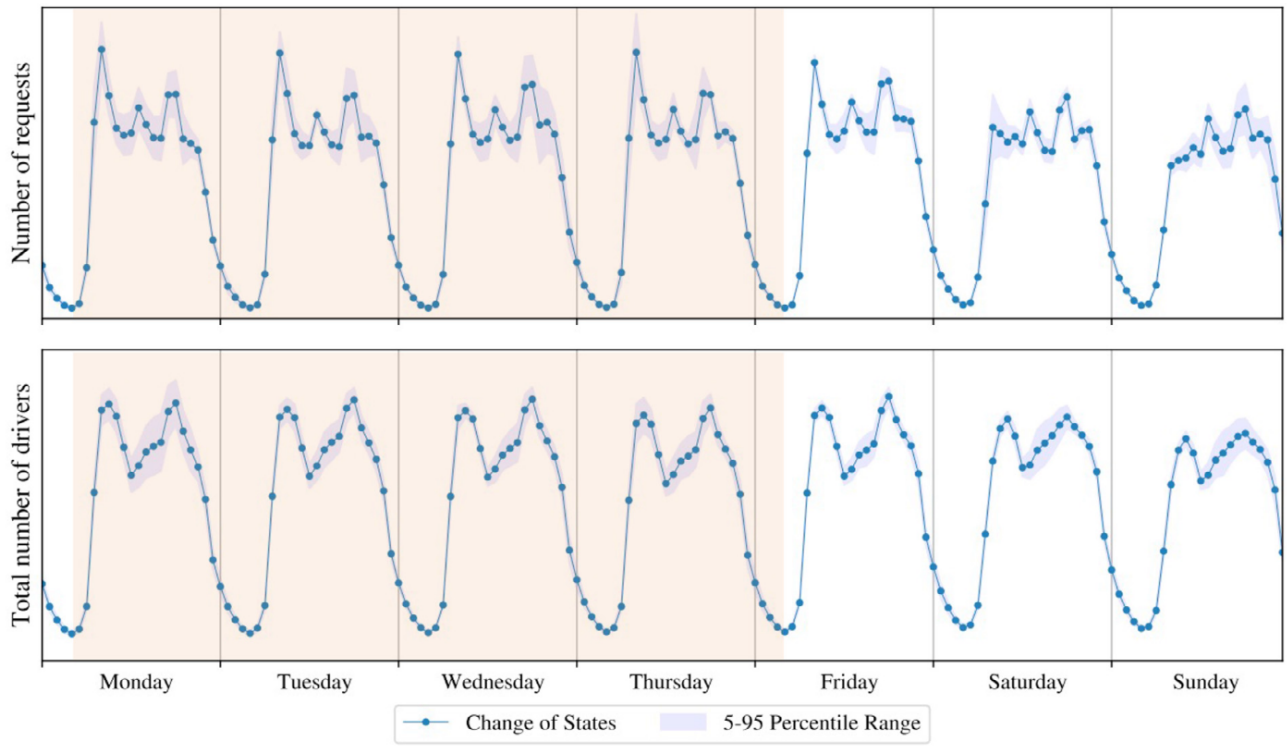

Fig. 5. Variation of supply and demand by hour of week, using the eight-week dataset from June 24 to August 11, 2019. The blue dotted lines record the mean values averaged over the corresponding metrics across all the five-min intervals within each hourly period, while the blue shaded area around denotes the 5th to 95th percentile range of respective distributions. The orange shaded area presents the target time period that comes into our empirical study, taking the weekdays from Monday 4:00 AM to Friday 4:00 AM. Note that the scalings of the two metrics are hidden for confidentiality. (For interpretation of the references to colors in this figure legend, the reader is referred to the web version of this article.)

where $\Lambda_{ij}^{c}(t)$ and $\Lambda_{i}^{b}(t)$ respectively represent the probability of an individual driver at region i repositioning to j and abandoning the service during period t. The platform's action of dispatching idle drivers $\mathcal{R}_{ij}(t)$ is represented as the influences posed on the probability of repositioning towards adjacent regions.

In this study, we assume that idle drivers' probabilities of zonal repositioning $\{\Lambda_{ij}^{c}(t)\}$ and service abandoning $\{\Lambda_{i}^{b}(t)\}$ are constants, unaffected by the changes of market condition. Consequently, the two components are calibrated by taking the sample average probabilities for each region in each period. However, we caution that in reality drivers do notice the real-time market condition and respond accordingly. Thus, the state-transition probabilities may need to be specified as the function of variable market conditions. Our framework is open for accommodating more sophisticated considerations of drivers' behaviors once the behaviors are modeled and calibrated in the future. See, e.g., Urata et al. (2020) for some initial effort.

3.3. Empirical validation of the parametric model

Summing up all the above specifications yields an empirical model customized for the ride-hailing system of interest. This section validates the model using the processed eight-week service dataset. We intend to examine how well the macroscopic model performs on describing a real system, given all the spatiotemporal discretization and the variety of approximations made.

Fig. 5 summarizes the weekly demand and supply pattern of the whole city across the eight weeks. The dotted blue lines show the average number of drivers in service (top) and requests being matched (bottom) respectively for five-min intervals in each hour of the week, while the surrounding shaded areas illustrate the 5th-95th percentile range of the corresponding distributions. As presented by the figure, both demand and supply exhibit well-shaped daily patterns with low scattering. The limited day-to-day variations suggest that after years of operations, the ride-sourcing service in this city has already reached a certain level of equilibrium. On the other hand, the substantial within-day variations in both state metrics call for a dynamic system model. To ensure the credibility of our calibrations and analysis, we choose to focus on the weekdays by narrowing the sample in each week to the period from Monday 4:00 AM to Friday 4:00 AM (as presented by the orange shared area in Fig. 5), resulting in a dataset of 32 repetitive service cycles. All the empirical analyses in the following are performed in such a daily setting that considers a typical service cycle starting at 4:00 AM of a weekday.

