Submitted to *Operations Research*
manuscript (Please, provide the manuscript number!)

The value of knowing drivers' opportunity cost in Ride Sharing systems

Consider a ride sharing platform, and a large population of strategic potential drivers, heterogeneous in terms of their opportunity costs, who choose whether or not to work for that platform. The platform is endowed with knowledge about the different drivers' opportunity costs. How can the platform implement a matching policy that uses this knowledge in order to improve system efficiency? Can such improvement be quantified? In this work we introduce an analytically-tractable mean field (fluid) model that accounts for the dynamic nature of drivers' spatial location, revenue, and availability status. Based on this model we compare drivers' equilibrium participation under two different matching policies. Our analysis leads to improvement bounds on the equilibrium performance: We show that a policy which utilizes knowledge about drivers' opportunity costs can perform up to two times better than a policy that does not do so, in terms of the number of drivers it attracts and in terms of the rate of matches it produces. We demonstrate by simulation that the mean field model provides an accurate approximation for a corresponding (stochastic) discrete model, in which the discussed improvement is observed empirically.

Key words: ...

History: ...

1. Background and motivation

Many companies in the sharing economy sector, and ride sharing platforms in particular, are putting tremendous efforts to retain their partners (specifically, drivers in the case of ride sharing), as well as recruiting new ones. These companies are providing potential and existing partners with monetary incentives to work for the platforms, for example by offering sign-up promotions. A prevalent method in the ride-sharing business to increase drivers' total work time is bonus programs that reward drivers for achieving certain productivity benchmarks. Though the precise format and motives of different bonus programs may vary, by and large, their high-level goal is similar, which is to increase drivers' supply in the (possibly time/location-specific) market.

Just as any other economic system, the fact that platforms can benefit from having a large service capacity initially seems obvious. However, a NYC Department of Transportation report (Department of Transportation (2019)) states, with regard to for-hire vehicles in the Manhattan core area, that

> *"...Uber, Lyft, Juno and Via [...] saturated the market with vehicles [...] causing drivers to spend over 40% of total work time empty and cruising for passengers. [...] this underutilization led to significant declines in driver income."*

Given that drivers spend close to half of their working time idling, it is perhaps surprising that platforms invest abundant resources, via bonuses and promotions, in increasing their potential supply. The rationale is given in part in Parrott and Reich (2018) – a report of the NYC Taxi and Limousine Commission:

> *"To achieve quick response times, the [ride-sharing] companies require many idle drivers to be available at any given moment and at many locations. This model creates a gap between the drivers' desires to maximize their earnings – by maximizing trips per working hour – and the companies' desire to minimize response times. In other words, the app business model works only if it keeps driver utilization low, which then keeps drivers' hourly pay low as well."*

In order to design efficient bonus and promotion programs, it is crucial to understand the value of drivers' commitment, i.e., to be able to quantify the benefit for the platform from having more drivers working for it in a given time. In this paper we interpret the latter quantity as the marginal value for the platform, either in terms of improved matching rate or in terms of response times, with respect to growth in working-drivers intensity.

What distinguishes drivers in ride sharing platforms from service providers in more traditional markets is that ride-sharing drivers are classified private contractors who get paid per job, as opposed to employees getting paid a fixed (global) salary. Thus, drivers are highly sensitive to changes in the system's overall performance. Performance measures, such as driver utilization, often exert impact on drivers' income – a fact notably highlighted by the Parrott and Reich (2018) and the Department of Transportation (2019) reports. In this paper, the main performance measure of interest will be the matching rate, i.e., the average number of passenger-driver matches produced by the platform per unit of time, which directly relates to per-driver income and utilization.

The matching rate is a byproduct of the balance between the demand side, namely potential passengers, and the supply side, which are the active drivers that are available for pick up. However, a simplistic supply-demand economic model that ignores the spatial dynamics of the system falls short of capturing the interplay between drivers utilization and response times as discussed in Parrott and Reich (2018). For example, a naive model for a ride sharing market considers a closed network with N drivers in circulation, with a pool of available drivers, ride requests arriving at rate λN per unit time, and a trip duration m for each ride (for concreteness, this can be thought of as a single-node BCMP queueing network). Each ride request upon arrival is matched arbitrarily with an available driver if the pool is not empty, otherwise the ride is lost. When N is large, this network behaves like fluid in circulation, where the proportion of fluid in the pool corresponds to the fraction of time drivers spend idling, and the demand flow is analogous to a pump that can drain fluid from the pool at rate not larger than λ. Such modeling approach will always lead to the conclusion that either (a) the pool is always empty, namely, all drivers are busy 100% of the time; or (b) the pool may not be empty but the pump is working at its maximal rate, in which case all the demand is filled. In case (b), the system is saturated with fluid (i.e., drivers), and increasing the amount of it will bring about any gains in terms of throughput. The data from the two reports regarding drivers idling time immediately contradict conclusion (a), while conclusion (b) falls short of explaining how ride sharing platforms benefit from saturating the market with drivers, albeit low per-driver utilization.

In the paper we introduce a novel modeling approach that resolves this conflict. Specifically, we introduce a spatial setting with pickup-time range – passengers whose ride requests cannot be filled immediately by an available driver within their pickup region are never matched and we consider them lost demand. We adopt a mean field approach that allows us to perform exact analysis to quantify the value for the platform, in terms of increased matching rate, and/or improved pickup range, from having an increased supply of drivers. Pickup time is a factor of significant importance, as mentioned in Parrott and Reich (2018):

> *"The [ride sharing] companies compete with each other primarily by minimizing passengers' wait times and, to a lesser extent, by decreasing fares."*

This means that for the platform to remain relevant in the market it must adhere to a pickup range competitive with that of other platforms, which is an underlying assumption of our model.

As one would expect, when potential demand grows, thereby increasing the matching rate (all other factors unchanged), drivers in the system spend less time idling and hence their income grows as well. However, a natural trade off is that a higher per-driver profit makes it more attractive to other potential drivers to work for the system, an act that generally imposes negative externalities on other drivers in the platform due to the arising competition between them. Thus, a question of economic equilibrium is brought about: Can one identify the supply level under which all active drivers meet their reservation wage?

Assuming that the platform is capable of employing matching policies that differentiate between drivers based on their heterogeneous characteristics, drivers' revenues and therefore the equilibrium supply level will depend on the matching policy. It is therefore of interest to compare different matching policies with respect to the equilibrium matching rates that each policy produces. In this work, we compare two policies subject to a certain pickup-time standard: One policy that treats all drivers symmetrically, ignoring all information about their income goals, and another policy which is geared towards helping drivers meet their income goals. The key takeaway that raises from the analysis is simple: Compared to a symmetric policy, a policy that takes into account drivers income goals in its matching decisions yields a more efficient distribution of the revenue among drivers and hence increases the number of active drivers in equilibrium. An increased number of drivers yields a better coverage of drivers over the city area and therefore reduces lost demand, hence it is more Pareto efficient. In the extreme, the matching rate can be improved by up to 100%. From a pickup range perspective, this means that if a platform switches from the first to the second policy, it can potentially reduce its pickup range (and hence its response time) without having to sacrifice potential demand.

2. The motivating discrete model

Prior to the formulation of the analytic model we describe in broad strokes the discrete simulation setup that motivates the mean field model to be analyzed. We study the discrete model by simulation and briefly explain how the simulation outcome supports the results accomplished in our mean field analysis.

2.1. Model entities

City: We consider a one-dimensional ring-shaped city which comprises a continuum of locations, represented by the unit interval $[0,1)$. The city being ring shaped implies that the distance between two locations $x, y \in [0,1)$ is given by $\text{dist}(x,y) = \min\{|x-y|, 1-|x-y|\}$. We assume that drivers' travel time between locations is a continuous increasing function of distance so there is a one-to-one correspondence between travel time and distance.

Drivers: In the city, there is an overall potential supply of $\Theta N > 0$ drivers, with $\Theta > 0$ being the potential-drivers' intensity. The scaling parameter N is interpreted as the *market size*. Each driver has an idiosyncratic opportunity cost rate (measured in money per unit-time) which is guaranteed to them in case they decide not to work for the platform. Driver's are heterogeneous in terms of their opportunity cost rates. To make insights sharp, we focus on the case where drivers are divided into 2 distinct classes, *high* (H) and *low* (L), which we refer to as types, each type $i \in \{L, H\}$ with potential number of drivers $\Theta_i N$ and opportunity cost rate κ_i, such that $\kappa_H > \kappa_L$, and $\Theta_L + \Theta_H = \Theta$.

Drivers are strategic entities who, before the system starts operating, can choose whether to work for the platform (i.e., *participate*) or not. Based on the system parameters and the platform's matching policy, which are common knowledge, each potential driver forms a belief about the participation of other potential drivers as well as the resulting revenue rate from working for the platform. Given their opportunity cost rate, they then make an irrevocable participation decision. The aggregate decision profile results in an endogenous total participation intensity $\theta = \theta_L + \theta_H$ where θ_i is the (endogenous) participation intensity per type $i \in \{L, H\}$. Accordingly, the effective total number of drivers who circulate the city is then given by θN, and similarly, $\theta_i N$ for each type $i \in \{L, H\}$.

While working, drivers alternate between states: *available* (idle) and *busy* (picking up/carrying a passenger). When carrying a passenger, a driver obtains revenue at rate of r per unit time, thus, r can equivalently be thought of as the revenue rate per unit distance traveling, so that a driver's revenue from carrying a passenger from a pickup location $x_p \in [0, 1)$ to their drop off $x_d \in [0, 1)$ is given by $r \cdot \text{dist}(x_p, x_d)$. When busy, drivers cannot be matched with new passengers. How the matching is formed and how it triggers alternations in drivers' states will be discussed shortly.

Passengers: Passengers generate the demand for service in the market by submitting ride requests into the system. Ride requests arrive following a Poisson process with rate $N\lambda > 0$, for some exogenous parameter $\lambda > 0$. Each request is associated with a single passenger and characterized by a pickup location $x_p \in [0, 1)$ and a drop-off location $x_d \in [0, 1)$. For simplicity, we assume that x_p and x_d are iid and uniformly distributed across the city's locations, $[0, 1)$. We refer to a passenger's location and their pickup location interchangeably.

A passenger has to be matched with a driver immediately at the moment of their arrival. When such matching is formed the ride request is carried out by the matched driver, and once completed, the passenger leaves the system. Yet, it may occur, depending on the system state, that a passenger cannot be matched with a driver, in which case the passenger leaves instantaneously, and no matching is formed.

Platform: The platform operates in a competitive market, in which passengers has different, comparable transportation possibilities. Thus, the platform has a rigid (exogenous) pickup range

$\frac{\delta}{2N}$: The latter assumption means that the platform matches a customer only with a driver not farther than $\frac{\delta}{2N}$ units distance away from a passenger's location. Because travel time as a function of distance is bijective, this implies that the platform adheres to a rigid pickup-time standard, which is increasing with δ. When a passenger arrives with pickup location x_p, all (if any) available drivers located in a $\frac{\delta}{N}$-long neighborhood centered at x_p constitute the set of candidates for the matching. Assuming such candidate driver exists, the matching policy that is set by the platform then determines which among the candidate drivers to match to that passenger. Loosely speaking, when N is the size of the market, the average distance between a passenger and their best outside transportation option is of order $\frac{1}{N}$. Thus, following the assumption that the platform's response time are competitive, the platform's pickup range is appropriately scaled.

While operating, the platform keeps track of each driver's state, which consists of the following attributes:

- the driver's availability status, i.e., *available* (idling) or *busy* (carrying out a ride);
- the driver's physical (spatial) *location* in the city;
- the driver's *accumulated revenue*.

Knowing also each driver's opportunity cost rate, the platform employs one of the following two matching policies: *MinRev*, which matches a ride with the candidate driver whose accumulated net revenue is the lowest; and *MinWeightRev*, that assigns the ride to the candidate driver whose net revenue rate normalized (or *weighted*) by their opportunity cost rate is the lowest.

Dynamics: For a given participation intensity θ there are θN drivers working for the platform, each driver at time $t = 0$ resides in an arbitrary location in the $[0, 1)$ interval. For the ease of exposition we assume that all drivers are initially available and start with 0 accumulated revenue. We describe the dynamics for the *MinRev* policy first.

Suppose a passenger arrives at time t_1 with pickup location x_p and drop-off location x_d. The platform then concerns all the candidate drivers, namely, drivers who by time t_1 are available and located in the interval $x_d \pm \frac{\delta}{2N}$. The platform chooses among the candidates the one who by time

t_1 had the minimal accumulated revenue, ties are broken arbitrarily. The matching is performed instantaneously. Once matched, the driver then changes their state from "available" to "busy". They then travel from their current location x to the pickup location x_p, and continue to the drop-off location x_p.Denoting by m to the mean duration of a ride, the mean revenue from a ride is then given by $r \cdot m$. Upon arrival to x_d, the driver becomes available again. To capture disturbance in available drivers' locations we assume that while available, drivers are making random movements in an $O\left(\frac{1}{N}\right)$-wide surrounding of their locations.

The matching mechanism for MinWeightRev is similar to that of MinRev, except that MinWeightRev differentiates between drivers based on their types. Recall that drivers are divided into two sets: drivers with outside opportunity κ_L and drivers with outside opportunity κ_H. For a given set of candidates, instead of comparing the actual accumulated revenue of candidates, MinWeightRev considers the *weighted* accumulated revenue, i.e., the accumulated revenue for each driver divided by their opportunity cost rate. It will then select the candidate driver with the minimal weighted accumulated revenue.

2.2. Equilibrium participation

Our main interest is in analyzing and comparing the participation intensities and the resulting matching rate under equilibrium induced by each of the two aforementioned policies. Equilibrium is characterized by a decision profile, namely a pair of participation intensities (θ_L, θ_H), such that when the system is in its steady state, no customer can benefit by deviating from their joining decision. Our equilibrium analysis is therefore valid under the assumption that drivers' states (availability, location and revenue rate) reach stationarity.

In order to make their decision, each driver has to asses their long run revenue rate from joining given the decision of others, thus, a game between drivers is brought about. However, exact analysis of the discrete model with a large yet finite market size N, is highly complex, mainly due to the dependencies between drivers' states. Thus, deriving performance measures that correspond to various decision profiles, and specifically the evaluation of drivers' revenue rate for every such profile, become impractical to carry out analytically. In this paper we therefore perform our analysis on a mean field (fluid) model, which is shown to provide a fairly accurate approximation for these performance measures when the value of N is large.

2.3. Simulation results

Below we present some results obtained in two separate simulation experiments, one for each of the two policies, MinRev and MinWeightRev. The purpose of these experiments is to demonstrate the equilibrium performance of each of the two policies under the same market condition. The setting consists of a market of size $N = 100$ and potential drivers intensity $\Theta = 1.25$ with per-type potential intensities $\Theta_L = 1$ and $\Theta_H = .25$ and corresponding opportunity costs $\kappa_L = .6$ and $\kappa_H = .8$. Thus, the potential market consists of 125 drivers with 80% being type-L and the rest type-H. The arrival rate and mean ride duration are given by $\lambda = m = 1$. In our simulations choose the bijection between travel times and distances so as to make the ride duration exponentially distributed. The working wage is $r = 1.04$, and the pickup range is $\delta = 5$.

First we simulate the system working under MinRev with total (effective) participation intensity $\theta = 1$ (corresponding to a total number of $\theta N = 100$ drivers) from time 0 until time $t = 10^3$, summing up to a total number of (roughly) $\lambda \times N \times t = 10^5$ arriving passengers during that period. The left panel of Figure 1 shows a histogram of the revenue rates of all the 100 participating drivers, from which it can be seen that the distribution is concentrated around its mean which is approximately .72. This strongly suggests that over time, the revenue rates of all drivers converge to the same value, which is expected due to symmetry between drivers.

In this particular setup the observed mean revenue rate (=.72) is between κ_L(=.6) and κ_H (=.8). This means that if all the 100 type-L drivers were to participate, they should expect a revenue rate higher then their opportunity cost. It is intuitive that if some type-H drivers were to participate as well (so as to make the total number of participating drivers larger than 100), then the revenue rate per driver would decrease, which is already below κ_H, suggesting that in equilibrium, $\theta = \theta_L = \Theta_L = 1$, and $\theta_H = 0$, i.e., 80% of the potential market participate under MinRev.

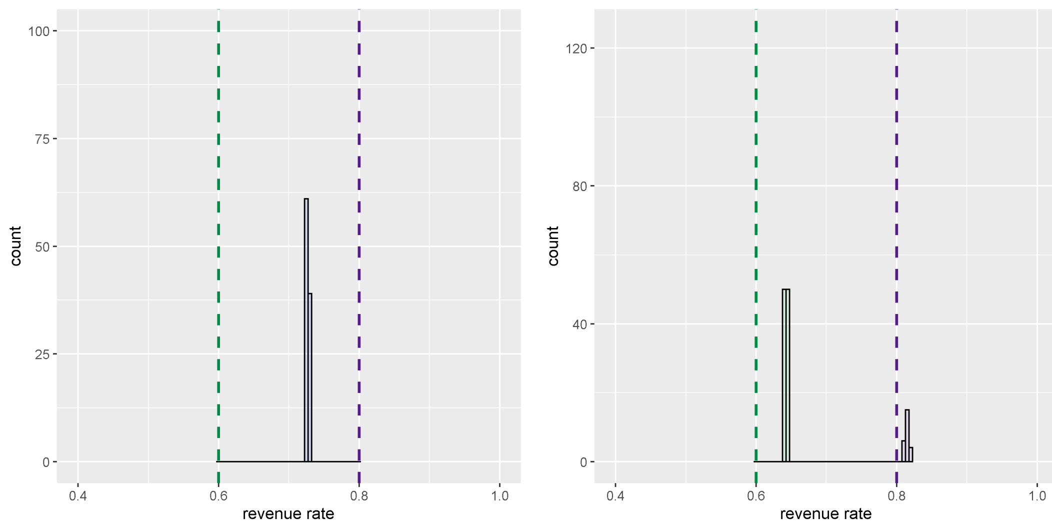

Figure 1 Revenue rate histogram in equilibrium for each policy. Dashed vertical lines correspond to $\kappa_L = 0.6$ (green) and $\kappa_H = 0.8$ (purple). In the left panel, the policy is MinRev and the number of drivers participating in equilibrium is 100 (corresponding to $\theta = 1$), with average per-driver revenue rate of .72. In the right panel, the policy is MinWeightRev and the number of drivers participating in equilibrium is 125 (corresponding to $\theta = 1.25$), with type-L (green) earning .64 and type-H (purple) earning .81 on average.

Next, we simulate the system under MinWeightRev with total participation intensity $\theta = 1.25$, namely, all 125 potential drivers participate, with 100 type-L drivers and 25 type-H (keeping all other parameters the same). As before, we simulate the system from time 0 to $t = 10^3$, and plot the histogram of revenue rates for each of the two types in the right panel of Figure 1. Remarkably, for both types, the average revenue rate is above their opportunity costs, with an average revenue rate of approximately $.64 > \kappa_L = .6$ for type L and $.81 > \kappa_H = .8$ for type H. Thus, under MinWeightRev, all the 125 potential drivers participate in equilibrium, leading to a 25% increase in drivers participation compared to MinRev.

For each arriving passenger, we count the number of candidates (i.e., drivers available in the passenger's pickup range at the time of their arrival), thus, generating roughly 10^5 samples. Figure 2 shows the corresponding histogram of these samples (blue bars), in each of the two experiments: In the left panel, the policy is MinRev with its equilibrium participation intensity $\theta = 1$, and in the right panel the policy is MinWeightRev with its equilibrium participation intensity $\theta = 1.25$. In both

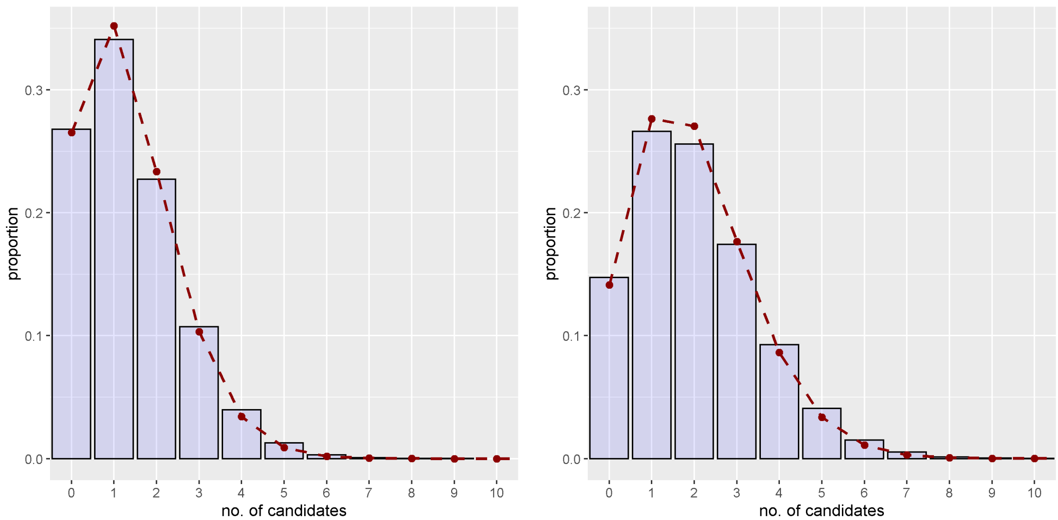

Figure 2 No. of candidates histogram in equilibrium for each policy. The fraction of lost rides (namely, rides with 0 candidates) is larger under equilibrium with MinRev (left panel) than with MinWeightRev. The red dashed curve is the probability distribution function of Poisson and its parameter is calculated based on the field approximation. on average.

panels of Figure 2, the red dashed line depicts the probability function of a Poisson distribution, whose mean is calculated using our analytic (mean field) model discussed in the following section. The strong similarity between the empirical histogram and the approximated Poisson probability function conforms to the intuition that the scaled locations of available drivers over the city area form a homogeneous spatial Poisson process, whose parameter is well-approximated using our mean field model's prediction.

Recall that the dynamics are such that a ride with 0 candidates is lost. By comparing the two histograms we observe that the proportion of lost rides in equilibrium under MinRev (left panel) is .26, whereas under MinWeightRev is .14. For this example we conclude that the increased equilibrium participation intensity under MinWeightRev significantly cuts down the demand loss by 46% compared to MinRev, resulting in 12% increase in the matching rate (from $.74 \times \lambda$ to $.86 \times \lambda$). In the paper, we show, based on our mean field analysis, that a factor 2 is an upper bound for the improvement in terms of the participation intensities and the matching rate.

3. Formulation, main result and relation to literature

In this section we introduce the baseline mean-field model that lays the groundwork for the subsequent comparison of the two policies. We present the relevant objects and the mean field dynamics that drive the progression of the system over time, and spell out the equilibrium conditions. We then state the papers' main result (Proposition 1) in terms of the mathematical objects and quantities defined.

To set up the model we introduce the primitives which are consistent with the notation and terminology introduced in Section 2. As explained, we assume a ring-shaped city represented by the real interval $[0,1)$, and two types of drivers, L and H, which differ by their opportunity cost rates. We denote by Θ_i the intensity of potential drivers of type $i \in \{L,H\}$, and let κ_i denote their opportunity cost rate, assuming w.l.o.g. that $\kappa_L < \kappa_H$. For convenience, we denote $\Theta = \Theta_L + \Theta_H$. Considering a type $i \in \{L,H\}$, we often use the term "i-drivers" to refer to drivers of that type, and we shall typically denote by j the opposite (i.e., complement) type, $j \in \{L,H\} \setminus \{i\}$.

The local intensity of passengers' arrivals is given by λ, which is equal across all locations in the city. In the mean field model, it is convenient to think about drivers and passengers as fluid particles. We let m denote the mean total duration of a ride (from the moment of the passenger's arrival until their drop off) and assume passengers' requested drop-off locations are uniformly distributed across the city. More precisely, suppose at a time epoch the amount of passengers riding in the system is given by y, then the instantaneous rate of passengers being dropped off at location $x \in [0,1)$ is $(y/m) \cdot dx$. The instantaneous total rate of revenue accumulation (of all busy drivers together) in that case is given by $y \cdot r$, where r is a positive constant. To avoid trivialities, we assume $\kappa_H < r$ (which also implies $\kappa_L < L$). This means that drivers are not expecting to make more revenue than they would have made if they were carrying passengers all the time without idling. Lastly, the parameter δ will be referred to as the *pickup range*.

Throughout the paper we will be interested in the (endogenous) quantity $\theta_i \in [0, \Theta_i]$ for each type $i \in \{L,H\}$ which represents the *effective* participation intensity of type-i drivers, namely the actual intensity of type-i drivers assuming a proportion $\frac{\theta_i}{\Theta_i}$ of them decide to participate. Thus, drivers' decision profile is characterized by a pair $(\theta_L, \theta_H) \in [0, \Theta_L] \times [0, \Theta_H]$. For a given pair $\{\theta_i\}_{i \in \{L,H\}}$ we shall typically denote $\theta = \theta_L + \theta_H$ and assume that $\theta > 0$. Our state space of choice is $\mathcal{Q}^2$, where $\mathcal{Q}$ is defined as the set of all differentiable, (weakly) increasing, positive real functions over the domain $[0, 1)$ mapping 0 to itself.