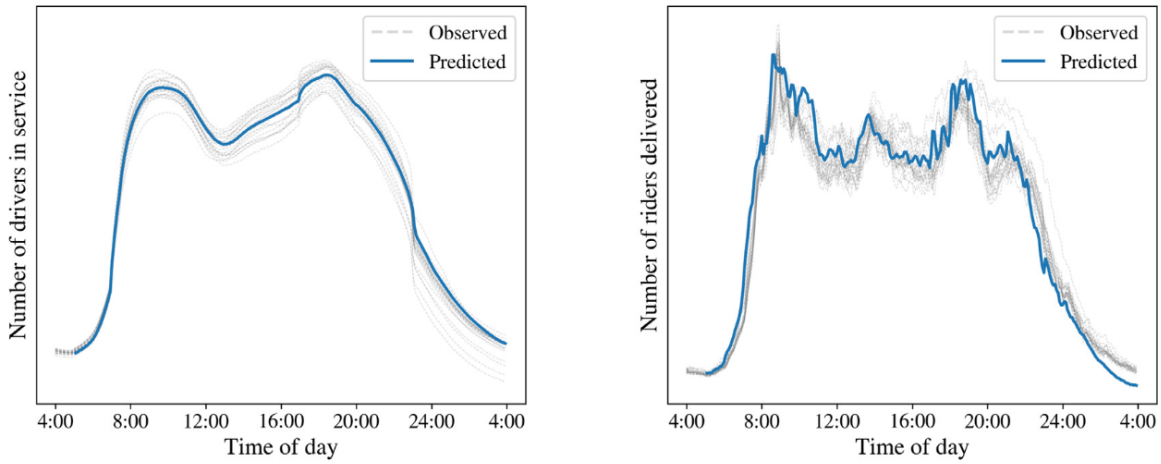

(a) Number of drivers in service (b) Number of riders delivered

Fig. 6. Comparisons of the predicted and observed system states on supply and demand. The blue solid line denotes the predicted values suggested by the empirical model, while the grey dotted lines show the real observations from the 32 daily samples.

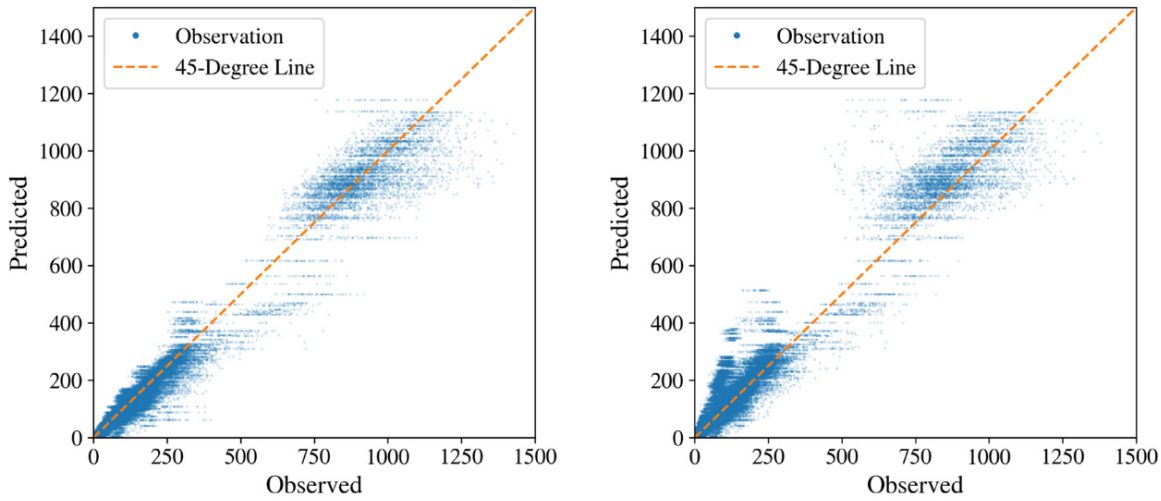

(a) Number of riders matched (b) Number of riders delivered

Fig. 7. Comparisons of the predicted and observed ridership by individual OD pairs in five-minute intervals. The two plots compare respectively the numbers of riders being (a) matched and (b) delivered on every single instance, with over 3 million instances in total. Note that each prediction corresponds to 32 observations across different days following the treatment we adopt for calibration. Accordingly, the horizontal spectrum at each level of predicted values is partly attributed to the market's natural fluctuations in addition to prediction errors.

Given each component calibrated following the above specifications, we now validate the proposed empirical model. We feed the currently realized trip demand into our calibrated spatiotemporal network and then compare the resultant system performance with the observed ones throughout the time of a day. We note that the fixed-demand scenario characterizes a special case that is easier to solve via a forward propagation scheme (see Section 4 for the details of the solution procedure). Fig. 6 presents the comparisons of the predictions (the blue solid line) from the empirical model and the 32 observations from the daily samples (the grey dotted lines). Two performance measures are compared, respectively being the number of active drivers in service (fleet supply) and the number of riders delivered to their destinations (system throughput) for each five-min interval. As can be clearly seen, the aggregate supply and demand predictions adhere fairly well with the observed counterparts. Meanwhile, Fig. 7 details the comparisons over disaggregate ridership upon individual OD pairs and time intervals. Each data point in the two plots connects the observed number of riders being matched and delivered, respectively, to its predicted value in each instance. The predictions again agree reasonably well with the observations on disaggregate service volumes, with the data points scattered closely beside the 45-degree line. Both comparisons on aggregate and disaggregate state measures validate our empirical model's capability to capture the macroscopic dynamics of a realistic system.

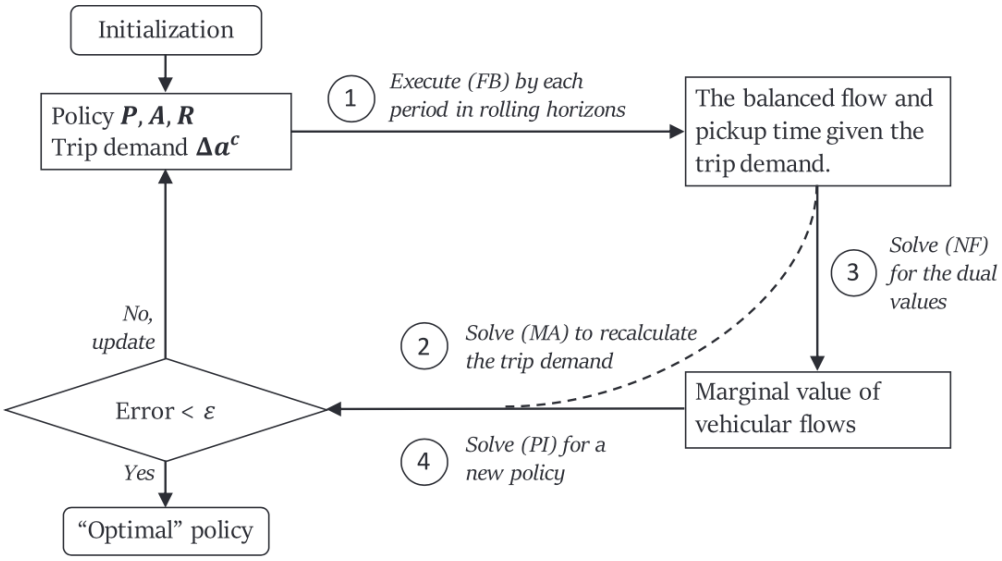

Fig. 8. Solution algorithm for policy optimization.

The counterfactual policy analysis in Section 5 will further demonstrate the utilization of this calibrated empirical model for prescribing optimal policies. Before that, we first introduce the solution procedure to solve the policy optimization problems formulated for various managerial actions.

4. Solution procedure for policy optimization

By virtue of the controllers embedded in the fluid model, it is conceptually straightforward to build optimization models over the system dynamics to seek policy improvements, i.e.,

$$\max_{\mathcal{P},\mathcal{A},\mathcal{R}} Z = \sum_{(i,j)\in\mathcal{W},\ t\in\mathcal{T}} \Phi_{ij}(t)\cdot m_{ij}(t) \quad \text{subject to the system equations (12)} \tag{16}$$

where a platform adjusts the policies in pricing $\mathcal{P}$, matching $\mathcal{A}$, and/or dispatching $\mathcal{R}$, to improve the system performance under certain managerial objective; and $\Phi_{ij}(t)$ denotes a measure of value taken for each request fulfilled. Note that such a multiplicative objective Z offers great flexibility and can be adapted to reflect a variety of managerial goals. For example, let $\Phi_{ij}(t) = 1$ reflects the goal of maximizing the system throughput, while prescribing $\Phi_{ij}(t)$ as $\mathcal{P}_{ij}(t)$ actually yields the case of a revenue-maximizing platform. The matching flows $\{m_{ij}(t)\}$ is derived subject to the relationships specified by the system dynamics. Simple as it appears though, the difficulty here is really about how to solve the resulting large-scale problem coupled with highly nonlinear constraints and possibly integer variables.

We thus devote this section to a heuristic algorithm that decomposes the whole problem into four subproblems or modules with each corresponding to a step in the proposed iterative procedure (see Fig. 8 for a graphical description of the procedure). Below we introduce these modules.

4.1. Mathematical modules in policy optimization

Four modules are presented, which are solved sequentially and iteratively to prescribe an optimal policy.

flow rebalancing module

Given the policy and trip demand, this module solves the bilateral matching and flow propagation problems in a forward propagation fashion. In other words, fixing the demand breaks the cyclic chains among the market conditions across periods, thereby enabling us to derive the flow propagation by rolling one-way through the spatiotemporal network. The following demonstrates the flow rebalancing (FB) process at interval t for any region $i \in \mathcal{N}$:

(FB) Initialize with given $q^c_{ij}(t)$, $q^d_i(t)$, and $\Delta d_{ij}(t)$

S1. Calculate the numbers of riders $n^c_{ij}(t)$ and drivers $n^d_i(t)$ participating in matching by

$$n^c_{ij}(t) = q^c_{ij}(t) + \Delta a^c_{ij}(t)$$

$$n_i^d(t) = q_i^d(t) + \sum_{j:(j,i)\in\mathcal{W}} \Delta d_{ji}(t) + \Delta e_i(t)$$

where the newly arrivals of riders $\Delta a_{ij}^c(t)$ and drivers $\Delta e_i(t)$ are supplied exogenously.

S2. Calculate the matching flows $\Delta m_{ij}^c(t)$ and $\Delta m_{ij}^d(t)$ as

$$\Delta m_{ij}^c(t) = \min\left(1,\ \frac{n_i^d(t) - q_i^0}{\sum_{k:(i,k)\in\mathcal{W}} \eta_{ik}^s(t) n_{ik}^c(t)}\right) \cdot n_{ij}^c(t) \cdot (1 + \mathcal{A}_{ij}(t))$$

$$\Delta m_{ij}^d(t) = \eta_{ij}^s(t) \cdot \Delta m_{ij}^c(t)$$

and update the queue length $q_{ij}^c(t+1)$ and $\hat{q}_i^d(t)$ through

$$q_{ij}^c(t+1) = n_{ij}^c(t) - \Delta m_{ij}^c(t)$$

$$\hat{q}_i^d(t) = n_i^d(t) - \sum_{j:(i,j)\in\mathcal{W}} \Delta m_{ij}^d(t)$$

where $\hat{q}_i^d(t)$ is an intermediate state created for the driver's queue.

S3. Calculate the zonal transition flows $\Delta c_{ij}(t)$ and then update the queue $q_i^d(t+1)$ by

$$\Delta c_{ij}(t) = \hat{q}_i^d(t) \cdot \Lambda_{ij}^c(t)$$

$$q_i^d(t+1) = \left(1 - \Lambda_t^b(t)\right) \cdot \hat{q}_i^d(t) - \sum_{j:(i,j)\in\mathcal{W}^c} \Delta c_{ij}(t) + \sum_{j:(j,i)\in\mathcal{W}^c} \Delta c_{ji}(t)$$

S4. Update the pickup time $w_{ij}^p(t)$ and the future deliveries $\Delta d_{ij}(s)$ via

$$w_i^p(t) = \alpha_i(t) \cdot [q_i^d(t)]^{\gamma_i(t)}$$

$$\Delta d_{ij}(s) = \Delta d_{ij}(s) + \mathbb{1}\left(t + w_i^p(t) + w_{ij}^l(t + w_i^p(t)) = s\right) \cdot \Delta m_{ij}^d(t), \qquad \forall s \geqslant t$$

S5. Transfer $q_{ij}^c(t+1)$, $q_i^d(t+1)$, and $\Delta d_{ij}(t+1)$ to the next period.

The above procedure can be implemented repetitively across space and time to recover the balanced flows of riders and drivers under a given trip demand and policy.

Additionally, based on the balanced flows, we define the service ratio $\bar{S}_{ij}^c(t)$ of each type of riders at each interval t as,

$$\bar{S}_{ij}^c(t) = \min\left(1,\ \frac{n_i^d(t) - q_i^0}{\sum_{k:(i,k)\in\mathcal{W}} \eta_{ik}^s(t) n_{ik}^c(t)}\right) \cdot (1 + \mathcal{A}_{ij}(t))$$

Also, the number of drivers available for matching $\bar{U}_i^d(t)$ is defined as

$$\bar{U}_i^d(t) = n_i^d(t) - q_i^0$$

Both notions as well as the balanced flows will be fed into the later module to produce the marginal values for various flows.

Matching assignment module

The first module determines the number of riders matched in each interval, i.e. $\{m_{ij}^c(t)\}$. But as mentioned, to fully recover the trip demand, we need the average waiting and pickup time of incoming riders, which further require the more detailed flow assignment $\{\pi_{ij}^c(t,s)\}$, i.e., the number of type-(i,j) riders who request for trips at period t and get matched at period s.