3.1. Mean field dynamics and equilibrium

We now describe the evolution of our mean field dynamical system through time as a set of differential equations, assuming that the participation intensities θ_L and θ_H (as well as θ) are known. To avoid using cumbersome notation, in some places we suppress the dependence on the participation intensities $\{\theta_i\}_{i \in \{L,H\}}$ and on the matching policy, yet all of the objects to be defined here should be understood as functions thereof.

At time $t \in [0, \infty)$, and for each type $i \in \{L, H\}$, let the function $Q_i(\cdot; t) \in \mathcal{Q}$ represent the proportion of type-i drivers (out of θ_i), who by time t were both available and located in the sub-interval $[0, x) \subseteq [0, 1)$. Thus, $Q_i(x; t)$ is a cumulative quantity, and for $x = 1$, $Q_i(1; t)$ expresses the fraction of available type-i drivers in the entire city, which in general is less than unity. Naturally, we assume $Q_i(0; t) = 0$ for all $t \in [0, \infty)$ and $i \in \{L, H\}$. The initial state $Q_i(\cdot; 0) \in \mathcal{Q}$ for each type $i \in \{L, H\}$ is assumed given. We use $Q_i'(x; t)$ to denote the derivative of $Q_i(x; t)$ with respect to x. Consistent with our interpretation of $Q_i(x; t)$, the term $\theta_i Q_i'(x; t)$ is understood as the density of available type-i drivers in the infinitesimal surrounding of location x.

Define

$$Q(x; t) = \frac{\theta_L}{\theta} Q_L(x; t) + \frac{\theta_H}{\theta} Q_H(x; t), \tag{1}$$

Then for each t, $Q(\cdot; t) \in \mathcal{Q}$, and it can be interpreted as the proportion of drivers, irrespective of types, who by time t were both available and located in the sub-interval $[0, x)$. Further define

$$R_i(t) = \frac{1}{t} \int_{u=0}^{t} r \cdot (1 - Q_i(1; u)) du, \tag{2}$$

Then $R_i(t)$ is the average revenue rate accumulated by type-i drivers, $i \in \{L, H\}$, calculated over a given (finite) time span, $[0, t]$. By contrast, when comparing the performance of the two policies we will be mainly interested in the stationary behavior of the system operating under each of the two different policies. Loosely speaking, the comparison of the policies hinges on the existence of a stationary system state, with the understanding that the latter reflects the limiting behavior of the system over a large time span (even though in the paper we do not prove such relation, but rather treat the stationary state as the solution concept of interest). Assuming such stationary state exists uniquely (this, in fact is proven in Lemma 2 for MinRev and in Lemma 5 for MinWeightRev) we define the stationary revenue rate R_i^* for each type $i \in \{L, H\}$ as the unique solution

$$\frac{d}{dt} R_i(t) = 0. \tag{3}$$

Not surprisingly, given the participation intensities in the MinRev policy, types do not play any role in the matching decision, and it is anticipated that $R_L^* = R_H^*$. The stationary effective arrival rate (or equivalently, the matching rate) to the system, λ^*, is defined as the unique solution of the system

$$\frac{d}{dt} \lambda(t) = 0, \tag{4}$$

where

$$\lambda(t) = \frac{1}{t} \int_{u=0}^{t} \frac{\theta}{m} (1 - Q(1; u)) du, \tag{5}$$

noting that these relations, together with (1) and (2) imply

$$\lambda^* = \frac{\theta_L R_L^* + \theta_H R_H^*}{rm}.$$

In the analysis, we use the same state representation to study both the MinRev and the MinWeightRev policies. However, in order to avoid confusion, when we discuss the MinWeightRev policy we shall use the "hat" indication $(\hat{\cdot})$, e.g., the notation Q_i will be replaced by $\hat{Q}_i$.

Collective revenue and dimension reduction Noticeably, our state space of choice $\mathcal{Q}^2$ does not account for the revenues accrued by individual drivers over time. This follows the observation that when the system is in its stationary state, symmetric drivers should all be accumulating revenue at the exact same rate (which is constant over time). Both policies are such that participating drivers of the same type are considered symmetric in the eyes of the platform, thus, for the purpose of analyzing the system in steady state it is enough to keep track of the stationary average revenue of type i, which is captured by the term R_i^*, $i \in \{L, H\}$.

With this in mind, setting about the mean field formulation we deviate from the discrete model described in Section 2 by slightly modifying the policy as follows: Instead of keeping track of each driver's revenue over time, we will *assume* that the accumulated revenue at each time epoch t is the same across all type-i drivers, and equals $R_i(t)$. This is done by grouping drivers of the same type and letting them work as a collective, in which the profit generated from each ride is shared evenly across all drivers. This modification promotes a substantial reduction of the state space dimension yet does not affect the stationary performance of the system.

In particular, the collective approach implies that when multiple drivers of different types become candidates for a ride, determining the matched driver's type is done by comparing only two values: $R_L(t)$ vs. $R_H(t)$ under MinRev, and similarly, $\hat{R}_L(t)/\kappa_L$ vs. $\hat{R}_H(t)/\kappa_H$ under MinWeightRev. This means that as long as one value is strictly larger than the other, the type with the smaller value will obtain strict priority in the matching over the other type. When these values are equal we resolve the tie using the concept of partial (or "probabilistic") priority: We let $S_i(t)$ denote the cumulative time until time epoch t that type i was granted a matching priority. If at time t, no type qualifies for strict priority then each type $i \in \{L, H\}$ is granted partial priority with proportion $S_i(t)/(S_L(t) + S_H(t))$. In the discrete setup, this corresponds to randomly selecting the prioritized type, with probability for each type equal to the historic proportion of time that type was prioritized. Note that assuming both types start with 0 revenue, this mechanism requires that we set initial values for $S_L(0)$ and $S_H(0)$.

MinRev policy The dynamics of the system operating under MinRev obey the following set of equations, for all $x \in [0,1)$, $t \in [0,\infty)$, and $i \in \{L,H\}$,

$$\frac{\partial Q_i(x;t)}{\partial t} = \frac{x}{m}(1 - Q_i(1;t)) - \frac{\lambda}{\theta_i}\int_{s=0}^{x}\left(1 - e^{-Q_i'(s;t)\theta_i\delta}\right) D_i(s;t)ds, \tag{6}$$

where

$$D_i(s;t) = e^{-Q_j'(s;t)\theta_j\delta} + \frac{dS_i(t)}{dt}\cdot\left(1 - e^{-Q_j'(s;t)\theta_j\delta}\right)$$

and

$$\frac{dS_i(t)}{dt} = \mathbf{1}\left(R_i(t) < R_j(t)\right) + \mathbf{1}\left(R_i(t) = R_j(t)\right)\cdot\frac{S_i(t)}{S_L(t) + S_H(t)}.$$

MinWeightRev policy The dynamics of the system operating under MinWightRev obey the following set of equations, for all $x \in [0,1)$, $t \in [0,\infty)$, and $i,j \in \{L,H\}$ s.t. $i \neq j$,

$$\frac{\partial \hat{Q}_i(x;t)}{\partial t} = \frac{x}{m}(1 - \hat{Q}_i(1;t)) - \frac{\lambda}{\theta_i}\int_{s=0}^{x}\left(1 - e^{-\hat{Q}_i'(s;t)\theta_i\delta}\right) \hat{D}_i(s;t)ds, \tag{7}$$

where

$$\hat{D}_i(s;t) = e^{-\hat{Q}_j'(s;t)\theta_j\delta} + \frac{d\hat{S}_i(t)}{dt}\cdot\left(1 - e^{-\hat{Q}_j'(s;t)\theta_j\delta}\right)$$

and

$$\frac{d\hat{S}_i(t)}{dt} = \mathbf{1}\left(\frac{\hat{R}_i(t)}{\kappa_i} < \frac{\hat{R}_j(t)}{\kappa_j}\right) + \mathbf{1}\left(\frac{\hat{R}_i(t)}{\kappa_i} = \frac{\hat{R}_j(t)}{\kappa_j}\right)\cdot\frac{\hat{S}_i(t)}{\hat{S}_L(t) + \hat{S}_H(t)}.$$

Equation (6) is a special case of (7) for $\kappa_i = \kappa_j$. An intuitive explanation for the drift terms in Equations (6) and (7) is given in Section 3.3. It should be mentioned that systems (6) and (7) fall into the category of non-smooth dynamical systems (see Kunze (2000)), and therefore their rigorous formulation is of the form of a differential inclusion problem. These technicalities are dealt with in more detail in the appendix (see APX-A).

Remarkably, given $\{\theta_i\}_{i\in\{L,H\}}$, we note that by summing up the terms $\theta_i/\theta \cdot \partial\hat{Q}_i(x;t)/\partial t$ in Equation (7) (and similarly the terms $\theta_i/\theta \cdot \partial Q_i(x;t)/\partial t$ in Equation (6)) for $i = L,H$, using the relation in (1) and the fact that $(d/dt)(\hat{S}_L + \hat{S}_H) = 1$, we produce a differential equation

$$\frac{\partial Q(x;t)}{\partial t} = \frac{x}{m}(1 - Q(1;t)) - \frac{\lambda}{\theta}\int_{s=0}^{x}\left(1 - e^{-Q'(s;t)\theta\delta}\right) ds. \tag{8}$$

This suggests that when the participation intensities $\{\theta_i\}_{i\in\{L,H\}}$ are given, the dynamics of the process $Q(x;t)$, and therefore the term λ^* as defined in (4), do not depend on the matching policy. It should be mentioned however that in equilibrium the terms $\{\theta_i\}_{i\in\{L,H\}}$ are endogenously formed as a consequence of drivers' strategic considerations, on which the matching policy has significant impact.

Equilibrium Conditions In equilibrium, no driver can increase their revenue by deviating from their decision. To stress the dependence of the stationary (per-type) revenue rate (defined in (3)) on the pair of participation intensities, we rewrite the former as a function of the latter, namely, $R_i^*(\theta_L,\theta_H)$ for each $i\in\{L,H\}$. Given the platforms' matching policy, a pair of participation intensities $\{\theta_i\}_{i\in\{L,H\}}$ is said to induce Nash Equilibrium in the mean field system, if for each $i\in\{L,H\}$ it satisfies (jointly):

$$\begin{cases} R_i^*(\theta_L,\theta_H) < \kappa_i \quad \Rightarrow \quad \theta_i = 0, \\ \\ R_i^*(\theta_L,\theta_H) > \kappa_i \quad \Rightarrow \quad \theta_i = \Theta_i. \end{cases} \tag{9}$$

The condition is identical for both policies (replacing R_i^* by $\hat{R}_i^*$ for MinWeightRev). We will also be interested in the stationary effective arrival rate (defined in (4)) resulting from each pair $\{\theta_i\}_{i\in\{L,H\}}$, thus, we write it as $\lambda^*(\theta)=\lambda^*(\theta_L+\theta_H)$.

3.2. Main result

The key findings of this work are summarized in the following Proposition 1, whose proof follows the analysis elaborated in the subsequent sections. This proposition argues that for each policy there exists a unique equilibrium, and provides bounds on the improvement of the equilibrium performance of the MinWeightRev policy compared to that of MinRev.

PROPOSITION 1. *For each given pair $\{\theta_i\}_{i\in\{L,H\}}$, the system defined in (6) with any initial state $\{Q_i(\cdot;0),S_i(0)\}_{i\in\{L,H\}}\in\{\mathcal{Q}\times\mathbf{R}_+\}^2$, and similarly the one in (7) with any initial state $\{\hat{Q}_i(\cdot;0),\hat{S}_i(0)\}_{i\in\{L,H\}}\in\{\mathcal{Q}\times\mathbf{R}_+\}^2$, admit a unique solution, and therefore are well defined. Additionally, the stationary revenues $\{R_i^*(\theta_L,\theta_H)\}_{i\in\{L,H\}}$ and $\{\hat{R}_i^*(\theta_L,\theta_H)\}_{i\in\{L,H\}}$, and the stationary* *effective matching rate* $\lambda^*(\theta_L + \theta_H)$*, all exist uniquely, thus, the equilibrium condition in (9) is well defined.*

Moreover, there exists a unique pair $\{\theta_i^e\}_{i\in\{L,H\}}$ *satisfying the equilibrium condition (9) for MinRev, and similarly, a unique pair* $\{\hat{\theta}_i^e\}_{i\in\{L,H\}}$ *exists satisfying (9) for MinWeightRev. Denote the* improvement ratios *in terms of the equilibrium participation intensities and effective arrival rate by*

$$\Phi = \frac{\hat{\theta}_L^e + \hat{\theta}_H^e}{\theta_L^e + \theta_H^e} \quad \text{and} \quad \Psi = \frac{\lambda^*(\hat{\theta}_L^e + \hat{\theta}_H^e)}{\lambda^*(\theta_L^e + \theta_H^e)}, \tag{10}$$

respectively. Then the equilibrium pairs corresponding with the two policies satisfy

1. $\theta_H^e > 0 \Rightarrow \theta_L^e = \Theta_L$ *and similarly,* $\hat{\theta}_H^e > 0 \Rightarrow \hat{\theta}_L^e = \Theta_L$
2. $\theta_L^e = \hat{\theta}_L^e$ *and* $\theta_H^e \leq \hat{\theta}_H^e$*;*
3. $\Phi \in [1, 2]$*, with* $\Phi = 2$ *if and only if* $R_L^*(\theta_L^e, \theta_H^e) = \kappa_H$ *and* $\hat{R}_L^*(\hat{\theta}_L^e, \hat{\theta}_H^e) \geq \kappa_L$*;*
4. $\Psi \in [1, 2)$*, with* $\lim_{\Theta\to 0} \Psi = 2$ *if and only if* $\lim_{\Theta\to 0} \Phi = 2$*.*

Discussion Item (1) in Proposition 1 tells us that both policies are such that in equilibrium, H-drivers participate only when all L-drivers participate. In MinRev, this behavior is expected due to drivers being treated in a symmetric manner, thereby receiving the same long-run revenue rate, which in turn makes MinRev more attractive to drivers with lower opportunity cost. The same conclusion is also true for MinWeightRev, however the explanation is less intuitive, because symmetry among drivers breaks after normalizing the revenue rates. We discuss this in more details in Section 5.2.

Given a pair of participation intensities, all other factors being equal, L-drivers are worse off under MinWeightRev compared to MinRev. However, item (2) interestingly states that when changing the policy from MinRev to MinWeightRev, the reduction in revenue rate that L-drivers incur is not too substantial so as to make less of them want to participate in equilibrium. In contrast, H-drivers do benefit from such a change, hence, more of them tend to participate. Consistent with item (3), it means that the participation intensity can only increase, i.e., $\Phi \geq 1$, which further implies that $\Psi \geq 1$ as argued in (4). Items (3) and (4) further provide bounds on how large the improvement is, in terms of the total participation intensity and in terms of the matching rate, and assert that these bounds are tight. In other words, depending on the parameters, MinWeightRev has the potential of processing twice as much throughput as MinRev, due to increased participation, assuming drivers in both policies are in equilibrium.

3.3. Intuitive explanation of mean field equations

We shall now provide some intuitive justification for the formulation of the two mean field systems defined in Equations (6) and (7). Due to their similarity, we shall focus in the explanation only on Equation (6) (namely the MinRev policy), and the same explanation will follow also for (7) by replacing the terms $R_i(t)$ with $\hat{R}_i(t)/\kappa_i$.

For any type $i \in \{L, H\}$, consider the positive drift term in Equation (6), $(1 - Q_i(1;t))\frac{x}{m}$. The term $1 - Q_i(1;t)$ relates to the type-i proportion of busy drivers at time t. Busy drivers complete their rides at rate $\frac{1}{m}$, after which they become available at the ride's drop-off location which is uniformly distributed. Hence, the probability that such busy driver will become available in a location in $[0, x)$ is x. Therefore, the total rate at which busy type-i drivers become available in $[0, x)$ is $(1 - Q_i(1;t))\frac{x}{m}$.

To explain the negative term in (6), and specifically the intuition that leads to the terms $\left(1 - e^{-Q_i'(s;t)\theta_i\delta}\right)$ and $D_i(s;t)$ in the integral, we consider the motivating discrete model described in Section 2. The term $\lambda \cdot \left(1 - e^{-Q_i'(s;t)\theta_i\delta}\right) D_i(s;t)$ captures the instantaneous matching rate of type-i drivers at time t around location s. In the discrete model, the term $1 - e^{-Q_i'(s;t)\theta_i\delta}$ is analogous to the probability that a passenger arriving at time t and location s will face at least one type-i candidate, and conditioned on that, $D_i(s;t)$ is the probability that this passenger will indeed be matched with a type-i driver. The explanation relies on the construction below.

Recall that in the discrete model, for a market of finite size N, the total number of drivers in the system is given by θN. Assume that at time t, the state of the system is given by $\{Q_i(\cdot;t)\}_{i\in\{L,H\}}$. Consider the set of of type-i available drivers, and suppose their locations are iid continuous random variables with cdf $\frac{Q_i(\cdot;t)}{Q_i(1;t)}$. Provided a ride request arrives at time t with pickup location s, the probability that any type-i driver is available within the pickup region $s \pm \frac{\delta}{2N}$ is given by

$$
\begin{aligned}
1-&\left(1-\frac{Q_i(s+\frac{\delta}{2N};t)-Q_i(s-\frac{\delta}{2N};t)}{Q_i(1;t)}\right)^{\theta_i N Q_i(1;t)} \\
&\cong 1-\left(1-\frac{Q_i(s;t)+Q_i'(s;t)\frac{\delta}{2N}-\left(Q_i(s;t)-Q_i'(s;t)\frac{\delta}{2N}\right)}{Q_i(1;t)}\right)^{\theta_i N Q_i(1;t)} \\
&= 1-\left(1-\frac{Q_i'(s;t)\delta}{NQ_i(1;t)}\right)^{\theta_i N Q_i(1;t)} \xrightarrow[N\to\infty]{} 1-e^{-Q_i'(s;t)\theta_i\delta}
\end{aligned}
$$

where the approximation follows the first order Taylor's expansion for Q_i as a function of x around the point s. Hence, $1-e^{-Q_i'(s;t)\theta_i\delta}$ is the probability that at least one type-i driver is a candidate for that ride. Furthermore, as an implication of Le Cam's theorem we have that for large N, the number of type-i candidates in the passenger's patience region is asymptotically Poisson distributed with parameter $\theta_i Q_i'(s;t)\delta$, independent of the number of candidates of the other type, j, $j \neq i$. Conditioned on the pickup region containing at least one available type-i driver, the probability the ride will be matched with a type-i driver is captured through $D_i(s;t)$. This conditional probability depends on type-i's priority status at time t, which we defined as $S_i(t)$. We partition into three complement cases:

- If $R_i(t) < R_j(t)$, then $dS_i(t)/dt = 1$ and type-i are granted priority over a type-j in the matching. This means that conditioned on a type-i candidate being available for a ride, the ride will be assigned to a type-i driver with probability 1.
- If $R_i(t) > R_j(t)$, then $dS_i(t)/dt = 0$. Thus, type-i drivers are only matched with a ride when no type-j candidates are around, and the latter event happens with probability $e^{-\hat{Q}_j'(s;t)\theta_j\delta}$.
- Lastly, when $R_i(t) = R_j(t)$, the priority is determined based on a random selection, with probability $S_i(t)/(S_L(t)+S_H(t))$ being in favor of type i.

Because passengers arrive with rate λ to the system at uniformly distributed locations, the instantaneous rate of rides matched to drivers of type i located in $[0,x)$ is $\lambda\int_{s=0}^{x} D_i(s;t)ds$. Equivalently, it is also the rate at which type-i drivers in $[0,x)$ change their state from available to busy. Dividing by θ_i then gives us the rate proportion of those out of the total population of drivers.

In the next sections we execute the analysis for the two policies based on the mean field formulation discussed here.

4. Modeling approach and related literature

The literature studying ride sharing platforms from an OR-MS perspective is extensive. The lion's share in this growing line of research is focused (either purely or jointly) on pricing the service, relocating drivers, and matching drivers and riders. Given the variety of existing models it is perhaps questionable whether adopting a new modeling approach will yield any significant contribution. The purpose of this section is to give a brief survey of existing relevant work, with emphasis on queueing networks models, and to justify our approach by explaining how it differs from the one taken in other papers. Extensive surveys of the OM literature with applications to the sharing economy is given in Benjaafar and Hu (2020) and in Hu (2019), chapters ? and ?. A summarized tutorial to that stream of literature is provided in Hu (2020).

In terms of modeling time-driven dynamics, probably the closest approach to ours is the one in Braverman et al. (2019), where the authors consider the problem of optimal relocation of empty cars. Their underlying model is a closed BCMP queueing network (which is similar to that Iglesias et al. (2019)) and its mean field (fluid) limit can be interpreted as a closed network of finitely-many nodes in which fluid particles (drivers) flow in circulation. In this model the distribution of available drivers across locations is atomic (as opposed to the non-atomic one here), and a passenger can only be picked up by a driver residing in their location. Consider for example the case where locations in the network are symmetric. Then, depending on the arrival rate, either all drivers are constantly busy, or all the demand is filled. This conclusion, which is not limited to the symmetric setup, is similar in essence to the one drawn by observing the naive model described in the introduction (see Section 1). The naive model is in fact the special, single-node case of the Braverman et al. (2019) model. Thus, albeit useful for studying empty car routing policies, this modeling approach is not appropriate for understanding the friction between pickup range and utilization.

Contrary to Braverman et al. (2019), Ozkan and Ward (2017) adapt a different approach, modeling a ride sharing system as an open queueing network. Rather than circulating the network, drivers arrive sequentially at different locations (represented by the different queues), and leave the system once matched or due to abandonment, whichever comes first. It is assumed that drivers' arrivals to each location form an independent stochastic process that does not depend on the system state. Service times correspond to passengers' arrivals, and pickup range is taken into account by allowing each passenger to be matched with a driver residing in one of several different sufficiently-close locations. The authors show how the fluid counterpart of their model can be used to devise and study the performance of sophisticated matching policies, some of which obtain asymptotic optimality (in terms of matchings produced). Unfortunately, since the model ignores ride duration and views drivers as short-lived entities, it does not offer a natural evaluation of variables such as drivers' long-run busy fraction and revenue rate, which are detrimental to our discussion of equilibrium. A similar modeling approach is followed also in Özkan (2020) to study joint pricing and matching decisions.

A variant of Braverman et al. (2019) is studied in Afeche et al. (2018), further accounting for drivers' strategic consideration: drivers can make participation as well as relocating decisions. The assumptions made in Afeche et al. (2018) regarding equilibrium participation of drivers is similar to those in here. Other similar-in-spirit models that consider strategic drivers in a queueing (or queueing-like) system include: Banerjee et al. (2015), in which the network is open (drivers may exit and enter) with a single location, and passengers arrivals depend on the price which is dynamically set by the platform; Taylor (2018) who studies the impact of passenger delay sensitivity and driver independence on prices and wages in a single queue setup; Bai et al. (2019), that build on Taylor (2018) to gain managerial insights on the interplay between prices, driver capacity and passenger waiting cost; and Bimpikis et al. (2019), who consider spatial pricing and drivers compensation as tools for balancing supply and demand over the (fluid) network, in a discrete time setting with price-sensitive passengers. All of the papers listed above do not model pickup range. Even though the underlying model in this paper is not a classical queueing model per se, it is closely related and inspired by this theory, and the strategic considerations of drivers raises strong connections to the literature about strategic queueing. An exhaustive survey of the theory and literature about queuing models with strategic agents can be found in Hassin and Haviv (2003) and Hassin (2016).

Relevant papers that depart from the queueing literature with applications to ride sharing platforms include Besbes et al. (2020) who study a two-stage Stackelberg game of a platform setting location-dependent ride prices, and drivers who react to these prices by making repositioning decisions. The authors employ a general spatial network structure that allows for both atomic as well as non atomic distributions of drivers across the city locations. Similar to Besbes et al. (2020), Cachon et al. (2017) also consider a Stackelberg game setting to compare different pricing schemes and employment contracts with respect to the firm's profit, drivers wage and passengers welfare. In a newsvendor setting, Gurvich et al. (2019) consider the problem faced by a firm, such as a ride sharing company, who manage its workforce, and self-scheduling strategic agents choose whether or not to work for that firm.

As briefly mentioned in Section 3, the discrete model introduced in Section 2 is intractable. Indeed, when only one driver participates, the stochastic process describing the availability status of the driver is that of an M/G/1/1 queue. The spatial component describing the driver's location, and the driver's revenue accumulation over time play no significant role in the analysis. Increasing the number of drivers beyond one immediately complicates the analysis to the point of intractability, mainly due to the dependencies that arise between drivers' availability, spatial location, revenue accumulation and effective matching processes. We therefore perform the analysis in this paper on a mean field model that well-approximates the discrete one in a large market regime.