Assume that riders of each OD type are matched with the FCFS discipline. Then, the assignment $\{\pi_{ij}^c(t,s)\}$ for given $\{\Delta a_{ij}^c(t)\}$ and $\{\Delta m_{ij}^c(t)\}$ solves the following linear program (MA) on each $(i,j) \in \mathcal{W}$:

$$\text{(MA)}\ \max_{\boldsymbol{\pi}^c} \sum_{t,s\in\mathcal{T}} \phi(s-t) \cdot \pi_{ij}^c(t,s)$$

$$\text{s.t.} \sum_{t\in\mathcal{T}} \pi_{ij}^c(t,s) \leqslant \Delta m_{ij}^c(s) \qquad \forall s \in \mathcal{T}$$

$$\sum_{s\in\mathcal{T}} \pi_{ij}^c(t,s) \leqslant \Delta a_{ij}^c(t) \qquad \forall t \in \mathcal{T}$$

$$\pi_{ij}^c(t,s) \geqslant 0, \text{with} \pi_{ij}^c(t,s) = 0 if t > s \qquad \forall t,\ s \in \mathcal{T}$$

where $\phi(s-t)$ represents an arbitrary value function that is strictly decreasing on the time duration $s-t$.

With $\{\pi_{ij}^c(t,s)\}$ ready, we then calculate the waiting time and pickup time for riders who request trips at t as per Eqs. (15), and apply these costs to Eq. (14) to update the trip demand.

Marginal effect module

Implementing the forward propagation scheme gives rise to the balanced flow pattern. However, with the flow information alone, it is difficult to identify where the operation bottlenecks are in the ride-hailing system. More importantly, we are blind to the global effects of the platform's interventions through local actions.

This module thus leverages a network flow model (NF) to make up the missing components by quantifying the marginal value of each movement/trip of drivers/riders, i.e., the change in the objective value caused by a unit change in a flow quantity. The model characterizes a linear program degenerated from the original policy optimization problem (16), i.e.,

$$\text{(NF)} \max_{\boldsymbol{q^c},\boldsymbol{q^d},\boldsymbol{\Delta d},\boldsymbol{\Delta m^c},\boldsymbol{\Delta m^d},\boldsymbol{\Delta c}} Z = \sum_{(i,j)\in\mathcal{W},\ t\in\mathcal{T}} \Phi_{ij}(t)\cdot m_{ij}^d(t) \tag{17a}$$

$$\text{s.t.}\quad q_{ij}^c(t+1) = q_{ij}^c(t) + \Delta a_{ij}^c(t) - \Delta m_{ij}^c(t), \quad \forall (i,j)\in\mathcal{W} \tag{17b}$$

$$\hat{q}_i^d(t) = q_i^d(t) + \sum_{j:(j,i)\in\mathcal{W}} \Delta d_{ji}(t) + \Delta e_i(t) - \sum_{j:(i,j)\in\mathcal{W}} \Delta m_{ij}^d(t), \quad \forall i\in\mathcal{N} \tag{17c}$$

$$q_i^d(t+1) = \hat{q}_i^d(t) - \Delta b_i(t) - \sum_{j:(i,j)\in\mathcal{W}^c} \Delta c_{ij}(t) + \sum_{j:(j,i)\in\mathcal{W}^c} \Delta c_{ji}(t), \quad \forall i\in\mathcal{N} \tag{17d}$$

$$\Delta d_{ij}(t) = \sum_{\tau\in\mathcal{T}} \mathbb{1}\left(\tau + w_{ij}^s(\tau) = t\right)\cdot \Delta m_{ij}^d(\tau), \quad \forall (i,j)\in\mathcal{W} \tag{17e}$$

$$\Delta m_{ij}^d(t) = \eta_{ij}^s(t)\cdot \Delta m_{ij}^c(t), \quad \forall (i,j)\in\mathcal{W} \tag{17f}$$

$$\Delta b_i(t) = \hat{q}_i^d(t)\cdot \Lambda_t^b(t), \quad \forall i\in\mathcal{N} \tag{17g}$$

$$\Delta c_{ij}(t) = \hat{q}_i^d(t)\cdot \Lambda_{ij}^c(t), \quad \forall (i,j)\in\mathcal{W}^c \tag{17h}$$

$$\Delta m_{ij}^c(t) \geq \bar{S}_{ij}^c(t)\cdot\left(q_{ij}^c(t) + \Delta a_{ij}^c(t)\right), \quad \forall (i,j)\in\mathcal{W} \tag{17i}$$

$$\sum_{j:(i,j)\in\mathcal{W}} \Delta m_{ij}^d(t) \leqslant \bar{U}_i^d(t), \quad \forall (i,j)\in\mathcal{W} \tag{17j}$$

All the constraints are essentially the flow propagation procedures specified in the flow rebalancing module. We only replace the matching component with the last two equations (i.e., 17 i,j), where $\bar{S}_{ij}^c(t)$ and $\bar{U}_i^d(t)$ are notions previously created based on the balanced flows. As one may notice, under such a construction, the balanced flow pattern solved by the forward propagation scheme is indeed the only solution that satisfies all the constraints of the system (17). Thus, the problem is readily solved.

Our purpose of solving this problem is to quickly retrieve the dual for different variables and constraints, which indicates the marginal value of various flows in the spatiotemporal network. For example, the dual to constraint (17 b) suggests the marginal value of increasing the trip demand from type-(i,j) riders at time t. Meanwhile, the variables $\delta_{ij}^c(t)$ and $\Delta c_{ij}(t)$'s dual can imply the marginal effect of speeding up the matching for type-(i,j) riders and the repositioning of idle drivers from region i to j at period t, respectively. All this information can be useful in analyzing pricing, matching, and repositioning strategies, and is thus collected and fed into the last module for policy improvement.

Policy improving module

The retrieval of marginal flow values enables us to formulate the policy improving problem as a quadratically constrained program that can be solved efficiently. To better present the idea, we define $\boldsymbol{f}$ as an augmented vector that spells all types of flows, including those of customer arrivals $\boldsymbol{\Delta a^c}$, matching $\boldsymbol{\Delta m}$, and repositioning of drivers $\boldsymbol{\Delta c}$, etc. Also, define $\boldsymbol{P}$ as a summary of all the existing policies considered. Let $\boldsymbol{\Delta f}$ and $\boldsymbol{\Delta P}$ denote the corresponding incremental changes in flow volumes and policies, respectively.