Mean field models for queueing networks in relation with their stochastic counterparts have been studied for instance by Tsitsiklis and Xu (2012), and by Xu and Hajek (2013) who further account for customers' strategic decision making. One challenging technical aspect of the model considered here is that the mean field dynamics are formulated as a non-smooth dynamical system (i.e., having a discontinuous drift), a challenge that is also dealt with in Tsitsiklis and Xu (2012). The rigorous way to treat dynamical systems of that kind is by reformulating them as a differential inclusion problem. Gast and Gaujal (2012) provide a concise introduction to differential inclusions and lay out a unified frame work for studying mean field limits of finite-dimensional queueing models with discontinuous drifts. For technical information about differential inclusion and non-smooth dynamical systems we refer the interested reader to Kunze (2000).

5. Analysis of mean field model

In this section we prove and discuss the first part of Proposition 1, that relates to the validity of our definition of the mean field system, and to the existence and uniqueness of equilibrium participation intensities, for each of the two policies. The analysis in this Section is carried separately for each policy, starting with MinRev.

5.1. The MinRev policy

The first milestone in the analysis of the mean field model is proving that the differential equations defined in (6) and (7) for a given pair of participation intensities $\{\theta_i\}_{i\in\{L,H\}}$ indeed define a unique mapping of time to system state. We state this in the following result:

LEMMA 1. *For every pair of participation intensities $\{\theta_i\}_{i\in\{L,H\}}$, and every pair of initial states $\{Q_i^0, S_i^0\}_{i\in\{L,H\}} \in \{\mathcal{Q}\times\mathbf{R}_+\}^2$, there exists a unique pair of mappings $\{Q_i\}_{i\in\{L,H\}}$, $Q_i:\mathbf{R}_+\to\mathcal{Q}$, with a corresponding pair of real functions $\{S_i\}_{i\in\{L,H\}}$, $S_i:\mathbf{R}_+\to\mathbf{R}_+$, that solve the initial value problem defined by Equation (6) for each $i\in\{L,H\}$ with initial value $Q_i(\cdot;0)=Q_i^0$, and $S_i(0)=S_i^0$.*

The proof is in Appendix APX-A.

Lemma 1 implies that our dynamical system is well defined given any initial state $\{Q_i^0, S_i^0\}_{i\in\{L,H\}}$. As a corollary we also have that for an initial state $Q^0=(\theta_L/\theta)Q_L^0+(\theta_H/\theta)Q_H^0$, there exists a unique solution $Q(x;t)$ for the problem defined in Equation (8). The next lemma states that this system of functions Q, Q_L and Q_H admits a unique steady state, which can be easily characterized. To simplify notation, when referring to steady state we omit the t argument in $Q_i(x;t)$ and $Q_i'(x;t)$ by writing $Q_i(x)$ and $Q_i'(x)$.

LEMMA 2. *The unique solution $\{Q_i^*\}_{i\in\{L,H\}}\in\mathcal{Q}^2$ to the mean field steady-state equations*

$$x\cdot(1-Q_i(1))\frac{\theta_i}{m}-\lambda\int_{s=0}^{x}D_i(s)ds=0, \qquad \forall x\in[0,1), i\in\{L,H\}, \tag{11}$$

is given by $Q_L^*(x)=Q_H^*(x)=qx$, where $q\in[0,1]$ is the unique solution to the (single variable) transcendental equation

$$(1-q)\frac{\theta}{m}-\lambda\cdot\left(1-e^{-q\theta\delta}\right)=0. \tag{12}$$

The corresponding stationary prioritization functions are given by $S_L^*(t) = s_L$ *and* $\hat{S}_H^*(t) = s_H$*, where for each* $i, j \in \{L, H\}$*,* $i \neq j$*,*

$$s_i = \frac{\theta_i}{\theta(1 - e^{-q\theta_i\delta})} - \frac{\theta_j e^{-q\theta_j\delta}}{\theta(1 - e^{-q\theta_j\delta})} \tag{13}$$

The proof is in Appendix APX-B.

From Lemma 2 it can be seen that in steady state, the functions $Q_L^{*'}(x) = Q_H^{*'}(x) = q$ are constant, thus, the density of available drivers at each location in the city is the same and equals θq. It also follows that by defining $Q^* = (\theta_L/\theta)Q_L^* + (\theta_H/\theta)Q_H^*$ as in Equation (8) we obtain $Q^*(x) = qx(= Q_i^*(x))$.

We notice that the value of q depends on the joining rates $\{\theta_i\}_{i\in\{L,H\}}$ only through their sum, θ. This is because MinRev does not differentiate between drivers of different types, thus the density of available drivers across the city depends on the total number of drivers but not on how they split into different types.

The constant q can be interpreted as the proportion of available drivers in the city in steady state, or, alternatively, the stationary proportion of time each driver spends being available. Interestingly, Equation (12) is related to the system's steady-state performance, in that it reflects a variant of Little's Law: If the long run proportion of available drivers is q, then the time-averaged number of passengers being served is $(1-q)\theta$, each of which remains in the system for m units of time on average. By Little's Law, $(1-q)\frac{\theta}{m}$ must be equal to passengers' effective arrival rate (namely the matching rate), which is given by the potential arrival rate λ, multiplied by the long-run proportion of accepted rides, $1 - e^{q\theta\delta}$, which is the same for all pickup locations.

Recall our definition of the long-run revenue R_i^*, as defined in (3). While busy, a driver generates revenue at rate $r > 0$. Consistent with our interpretation of q as the long-run availability fraction, we have that $R_L^* = R_H^* = (1-q)r$.

Equilibrium under MinRev For the purpose of performing equilibrium analysis, with a slight abuse of notation we shall refer to the steady-state proportion of available drivers q as a function of the total participation intensity θ, and denote it by $q(\theta)$. We also reintroduce the long-run revenue rate, $R^*(\theta) = (1-q(\theta))r$, which as previously explained is equal for both drivers' types under MinRev, namely $R_i^*(\theta_L, \theta_H) = R^*(\theta)$ for $i \in \{L, H\}$, assuming $\theta = \theta_L + \theta_H$.

The next property of $q(\theta)$ plays an important roll in the discussion about uniqueness of the equilibrium solution:

Lemma 3. *The steady-state proportion of available drivers, $q(\theta)$ is increasing in θ.*

The proof is in Appendix APX-C.

We note, since the long-run revenue rate R^* is decreasing with $q(\theta)$, by Lemma 3 it is also decreasing with θ.

Rearranging Equation (12) and taking the limit as $\theta \to 0$ we get

$$1 - q(0) = \lambda m \cdot \lim_{\theta \to 0} \frac{1 - e^{q(0)\theta\delta}}{\theta} = \lambda m \cdot q(0)\delta$$

from which we derive a lower bound for the availability fraction, $\underline{q} \equiv q(0) = 1/(1+\lambda\delta m)$, and therefore also an upper bound for the revenue rate per driver $\overline{R} = (1-\underline{q})r = \lambda\delta m r/(1+\lambda\delta m)$. Being the highest revenue rate possible, $\overline{R}$ can be interpreted as the revenue rate of a driver in the hypothetical case where no other driver joins the system. Intuitively, in the discrete model with finite N, such driver, when available, will be matched with the first ride that arrives within their pickup surrounding, which is of size δ/N. The average time until the arrival of such ride is $\lambda N \cdot \delta/N = (\lambda\delta)^{-1}$. Hence, the expected time this driver spends in a cycle between subsequent jobs consists of the mean availability duration, $(\lambda\delta)^{-1}$, and the total ride's duration m. Their revenue rate is therefore given by $r \cdot m/((\lambda\delta)^{-1} + m) = \overline{R}$.

To avoid trivialities when dealing with equilibrium joining rates, we make the assumption that $\kappa_L < \kappa_H < \overline{R}$, which ensures that a non-zero proportion of drivers choose to join in equilibrium.

Recall that $\kappa_L < \kappa_H$, thus, substituting R^* for both types in the equilibrium condition (9), it is clear under MinRev, any type-H driver decides to participate, then all type-L drivers should participate too. It is therefore natural to characterize drivers' equilibrium participation intensity using a single value, $\theta^e \in (0, \Theta_L + \Theta_H]$, such that if $\theta^e \le \Theta_L$, only type-L drivers join, and if $\theta^e > \Theta_L$, all type-L join in addition to $\theta^e - \Theta_L$ type-H drivers.

Proposition 2. *The equilibrium participation intensity under MinRev, denoted by θ^e, is unique and satisfies exactly one of the following:*

1. $\theta^e < \Theta_L$ *and* $R^*(\theta^e) = \kappa_L$
2. $\theta^e = \Theta_L$ *and* $\kappa_L \le R^*(\theta^e) \le \kappa_H$
3. $\theta^e \in (\Theta_L, \Theta_L + \Theta_H]$ *and* $R^*(\theta^e) \ge \kappa_H$

The proof is in Appendix APX-D.

Proposition 2 stems from the monotone behavior of the revenue rate R^* with respect to θ. The revenue rate decreasing in θ implies that the game dynamics are of the Avoid-the-Crowed (ATC) kind – the more drivers participating the less beneficial it becomes to do so. This observation may seem obvious, due to the fact that when more drivers participate, each driver's share of the total revenue decreases. While being true, this argument alone does not adequately rationalize the ATC dynamics, because the total revenue generated in the system, which is proportional to the matching rate $1 - e^{q\theta\delta}$, is increasing in θ. Nevertheless, it is implied by Lemma 3 that the negative impact of θ on the per-driver revenue rate is dominating. In the discrete setting with finite N, this is intuitive: As long as two drivers are located δ/N units away from each other, they do not interact and therefore are indifferent to the existence of one another. But once they get closer than δ/N units from each other, it is possible that the two drivers will have to compete for the same ride, in which case it is of interest for each one of them that the other would not participate.

The three different cases in Proposition 2 correspond with each of the three scenarios respectively: (1) The system under equilibrium can generate only enough revenue so that some type-L drivers decide to join, these drivers are indifferent between participating or not, and all type-H drivers are better off not working for the platform; (2) There is enough revenue in equilibrium to attract all type-L drivers, but not enough to attract type-H drivers; (3) All type-L drivers (strictly) prefer to participate, and some (possibly all) type-H drivers participate too.

5.2. The MinWeightRev policy

In this section we study the the MinWeightRev policy, to derive results paralleling Lemmas 1–2 and Proposition 2 as obtained for MinRev. Once again we assume at first we are given a fixed pair of participation intensities, and later on we solve for the endogenous participation intensities that induce equilibrium.

LEMMA 4. *For every pair of participation intensities $\{\theta_i\}_{i\in\{L,H\}}$, and every pair of initial states $(\hat{Q}_L^0, \hat{Q}_H^0) \in \mathcal{Q}^2$, there exists a unique pair of functions $\hat{Q}_L(x;t)$ and $\hat{Q}_H(x;t)$ solving the initial value problem defined by Equation (7) with initial value $Q_i(x;0) = Q_i^0(x)$ for each $i \in \{L, H\}$.*

The proof is in Appendix APX-E.

Similarly to MinRev, given a pair of initial states $(\hat{Q}_L^0, \hat{Q}_H^0)$ we define $\hat{Q}^0 = (\theta_L/\theta)\hat{Q}_L^0 + (\theta_H/\theta)\hat{Q}_H^0$, and $\hat{Q}(x;t)$ as the unique solution for the problem defined in Equation (8). Next we characterize the invariant state (omitting the argument t as done in Lemma 2).

LEMMA 5. *The unique stationary solution $(\hat{Q}_L^*, \hat{Q}_H^*) \in \mathcal{Q}^2$ to the mean field steady-state equations*

$$x \cdot (1 - \hat{Q}_i(1))\frac{\theta_i}{m} - \lambda \int_{s=0}^{x} \hat{D}_i(s) ds = 0, \qquad \forall x \in [0,1), i \in \{L, H\}, \tag{14}$$

is given by $\hat{Q}_L^*(x) = \hat{q}_L x$ and $\hat{Q}_H^*(x) = \hat{q}_H x$, where $(\hat{q}_L, \hat{q}_H, \hat{q}) \in [0,1]^3$ uniquely solve the system

$$\begin{cases} \theta_L \hat{q}_L + \theta_H \hat{q}_H = \theta \hat{q}, \\ 1 - \hat{q} = \lambda m \cdot (1 - e^{-\hat{q}\theta\delta})/\theta, \\ 1 - \hat{q}_H = \min\left\{(1-\hat{q}_L)\kappa_H/\kappa_L,\, \lambda m \cdot (1 - e^{-\hat{q}_H\theta_H\delta})/\theta_H\right\}. \end{cases} \tag{15}$$

The corresponding stationary prioritization functions are given by $\hat{S}_L^*(t) = \hat{s}_L$ and $\hat{S}_H^*(t) = \hat{s}_H$, where $(\hat{s}_L, \hat{s}_H) \in [0,1]^2$ satisfy $\hat{s}_L = 1 - \hat{s}_H$, and

$$\hat{s}_H = \begin{cases} \frac{\alpha}{1 - e^{-\hat{q}_H\theta_H\delta}} - \frac{(1-\alpha)e^{-\hat{q}_L\theta_L\delta}}{1 - e^{-\hat{q}_L\theta_L\delta}} & \text{if } \frac{1-\hat{q}_H}{\kappa_H} = \frac{1-\hat{q}_L}{\kappa_L}, \\ 1 & \text{otherwise}, \end{cases} \tag{16}$$

with

$$\alpha = \frac{\theta_H\kappa_H}{\theta_L\kappa_L + \theta_H\kappa_H}.$$

The proof is in Appendix APX-F.