The policy improving module then solves the program (PI) below for $\boldsymbol{\Delta P}$:

$$\text{(PI)} \quad \max_{\boldsymbol{\Delta P}, \boldsymbol{\Delta f}} \ \boldsymbol{\psi_f} \cdot \boldsymbol{\Delta f} + \boldsymbol{f} \cdot \boldsymbol{\Psi}(\boldsymbol{\Delta P}) \tag{18a}$$

$$\text{s.t.} \quad \boldsymbol{\Delta f} = \frac{\boldsymbol{\partial f}}{\boldsymbol{\partial P}} \cdot \boldsymbol{\Delta P} \tag{18b}$$

$$\|\boldsymbol{K} \cdot \boldsymbol{\Delta f}\| \leqslant \boldsymbol{\nu} \tag{18c}$$

where $\boldsymbol{\psi_f}$ denotes the marginal values for different flows obtained from the previous module; $\Psi(\cdot)$ specifies a linear function that captures the changes of values for each unit of driver and rider flows due to the updates of the policy ΔP. The two terms $\boldsymbol{\psi_f} \cdot \boldsymbol{\Delta f}$ and $\boldsymbol{f} \cdot \boldsymbol{\Psi}(\boldsymbol{\Delta P})$ in (18 a) capture the consequences on the original optimization objective (16) from the incremental and existing flows, respectively. The constraint (18 b) characterizes a possibly huge system of linear equations that connect the changes of flows to the policy adjustments. The Jacobian matrix $\frac{\boldsymbol{\partial f}}{\boldsymbol{\partial P}}$ here is taken as the constant that valued at the current system state $(\boldsymbol{f}, \boldsymbol{P})$. Note that the derivations of the Jacobian matrix may involve some efforts of system approximations, as not all the relationships are explicitly functional. Since all the impacts are evaluated with the first-order approximations, quadratic regularization constraints (18c) are imposed to ensure the local search of a new policy, where $\boldsymbol{K}$ is a coefficient matrix and $\boldsymbol{\nu}$ is a set of bounds. Some fine-tuning of both parameters could significantly accelerate the search of improving policies in implementation.

4.2. Solution procedure

Given the above deductions, we now present the following solution procedure to solve the policy optimization (see Fig. 8 for the corresponding flow chart):

Initialize the policy $\boldsymbol{\mathcal{P}}, \boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{R}}$ and trip demand $\boldsymbol{\Delta a^c}$

S1. Derive the balanced flow pattern and pickup time by executing (FB) in forward propagation, and calculate the objective value $\hat{Z}$ given by Eq. (16).
S2. Obtain the matching assignment $\{\pi_{ij}^c(t,s)\}$ by solving (MA) and then update trip demand $\Delta\hat{a}^c$ for Step 5.
S3. Solve the network flow problem (NF) to retrieve the marginal value of different actions/policies.
S4. Obtain the new policy $\{\hat{\boldsymbol{\mathcal{P}}}, \hat{\boldsymbol{\mathcal{A}}}, \hat{\boldsymbol{\mathcal{R}}}\}$ by solving (PI).
S5. If $\|\boldsymbol{\Delta a^c} - \boldsymbol{\Delta \hat{a}^c}\| < \varepsilon_a$, then
- Update the "optimal" policy with $\{\hat{\boldsymbol{\mathcal{P}}}, \hat{\boldsymbol{\mathcal{A}}}, \hat{\boldsymbol{\mathcal{R}}}\}$, if $\hat{Z}$ is higher than the best value achieved so far,
- Stop the search and return the "optimal" policy, if $\kappa_p\|\boldsymbol{\mathcal{P}} - \hat{\boldsymbol{\mathcal{P}}}\| + \kappa_a\|\boldsymbol{\mathcal{A}} - \hat{\boldsymbol{\mathcal{A}}}\| + \kappa_r\|\boldsymbol{\mathcal{R}} - \hat{\boldsymbol{\mathcal{R}}}\| < \varepsilon_P$;

Otherwise, update the current policy and demand $\{\boldsymbol{\mathcal{P}}, \boldsymbol{\mathcal{A}}, \boldsymbol{\mathcal{R}}, \boldsymbol{\Delta a^c}\}$ using its linear combination with $\{\hat{\boldsymbol{\mathcal{P}}}, \hat{\boldsymbol{\mathcal{A}}}, \hat{\boldsymbol{\mathcal{R}}}, \boldsymbol{\Delta \hat{a}^c}\}$ and repeat the process from S1.

where $\boldsymbol{\kappa}$ in the last step characterizes a set of weights and $\boldsymbol{\varepsilon}$ is a given tolerance. Given the complexity of the procedure, its convergence properties are not yet theoretically investigated. However, for our all numerical experiments conducted for counterfactual analyses in the subsequent section, the procedure has always converged efficiently.

5. Counterfactual analysis on pricing policy

To exemplify and demonstrate the applicability of our framework, this section adopts the above solution procedure and conducts counterfactual analyses on pricing policy. The potentials of four different pricing schemes on optimizing system performance are respectively examined and compared. The baseline is the status-quo pricing scheme with a fixed price structure, regardless of riders' location and time. The second scenario considers a time-of-day (dynamic) pricing that at each period applies a surge multiplier to adjust the price proportionately for riders. On top of dynamic pricing, the third and fourth scenario further varies the price by riders' origins (spatial pricing) and OD pairs (OD-specific pricing), respectively. For the latter three scenarios, the optimal pricing strategies are investigated to achieve different goals of service management.

Our counterfactual experiments are conducted in the morning (6:00-10:00 AM) of the averaged weekday (see Section 3.3 for the production of the "typical" weekday). The constrained trip fares are allowed to vary in the range between a 15% discount to a 50% surge based on the status-quo prices. The proposed procedure has converged successfully

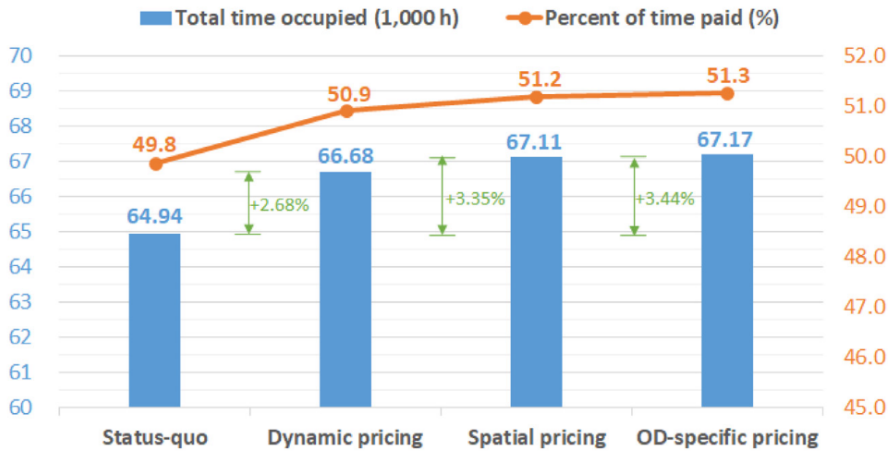

Fig. 9. Comparison of the optimal pricing schemes under throughput maximization.

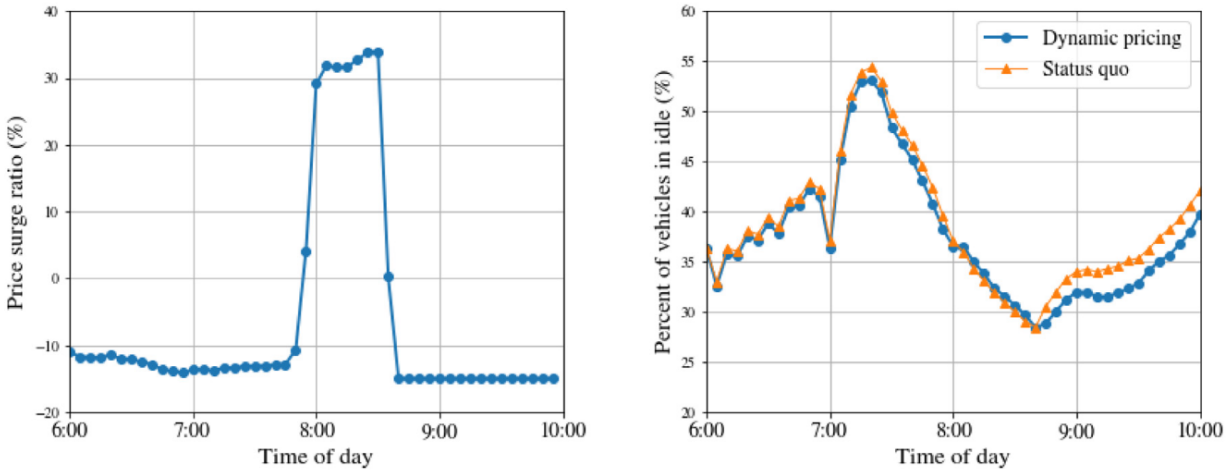

(a) Price surge ratio (%) (b) Percent of vehicles in idle (%)

Fig. 10. Dynamic pricing for throughput maximization in the morning.

in all these numerical cases within four to five hundreds of iterations, each taking about six seconds to execute with our GAMS code on a desktop with Intel i5-6600 CPU@3.30Ghz and 8GB memory.

5.1. Throughput-maximizing pricing policies

We first examine the policies that maximize the system throughput, which is defined as the number of customers delivered to their destinations per unit of time, an important metric for the production of a ride-hailing system. It is also directly related to the utilization of ride-hailing vehicles. For a system with relatively homogeneous trip length among passengers, the system throughput is proportional to the percent time that in-service vehicles are occupied and paid.

Fig. 9 first compares the system throughput (measured by the total time being occupied or paid) yielded by four pricing scenarios. The blue columns represent the total occupied time of drivers, while the orange polyline presents their percent time being paid during the investigated period. As we can see in Fig. 9, both measures improve constantly as the pricing policy becomes more differentiated. Specifically, the ride-hailing system will be able to accommodate 2.68% more passenger trip time by allowing the price to vary in time, and the number further rises to more than 3% as additional flexibility is provided for handling the spatial heterogeneity. Meanwhile, drivers in service also achieve higher utilization as the percent of time paid increases from 49.8% to 51.3%. These improvements imply the potentials of the system to accommodate more passenger trips and utilize the supply resources more efficiently through price adjustments.

To help us better understand the platform's strategies, Fig. 10 and 11 respectively visualize the surge ratios (or adjustment ratios) of prices for maximizing the system throughput under dynamic and spatial pricing. The surge ratios suggested by both figures indicate the necessity of context-dependent strategies in service management, as the market condition varies drastically in space and time. Fig. 10a shows that a throughput-maximizing platform would lower the prices in non-peak periods, while lifting them during the rush hours. For non-peak periods, the system with sufficient supply retains extra