Observing the first equation in (15), it follows that by defining $\hat{Q}^* = (\theta_L/\theta)\hat{Q}^*_L + (\theta_H/\theta)\hat{Q}^*_H$ we have that $\hat{Q}^*(x) = \hat{q}x$. We also note that the second equation in (15) is equivalent to Equation (1), thus the value of $\hat{q}$ (and hence $\hat{Q}^*$) depends on the participation intensities only through their sum, θ. This is because MinWeightRev is a "work conserving" policy, in the sense that it always matches a passenger if a driver is available within the pickup region. Thus the proportion of busy drivers (irrespective of types) depends on the total participation intensity θ but not on how it is partitioned.

The third equation in (15) is slightly more intricate. Each term in the choice set of the minimization in the right hand side corresponds to a different implication:

1. If $1 - \hat{q}_H = (1 - \hat{q}_L)\kappa_H/\kappa_L$, then

$$\frac{\hat{R}^*_H}{\kappa_H} = \frac{1 - \hat{q}_H}{\kappa_H} r = \frac{1 - \hat{q}_L}{\kappa_L} r = \frac{\hat{R}^*_L}{\kappa_L}.$$

This means that in stationarity, the weighted revenues of both types are equal, thus the busy fractions of the two types are equally proportional to their opportunity costs.

2. If $1 - \hat{q}_H = \lambda m \cdot (1 - e^{-\hat{q}_H \theta_H \delta}) / \theta_H$, then by rearranging we have

$$(1 - \hat{q}_H)\frac{\theta_H}{m} - \lambda \cdot \left(1 - e^{-\hat{q}_H \theta_H \delta}\right) = 0,$$

hence the solution $\hat{q}_H$ for the given participation intensities (θ_L, θ_H) is similar to the solution q defined by Equation (12) under MinRev given the participation intensities $(0, \theta_H)$. This can be interpreted as a situation in which H-drivers are granted priority over L-drivers in the matching. In other words, there exists some time epoch $\tilde{t}$ such that for all $t > \tilde{t}$, $\hat{R}_H(t)/\kappa_H < \hat{R}_L(t)/\kappa_L$, which, by Equation 7 implies the $\hat{D}_H(s;t) = 1 - e^{-\hat{Q}'_H(s;t)\theta_H \delta}$. It further implies that $\hat{R}^*_H/\kappa_H < \hat{R}^*_L/\kappa_L$.

MinWeightRev is designed so as to reduce the imbalances between drivers' actual revenues and their income goals, by matching a ride with the candidate whose accumulated revenue is the furthest away from their goal. These imbalances arise due to the heterogeneity in opportunity cost rates (and in the discrete setup, also due to the stochastic nature of the process). It therefore seems natural to anticipate that as time approaches infinity, the weighted revenue rates of both types of drivers converge to a common limit. Somewhat counter-intuitively, Case 2 in the explanation above implies that this conjecture is false in general.

To shed light on this result, imagine a hypothetical extreme situation where κ_L is negligibly small compared to κ_H. Suppose the participation intensities are fixed and consider one singled-out type-L driver in the discrete system with market size N. Provided this driver is available, passengers arrive within this driver's pickup region at rate $\lambda\delta$. For any fixed joining rate θ, every such passenger has a non-infinitesimal probability (invariant of κ_H) that the *only* available candidate is the tagged type-L driver, in which case this driver is selected for the matching. This means that type-L drivers' revenue is bounded from below by a non-zero value, irrespective of κ_L. Therefore, this driver obtains a non-negligible rate of revenue, and because κ_L is small, this driver's *normalized* revenue is arbitrarily large, in particular, larger than that of a type-H driver. This inequality can prevail even after factoring in equilibrium conditions: In equilibrium, when κ_L is sufficiently small, type-L drivers choose to participate. Because the normalized revenue rate of type-H drivers is bounded by $\overline{R}/\kappa_H$, when κ_H is fixed, one can set κ_L small enough such that $\hat{R}_L^*/\kappa_L > \overline{R}/\kappa_H \geq \hat{R}_H^*/\kappa_H$.

The above intuition can be mathematically verified, observing that when the ratio κ_H/κ_L is sufficiently large (that is, larger than a finite threshold), the third equation in (15) takes the form

$$1 - \hat{q}_H = \lambda m \cdot \left(1 - e^{-\hat{q}_H \theta_H \delta}\right),$$

and the solution to system (15) in that case does not depend on the opportunity costs. Such solution prescribes that in the stationary states, the platform should endow type-H drivers with full priority in the matching over type L.

We stress that the discrepancy in normalized revenues between drivers of different types does not require an extreme choice of parameters to occur. The characterization of the equilibrium participation intensities in Proposition 3 below points out that this phenomenon directly impacts the formation of equilibrium.

Equilibrium under MinWeightRev As a corollary of Lemma 5, we have that $(1-\hat{q}_H)/\kappa_H \leq (1-\hat{q}_L)/\kappa_L$, and therefore $\hat{R}^*_H/\kappa_H \leq \hat{R}^*_H/\kappa_L$, meaning type-$L$ drivers have stronger incentives to participate than type H. As in Section 5.1, we again exploit this structure by describing the participation intensity using a single value, $\theta \in (0, \Theta_L + \Theta_H]$, such that type-$H$ drivers participate iff $\theta > \Theta_L$. Using this characterization of the participation intensity we express the dependence of $\hat{q}_L, \hat{q}_H$ and $\hat{q}$ on θ through writing $\hat{q}_L(\theta), \hat{q}_H(\theta)$ and $\hat{q}(\theta)$. Then, we have the following lemma:

Lemma 6. *The steady-state, type-dependent proportions of available drivers, $\hat{q}_L$ and $\hat{q}_H$ are increasing functions of θ.*

The proof is in Appendix APX-G.

We further reintroduce the revenue rates as functions of θ, $\hat{R}^*_L(\theta)$ and $\hat{R}^*_H(\theta)$, which by Lemma 6 are decreasing in θ.

Proposition 3. *The equilibrium participation intensity under MinWeightRev, denoted by $\hat{\theta}^e$, is unique and satisfies exactly one of the following:*

1. $\hat{\theta}^e \leq \Theta_L$ *with* $\hat{R}^*_L(\hat{\theta}^e) \geq \kappa_L$ *and* $\hat{R}^*_H(\hat{\theta}^e) < \kappa_H$
2. $\hat{\theta}^e \in (\Theta_L, \Theta_L + \Theta_H]$ *with* $\hat{R}^*_L(\hat{\theta}^e) \geq \kappa_L$ *and* $\hat{R}^*_H(\hat{\theta}^e) \geq \kappa_H$

The proof is in Appendix APX-H.

6. Policy comparison

Equipped with Propositions 2 and 3 derived in Section 5 we can compare the equilibrium performance of the two policies. Recall the equilibrium participation intensities, $\theta^e, \hat{\theta}^e \in (0, \Theta_L + \Theta_H]$, corresponding with the two policies, MinRev and MinWeightRev respectively. For consistency, we denote

$$\theta^e_L = \min\{\theta^e, \Theta_L\} \quad \text{and} \quad \theta^e_H = \{\theta^e - \Theta_L\}^+,$$

and similarly,

$$\hat{\theta}^e_L = \min\{\hat{\theta}^e, \Theta_L\} \quad \text{and} \quad \hat{\theta}^e_H = \{\hat{\theta}^e - \Theta_L\}^+,$$

Lemma 7. *The equilibrium participation intensity for MinRev is never greater than that of MinWeightRev, namely,* $\theta^e \leq \hat{\theta}^e$*, with strict inequality iff* $\theta^e > \Theta_L$*.*

The proof is in Appendix APX-I

Lemma 7 suggests that by moving from MinRev policy to MinWeightRev, the platform increases the number of participating drivers. Following our characterization of the equilibrium participation intensities, it also follows that the participation intensity *per type* increases as well, $\theta_i^e \leq \hat{\theta}_i^e, i \in \{L, H\}$.

Recall that for a fixed participation intensity θ, the matching rate is given by $\lambda^*(\theta)$ as defined in (4). A corollary of Lemmas 2 and 5 is that for both policies in steady state, the latter can be expressed as

$$\lambda^*(\theta) = \lambda \cdot (1 - e^{q(\theta)\theta\delta})$$

where $q(\theta)$ is the (overall) availability proportion (replacing $q(\theta)$ by $\hat{q}(\theta)$ for MinWeightRev). It therefore follows from Lemma 3 that $\lambda^*(\theta)$ is increasing in θ. As a consequence, together with Lemma 7 above we have that $\lambda^*(\hat{\theta}^e) \geq \lambda^*(\theta^e)$, i.e., in equilibrium MinWeightRev yields higher matching rate than MinRev. We next prescribe bounds on this improvement, showing that both the participation intensity and the matching rate can potentially increase by a multiplicative factor of 2.

Recall from (10) the definitions of the improvement ratios, Φ and Ψ, for the participation intensity and the matching rate, respectively which can now be written as

$$\Phi = \frac{\hat{\theta}_L^e + \hat{\theta}_H^e}{\theta_L^e + \theta_H^e} = \frac{\hat{\theta}^e}{\theta^e} \quad \text{and} \quad \Psi = \frac{\lambda^*(\hat{\theta}^e)}{\lambda^*(\theta^e)} = \frac{1 - e^{\hat{q}(\hat{\theta}^e)\hat{\theta}^e\delta}}{1 - e^{q(\theta^e)\theta^e\delta}}$$

Clearly, by Lemma 7, both Φ and Ψ are greater or equal to one. With regard to the above ratios we can state the following result:

Proposition 4. *The participation intensity and the matching rate improvement ratios,* Φ *and* Ψ*, satisfy:*

1. $\Phi \in [1, 2]$*, with* $\Phi = 2$ *if and only if* $R_L^*(\theta^e) = \kappa_H$ *and* $\hat{R}_L^*(\hat{\theta}^e) \geq \kappa_L$*;*

2. $\Psi \in [1,2)$, *with* $\lim_{\Theta\to 0}\Psi = 2$ *if and only if* $\lim_{\Theta\to 0}\Phi = 2$.

The proof is in Appendix APX-J

We can in fact construct a mechanism that specifically chooses the system primitives so as to obtain the bound $\Phi = 2$. To this aim we first arbitrarily set values to λ, m, δ and Θ_L, and choose Θ_H such that $\Theta_H \geq \Theta_L$. To simplify the explanation assume that $\Theta_H = \Theta_L = \Theta/2$. Using Lemma 2 we compute $R^*(\Theta/2)$, i.e. the revenue rate assuming that all type-L drivers (and only them) participate, and set $\kappa_H = R^*(\Theta/2)$. Assuming $\kappa_L < \kappa_H$, this ensures that in equilibrium under MinRev, all type-L drivers participate. Type-H are "on the verge" of participating, however, because $R^*(\theta)$ is decreasing in θ (by Lemma 3), they do not do so under MinRev. We then choose an arbitrarily small value for κ_L and solve for the equilibrium revenue rates using Equation (15), assuming that all drivers join. Specifically, we choose κ_L such that the solution to (15) does not depend on the opportunity costs. Under this assumption, type-H drivers obtain matching priority in steady state under MinWeightRev, and therefore are indifferent to the existence of type-L in the system. The revenue rate that type-H drivers obtain when they all participate is exactly κ_H, namely $\hat{R}^*_H(\Theta) = R^*(\Theta/2) = \kappa_H$, making $\Theta/2$ the equilibrium participation intensity for type-H under MinWeightRev. Assuming that κ_L is chosen such that the solution to (15) is independent of κ_H/κ_L, we ensure that $\hat{R}^*_L(\Theta) \geq \kappa_L$, facilitating the participation of type-L drivers in this equilibrium. Thus, the equilibrium participation intensity grows from $\Theta/2$ under MinRev to Θ under MinWeightRev.