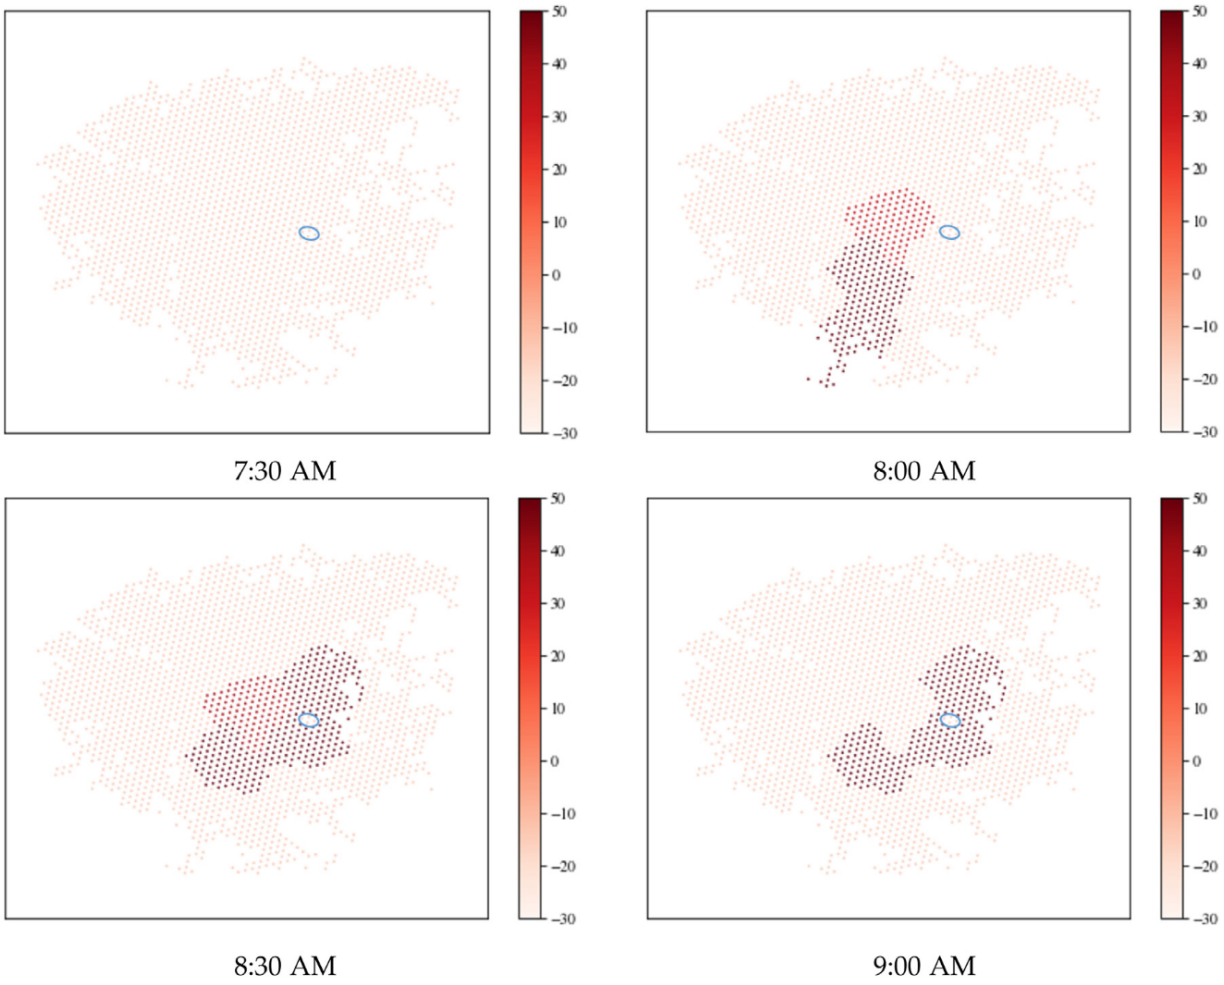

Fig. 11. Surge ratios of the throughput-maximizing spatial pricing during the morning peak. The blue circled area marks a railway station, i.e. one of the three service spots identified.

capacity to accommodate more trip of rides. By lowering down the prices, the platform aims to attract and serve more customers. In contrast, the platform opts to increase the price to suppress the demand for achieving higher throughput during the peak period, i.e., 8:00-8:30 AM, for this city. Such a strategy is actually consistent with the suggestion discussed in the literature to address the “wild goose chases” phenomenon, where a platform may need to buffer the idle vehicle accumulations by suppressing the demand to avoid the enormous supply wastes in devastatingly long pickups during the peak periods (see Castillo et al., 2018; Xu et al., 2020 for specific discussions). Fig. 10b evidences such reasoning by comparing the percent of idle vehicles for dynamic pricing and status quo, respectively. It is observed that the dynamic pricing scheme lowers the percent of vehicles in idle during most of the time except for the peak period, where the value stays constantly higher than the corresponding level at the status quo. The same contrast is again captured in Fig. 11, which presents four representative instances of price-surge maps with darker colors marking higher surge ratios. As can be seen, the platform may still lower the price in most of the time and areas to increase the system throughput. However, for the peak period, a few regions in the central downtown need price surges. Since lower prices are suggested in most cases by the above implementation, the service revenue of the platform will not increase despite the efficiency gains. As drivers expect higher revenue to offset the wear and tear caused by additional service provision, the enhancements come at the cost of a platform aggressively subsiding the system, which may not be sustainable in the long run. A mechanism is thus needed to fairly allocate the costs and benefits of efficiency gains among stakeholders, making the schemes more competitive for implementation.

Besides, it is interesting to note that the dynamic policy specifies a distinct pricing strategy at a service spot versus its surroundings. As marked in Fig. 11, while the surrounding area raises prices during peak hours in response to the supply shortage, the price at the railway station remains discounted throughout, implying the invariably oversupplied condition therein. In fact, drawn by the relatively high profitability of long-haul trips, many ride-hailing drivers prefer to "hunt" at service spots, enduring a long wait, rather than serving more frequent but shorter trips in urban areas (Yang et al., 2010).

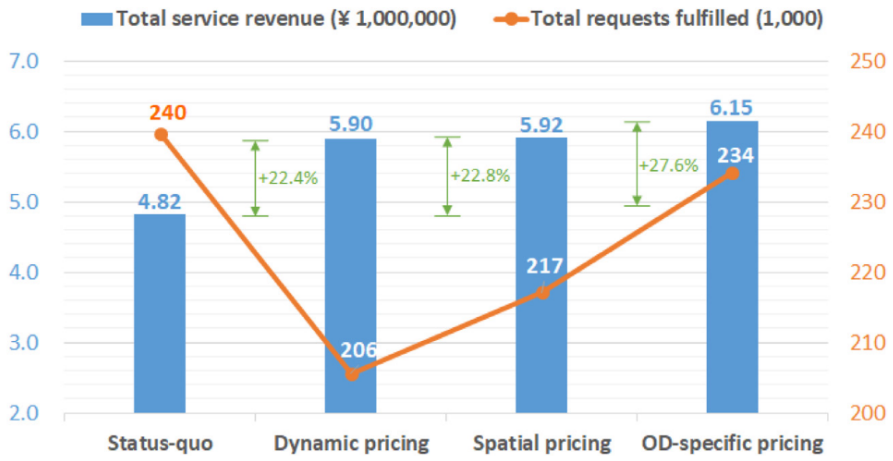

Fig. 12. Comparison of the optimal pricing schemes under revenue maximization.

It is thus critical to distinguish the spatial and spot-type service regions to facilitate specific operations that cater to the contrasting market characteristics and conditions.

5.2. Revenue-maximizing pricing policies

We proceed to investigate another case where the platform optimizes the pricing policy aiming to maximize the total service revenue. Two performance metrics are compared in Fig. 12 for the four schemes. The blue columns now represent the total service revenue of the ride-hailing market, while the orange polyline presents the total number of requests fulfilled. As clearly shown by the figure, the platform can benefit substantially from higher levels of pricing differentiation. Only allowing price discrimination in time could give rise to a 22.4% increment in service revenue, while the optimal OD-specific pricing yields 27.6% more revenue than the current fixed-price condition. On the other hand, the number of customers served does not necessarily increase amid the revenue boosts. As we can see from the figure, by allowing a revenue-maximizing platform to implement time-dependent pricing could incur a 14% reduction in the number of requests fulfilled. The platform can strategically raise the price to serve fewer customers but receive higher revenue. In fact, in our analyses the platform would opt to price at the maximum surge ratio in most of the time and areas with only a few exceptions. our counterfactual analysis suggests that the current pricing structure adopted by the platform does not maximize its revenue for this particular market and doing so via price differentiation in time and space will compromise social welfare and consumers' surplus of the market.

6. Conclusions and future research

In this paper, we have proposed a generalized fluid framework for modeling ride-hailing systems and demonstrated its applicability by customizing and calibrating the framework for a real system operated by Didi Chuxing. The customized model was further applied to prescribe optimal pricing policies for the system, and the results reveal the potentials of pricing policies for improving the performance of the system. The model can be readily applied to explore other operational strategies such as matching and dispatching.