Proposition 4 also asserts that using this mechanism, if λ is initially set small enough, the doubled participation intensity will result in nearly 100% improvement in terms of matching rate. This is due to the fact that around a point $x = 0$, the value of $1 - e^{-x}$ can be approximated by x itself, thus $(1 - e^{-2x})/(1 - e^{-x}) \approx 2$. Letting Θ diminish and keeping $\Phi = 2$ using the mechanism above will result in Ψ approaching the value 2. Similar limiting behavior is exhibited by taking $\delta \to 0$.

Lemmas 1–7 together with Propositions 2–4 constitute a full proof of our main result Proposition 1. In the subsequent section we support these analytic findings numerically and compare them by simulation to the performance of the discrete model described in Section 2.

7. Numerical results and simulation

In this section we conduct an extensive numerical study to support the analytic results of Section 5. Unless explicitly said differently, we will assume throughout this section that the following model parameters satisfy $\lambda = m = \delta = r = 1$. We will assume, as in Section 5.2, that given the total participation intensity $\theta \in (0, \Theta_L + \Theta_H]$, the type-dependent participation intensities satisfy $\theta_L = \max\{\theta, \Theta_L\}$ and $\theta_H = [\theta - \Theta_L]^+$.

It is worth noticing that for this specific parameters set of choice, the revenue rate for a driver (under any policy and for any participation profile) is bounded by $\overline{R} = r \cdot m/((\lambda\delta)^{-1} + m) = .5$. Furthermore, because $r = m = 1$, the matching rate equals the total revenue rate in the system. That is, for every θ, $\lambda^*(\theta) = \theta \cdot R^*(\theta)$ under MinRev, and under MinWeightRev, $\lambda^*(\theta) = \theta_L \cdot \hat{R}_L^*(\theta) + \theta_H \cdot \hat{R}_H^*(\theta)$.

7.1. Validation and justification of theoretic results

The purpose of the first example is to visually illustrate the theoretic results of Propositions 2 and 3, namely the monotonicity of the revenue rates and uniqueness of equilibrium. We consider the pair of opportunity costs $\kappa_L = .35$ and $\kappa_H = .45$, with potential intensities $\Theta_L = \Theta_H = 1$. Figure 3 depicts the change in revenue rates R^* (black), $\hat{R}_L^*$ (green) and $\hat{R}_H^*$ (purple) as functions of θ. The dashed horizontal lines correspond with the values of κ_L (green) and κ_H (purple), and the dashed vertical line correspond with the values of Θ_L (green) and Θ_H (purple). Noticeably, when $\theta \le \Theta_L$, only type-L drivers participate and the two policies are equivalent, thus we will be mainly focusing on $\theta > \Theta_L$. As implied by Lemmas 3 and 6, all three functions, R^*, $\hat{R}_L^*$ and $\hat{R}_H^*$, are decreasing with θ.

At $\theta = \Theta_L = 1$, the revenue for L-drivers (under both policies) takes the value $R^*(\Theta_L) = .433$, thus satisfying $\kappa_L < R^*(\Theta_L) < \kappa_H$, meaning that $\theta^e = \Theta_L = 1$ is the equilibrium participation intensity under MinRev. Interestingly, for MinWeightRev, it can be seen that

$$\lim_{\theta \to \Theta_L^-} \hat{R}_H^*(\theta) = 5 = \overline{R}.$$

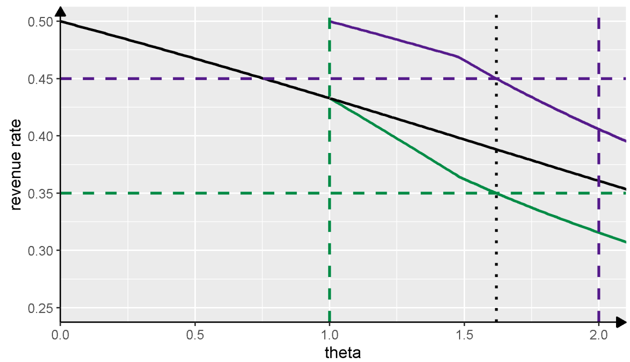

Figure 3 Revenue rates as a functions of participation intensity, θ. All three functions, $R^*(\theta)$ (black), $\hat{R}^*_L(\theta)$ (green) and $\hat{R}^*_H(\theta)$ (purple) are decreasing. Equilibrium under MinRev is induced by $\theta^e = \Theta_L = 1$, and under MinWeightRev (vertical dotted black line) by $\hat{\theta}^e = 1.619$

This is because for the specific choice of parameters here, when $\theta_H (= [\theta - \Theta_L]^+)$ is small, H-drivers obtain strict priority in the matching, thus when their participation intensity is arbitrarily small, they approach the maximal revenue rate possible. The non-smoothness of $\hat{R}^*_H(\theta)$ and $\hat{R}^*_L(\theta)$ around the point $\theta = 1.479$ is the point where MinWeightRev transitions from granting H-drivers with strict priority ($\hat{s}_H = 1$) to partial priority ($\hat{s}_H < 1$), in which case we have that the normalized revenue rate is equal for both types. Marked by a black dotted vertical line, the point $\theta = 1.619$ is where both types break even, namely $\hat{R}^*_L(\theta)/\kappa_L = \hat{R}^*_H(\theta)/\kappa_H = 1$, indicating that this is equilibrium: $\hat{\theta}^e = 1.619$. In this example we conclude that $\Phi = 1.619/1 = 1.619$, and further elementary calculations yield $\Psi = .629/.433 = 1.453$.

For further validation of our numeric results, and for the purpose of justifying our mean field formulation, we conduct a simulation study concerning the above setup. Elaborate explanation of the simulation process is given in Section 2. We plot empirical observations of the system state with its mean field counterparts at different points in time. As an input to our simulation process we need to set the market size, N, and we experiment with different values of N.

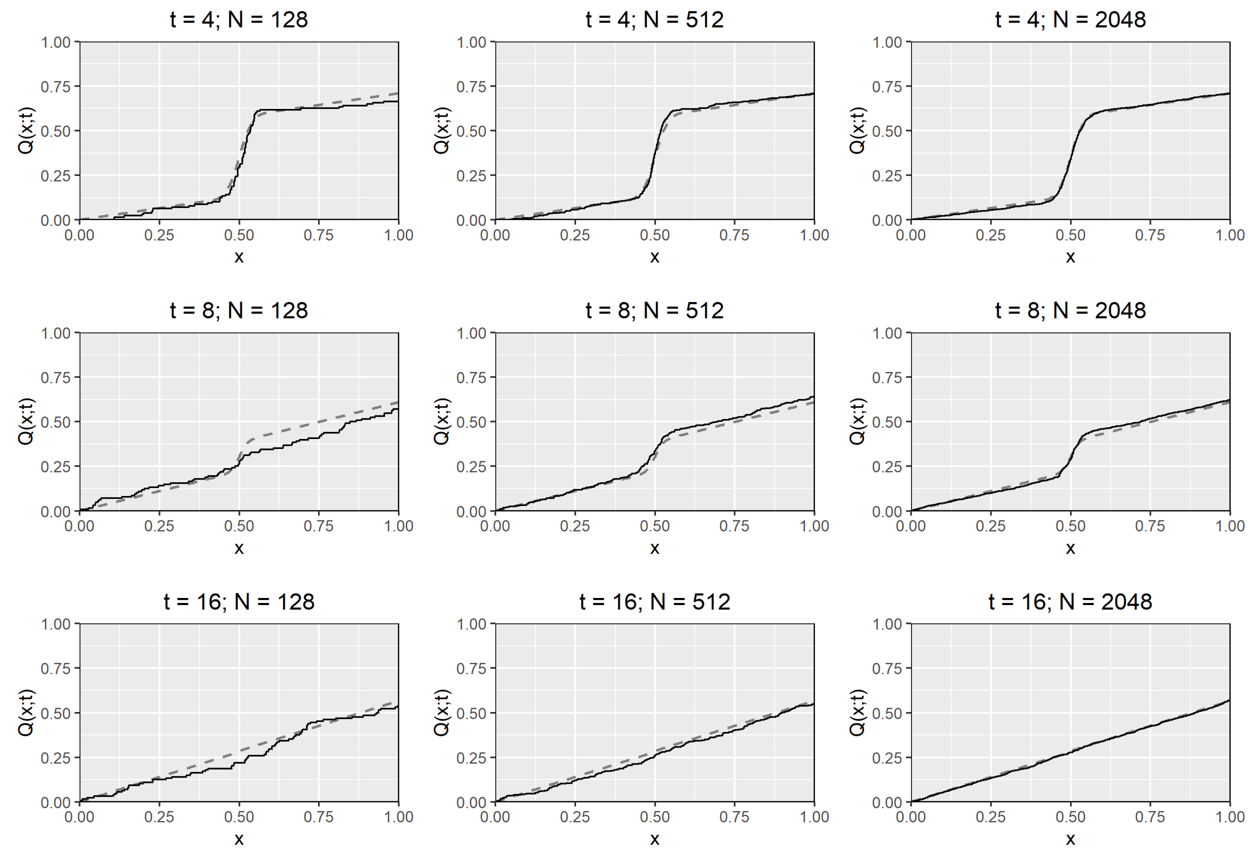

Figure 4 Simulated state (solid black) and mean field state (dashed gray) for various values of N and t. The simulated state approaches the mean field state for every t as N increases, and the latter approaches its stationary state as t increases.

In the simulation, we assume that all drivers are initially available with 0 revenue and their locations are drawn independently from a distribution whose density, f, for $x \in [.4, .6]$ is given by:

$$f(x) = \begin{cases} 100 \cdot (x - .4) & \text{if } x \in [0.4, .5), \\ 10 - 100 \cdot (x - .5), & \text{if } x \in [.5, .6], \end{cases}$$

and is zero for any $x \notin [.4, .6]$. Thus f forms an isosceles triangle with height 10 (and base length .2) centered at $x = .5$. For both types, the initial state for the mean field system is the cdf of that distribution.

First we consider the MinRev policy at its equilibrium point, $\theta^e = 1$. Note that in this case only L-drivers participate, thus $Q_L(x;t)$ and $Q(x;t)$ coincide (for all t). Figure 4 depicts the simulated