We believe that the proposed fluid framework has made a good balance between mathematical tractability and behavioral realism to allow service providers to answer what-if questions quickly and optimize their strategic operation decisions. The current framework mainly applies to a monopolistic market where a ride-hailing platform manages to maximize its own business goals. A meaningful extension in the future is to model the competitive ride-hailing market with multiple service providers. It would be interesting to incorporate riders and drivers' strategic behaviors on, e.g., multi-homing and information-sharing, and see how each single platform could react accordingly given its market power. Further, the framework can also be extended to capture the interaction between ride-hailing services with other modes of transportation such as private cars and transit and facilities such as parking and dedicated pick-up/drop-off locations for ride-hailing services. An important consideration in this extension is to capture the congestion effect of in-service ride-hailing vehicles, which can be achieved by leveraging recent developments in network macroscopic fundamental diagrams, as demonstrated by, e.g., Xu et al. (2017) and Ramezani and Nourinejad (2018). The resulting framework will be able to support the planning practice and policy making of governmental agencies in applications such as regulation of ride-hailing services, congestion pricing, and leveraging ride hailing to solve the first mile/last mile problem of transit systems.

CRediT authorship contribution statement

Zhengtian Xu: Methodology, Investigation, Writing - original draft. **Yafeng Yin:** Conceptualization, Methodology, Writing - review & editing, Supervision, Funding acquisition. **Xiuli Chao:** Methodology, Funding acquisition. **Hongtu Zhu:** Methodology, Funding acquisition. **Jieping Ye:** Methodology, Funding acquisition.

Acknowledgement

We would like to thank Dr. Shikai Luo and Mr. Qichao Wang from Didi Chuxing for their professional and invaluable support on empirical data preparation and model calibrations. The work described in this paper was partly supported by research grants from the National Science Foundation (CMMI-1740865 and CMMI-1854684) and Didi Chuxing.

References

Arnott, R., 1996. Taxi travel should be subsidized. J. Urban Econ. 40 (3), 316–333.
Ban, X.J., Dessouky, M., Pang, J.-S., Fan, R., 2019. A general equilibrium model for transportation systems with e-hailing services and flow congestion. Transportation Research Part B: Methodological 129, 273–304.
Bimpikis, K., Candogan, O., Saban, D., 2019. Spatial pricing in ride-sharing networks. Oper. Res. 67 (3), 744–769.
Braverman, A., Dai, J.G., Liu, X., Ying, L., 2019. Empty-car routing in ridesharing systems. Oper. Res. 67 (5), 1437–1452.
Buchholz, N., 2019. Spatial equilibrium, search frictions and dynamic efficiency in the taxi industry. Available at https://scholar.princeton.edu/sites/default/files/nbuchholz/files/taxi_draft.pdf. Accessed on July 5, 2020.
Castillo, J. C., Knoepfle, D. T., Weyl, E. G., 2018. Surge pricing solves the wild goose chase. Available at SSRN: https://ssrn.com/abstract=2890666. Accessed on December 15, 2019.
Chen, X.M., Zheng, H., Ke, J., Yang, H., 2020. Dynamic optimization strategies for on-demand ride services platform: surge pricing, commission rate, and incentives. Transportation Research Part B: Methodological 138, 23–45.
Daganzo, C.F., 1978. An approximate analytic model of many-to-many demand responsive transportation systems. Transp. Res. 12 (5), 325–333.
Ke, J., Xiao, F., Yang, H., Ye, J., 2019. Optimizing online matching for ride-sourcing services with multi-agent deep reinforcement learning. Available at arXiv preprint:1902.06228. Accessed on December 15, 2019.
Newell, G.F., 1971. Applications of queueing theory. Chapman and Hall, London.
Nourinejad, M., Ramezani, M., 2020. Ride-sourcing modeling and pricing in non-equilibrium two-sided markets. Transportation Research Part B: Methodological 132, 340–357.
Ramezani, M., Nourinejad, M., 2018. Dynamic modeling and control of taxi services in large-scale urban networks: a macroscopic approach. Transportation Research Part C: Emerging Technologies 94, 203–219.
Salanova, J.M., Estrada, M., Aifadopoulou, G., Mitsakis, E., 2011. A review of the modeling of taxi services. Procedia-Social and Behavioral Sciences 20, 150–161.
Shou, Z., Di, X., Ye, J., Zhu, H., Zhang, H., 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.
Urata, J., Xu, Z., Ke, J., Wu, G., Yin, Y., Yang, H., Ye, J., 2020. Learning ride-sourcing drivers customer-searching behavior: a dynamic discrete choice approach. Manuscript Submitted for Publication.
Wang, H., Yang, H., 2019. Ridesourcing systems: a framework and review. Transportation Research Part B: Methodological 129, 122–155.
Xu, Z., Chen, Z., Yin, Y., Ye, J., 2019. Equilibrium analysis of urban traffic networks with ride-sourcing services. Available at SSRN: https://ssrn.com/abstract=3422294. Accessed on July 5, 2020.
Xu, Z., Yin, Y., Ye, J., 2020. On the supply curve of ride-hailing systems. Transportation Research Part B: Methodological 132, 29–43.
Xu, Z., Yin, Y., Zha, L., 2017. Optimal parking provision for ride-sourcing services. Transportation Research Part B: Methodological 105, 559–578.
Yang, H., Fung, C., Wong, K.I., Wong, S.C., 2010. Nonlinear pricing of taxi services. Transportation Research Part A: Policy and Practice 44 (5), 337–348.
Yang, H., Qin, X., Ke, J., Ye, J., 2020. Optimizing matching time interval and matching radius in on-demand ride-sourcing markets. Transportation Research Part B: Methodological 131, 84–105.
Yang, H., Shao, C., Wang, H., Ye, J., 2020. Integrated reward scheme and surge pricing in a ridesourcing market. Transportation Research Part B: Methodological 134, 126–142.
Yang, H., Wong, S., 1998. A network model of urban taxi services. Transportation Research Part B: Methodological 32 (4), 235–246.
Zha, L., Yin, Y., Du, Y., et al., 2018. Surge pricing and labor supply in the ride-sourcing market. Transportation Research Part B: Methodological 117 (Part B), 708–722.
Zha, L., Yin, Y., Xu, Z., 2018. Geometric matching and spatial pricing in ride-sourcing markets. Transportation Research Part C: Emerging Technologies 92, 58–75.
Zha, L., Yin, Y., Yang, H., 2016. Economic analysis of ride-sourcing markets. Transportation Research Part C: Emerging Technologies 71, 249–266.
Zhang, K., Chen, H., Yao, S., Xu, L., Ge, J., Liu, X., Nie, M., 2019. An efficiency paradox of uberization. Available at SSRN: https://ssrn.com/abstract=3462912. Accessed on December 15, 2019.