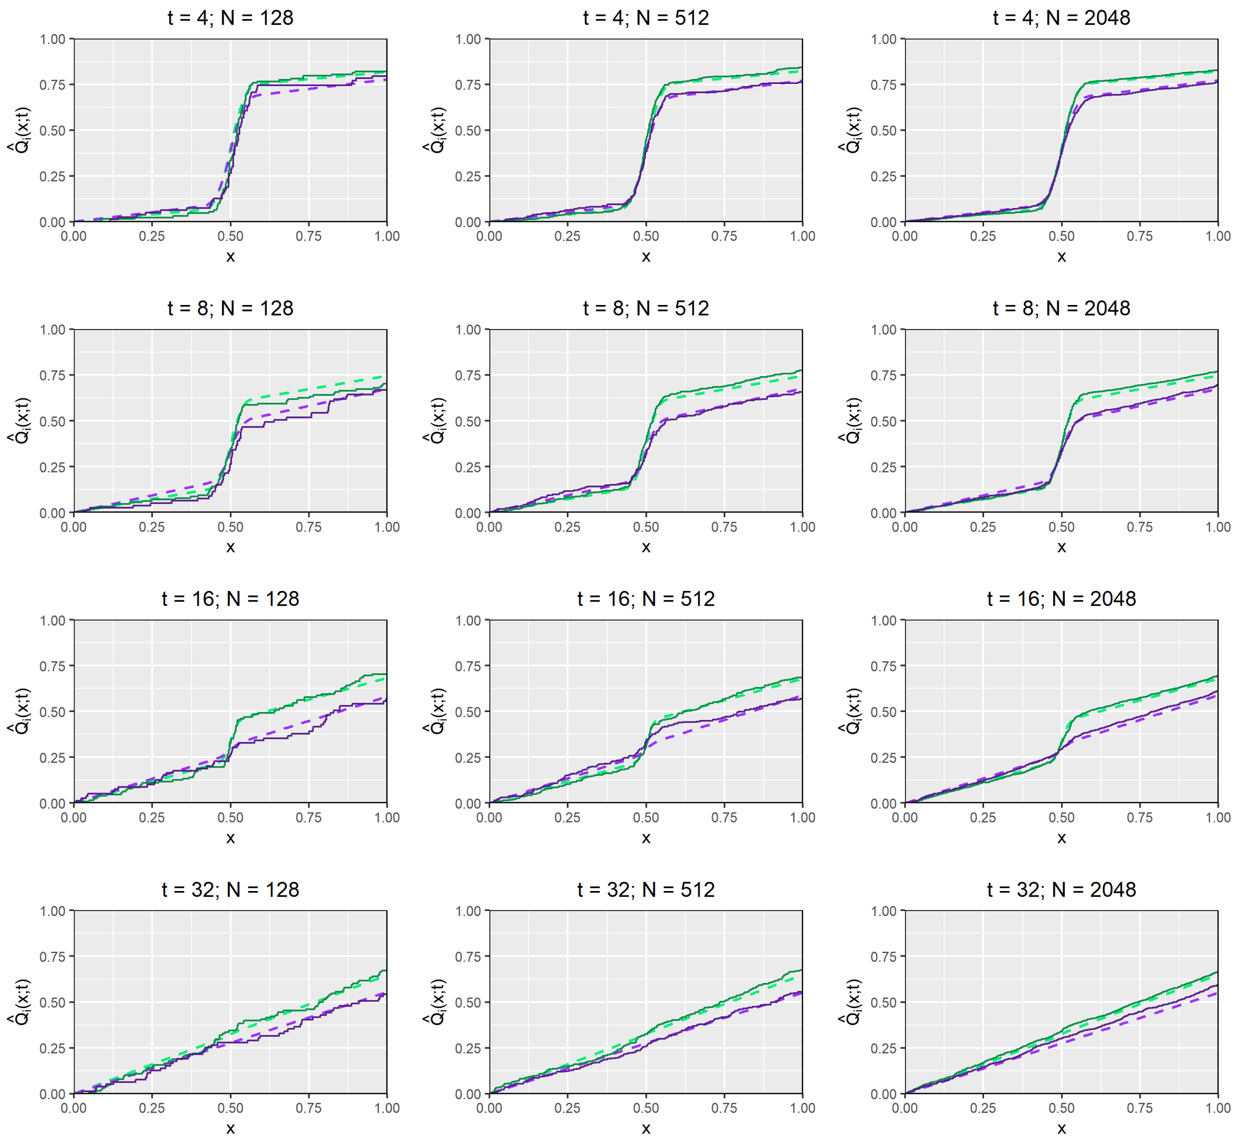

Figure 5 Simulated state (solid black) and mean field state (dashed gray) for various values of N and t. The simulated state approaches the mean field state for every t as N increases, and the latter approaches its stationary state as t increases.

state (solid black line) and its mean field counterpart $Q(x;t)$ (dashed gray line), across all $x \in [0,1)$, for all combinations of $N \in \{2^7, 2^9, 2^{11}\}$ and $t \in \{2^2, 2^3, 2^4\}$. It can be seen that as the market size N scales up, the two curves become arbitrarily close to each other, for all t. Moreover, for large t, the mean field state approaches a linear function with slope .567, which agrees with the unique steady state characterized by Lemma 2.

Next we apply the same methodology to study the system under MinWeightRev, with participation intensity equal to its corresponding equilibrium $\hat{\theta}^e = 1.619$ (implying $\hat{\theta}^e_L = 1$ and $\hat{\theta}^e_H = .619$). We simulate the system for all combinations of $N \in \{2^7, 2^9, 2^{11}\}$ and $t \in \{2^2, 2^3, 2^4, 2^5\}$ (with the implementation of collective revenue accumulation as discussed in Section 3). The solid curves in Figure 5 depict the simulated state for each type, green for L-drivers and purple for H-drivers. The dashed curves correspond to the parallel state of the mean field system, $\hat{Q}_L(x;t)$ (green) and $\hat{Q}_H(x;t)$ (purple). We conclude that the qualitative results regarding the accuracy of the mean field approximation and the convergence to the stationary state apply to MinWeightRev as well.

7.2. Extreme improvement example

This example is tailored to demonstrate the extreme factor-2 improvement in participation that can be obtained when moving from MinRev to MinWeightRev. We consider the pair of opportunity costs $\kappa_L = .27$ and $\kappa_H = .433$, with potential intensities $\Theta_L = \Theta_H = 1$. Intentionally, we chose κ_H such that it equals the revenue rate when only L-drivers participate (see Example 7.1), that is, $\kappa_H = R^*(\Theta_L)$. This is in correspondence with the conditions of the first item of Proposition 4. Similar to the description of Figure 3, the left panel of Figure 6 shows the revenue rates as functions of the participation intensity. It can be seen that under MinRev (black curve) at $\theta = \Theta_L = 1$, the revenue is above κ_L, yet H-drivers break even, i.e., $R^*(\theta) = \kappa_H > \kappa_L$. This means that if a non-zero intensity of H-drivers participate, their revenue will be strictly less than their opportunity cost, hence this is equilibrium: $\theta^e = \Theta_L = 1$. Under MinWeightRev, H-drivers break even at $\theta = \Theta_L + \Theta_H = 2$ (while L-drivers revenue is still above κ_L), meaning that $\hat{\theta}^e = 2$. It therefore follow that $\Phi = 2/1 = 2$, and by calculating the equilibrium matching rates we get $\Psi = .703/.433 = 1.624$.

By considering smaller potential intensities we can construct a similar example so as to make Ψ arbitrarily close to 2. For example, by changing the potential intensities to $\Theta_L = \Theta_H = .1$ and the opportunity costs to $\kappa_L = .48$ and $\kappa_H = .494$, we obtain $\theta^e = \Theta_L = .1$ and $\hat{\theta}^e = \Theta_L + \Theta_H = .2$, thus $\Phi = .2/.1 = 2$, and it can be shown that $\Psi = .975/.494 = 1.974$. This is illustrated in the right panel of Figure 6.

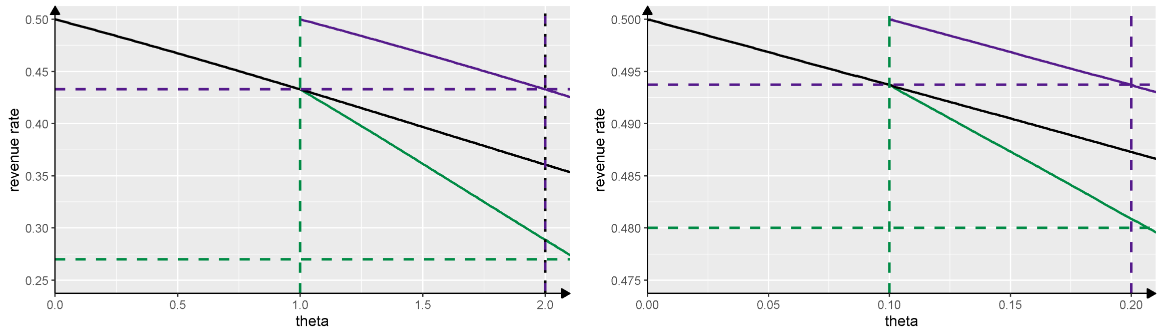

Figure 6 Revenue rates as a functions of participation intensity, θ. Left panel considers $\Theta_L = \Theta_H = 1$ and $\kappa_L = .27 < \kappa_H = .433$, and right panel considers $\Theta_L = \Theta_H = .1$ and $\kappa_L = .48 < \kappa_H = .494$. In both panels, equilibrium under MinRev is induced by $\theta^e = \Theta_L$, and under MinWeightRev by $\hat{\theta}^e = \Theta_L + \Theta_H$

7.3. Pickup range improvement

Propostion 4 as illustrated by the previous example suggests that the most significant improvement in terms of the matching rate ratio Ψ is obtained when the potential intensity Θ is small, hence the absolute values of the equilibrium matching rates are small too. When equilibrium under MinRev is such that the matching rate is already high, switching to MinWeightRev will not make significant impact on utilization. However, the increased participation intensity obtained by MinWeightRev allows the platform to improve its pickup range without forgoing more potential demand. We demonstrate this idea in the following example.

Suppose the potential market for drivers is characterized by $\Theta_L = 2$ and $\Theta_H = 1$, with opportunity costs $\kappa_L = .37$ and $\kappa_H = .5$. Given the matching policy, the platform has to set the minimal δ so as to meet a matching rate goal of $.99 \times \lambda$, meaning 99% of potential passengers are matched in equilibrium under the predetermined policy. Figure 7 depicts various values θ and their corresponding δ such that the resulting matching rate is .99.

When the policy is MinRev, this goal is obtained by setting $\delta = 4.559$, and the corresponding equilibrium is given by $\theta^e = \Theta_L = 2$ for which $R^*(\Theta_L) = .495 \in (\kappa_L, \kappa_H)$. Left panel of Figure 8 depicts the revenue rate as a function of θ. Under MinWeightRev with $\delta = 3.05$, the resulting equilibrium $\hat{\theta}^e = 2.5$ (see Figure 8 right panel, dotted black line), which also obtains a matching

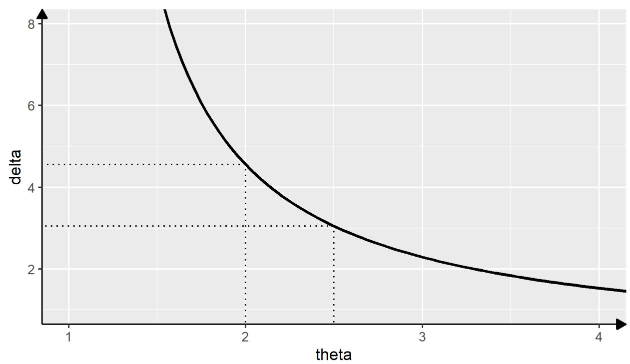

Figure 7 Pickup range required for .99 matching rate, as a function of participation intensities. To guarantee a matching rate $\lambda^* = .99 \times \lambda$, when $\theta = 2$, the required pickup range is $\delta = 4.559$, whereas for $\theta = 2.5$, the required pickup range is $\delta = 3.05$ which is roughly 33% smaller.

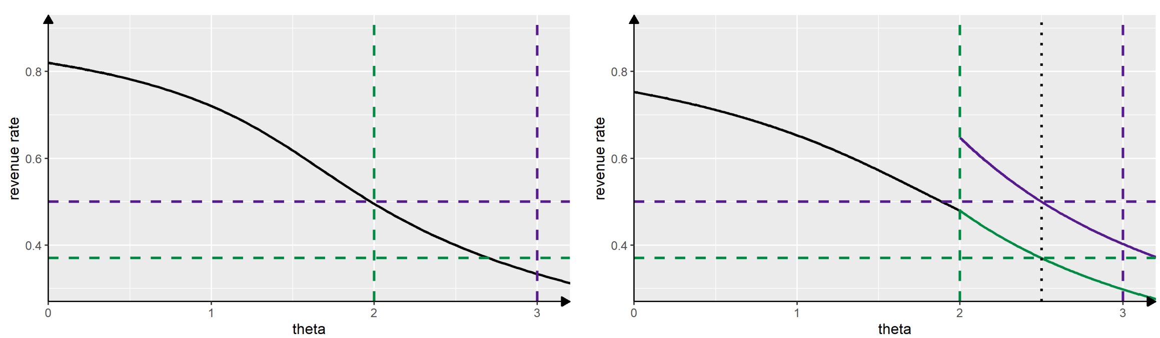

Figure 8 Revenue rates as a functions of participation intensity, θ. In the left panel, the policy is MinRev, and $\delta = 4.559$. The corresponding equilibrium is $\theta^e = \Theta_L = 2$. In the right panel, the policy is MinWeightRev, $\delta = 3.05$ and the corresponding equilibrium is $\hat{\theta}^e = 2.5$ (black dotted line).

rate of .99, as shown in Figure 7. By switching to MinWeightRev the platform can reduce its pickup range by 33%, from 4.559 to 3.05.

8. Concluding remarks

In this paper we suggest a novel modeling approach to study the behavior of nomadic agents that are constantly moving across locations in a network and interact with each other through matching mechanism controlled by a centralizing platform. Our model accounts for the possibility of rides being lost due to temporal and spatial mismatches between drivers' supply and passenger's demand at specific locations of the network. Yet it is simple enough to allow stationary-state and equilibrium analysis to be carried out, for two different (yet closely related) policies. We derive analytic results concerning the uniqueness of equilibrium and bounds on the improvement in equilibrium performance obtained by changing the policy from the symmetric MinRev policy to the slightly more sophisticated MinWeightRev. In particular, we show that the MinWeightRev policy, in which drivers with different opportunity costs are treated differently by the policy, yields to more efficient allocation of revenue, which in turn increases the equilibrium participation of drivers, thus, generating more matches and increasing system efficacy.

Not surprisingly, it is shown that in terms of revenue-to-goal ratio, drivers with lower opportunity cost are in a (weakly) favorable position compared to those with higher opportunity cost, under both policies discussed. What is perhaps less intuitive is that even though MinWeightRev is designed to eliminate this discrepancy, under certain market conditions it may not be possible to do so, assuming that the platform do not reject passengers when it can potentially match them to available drivers.

We justify the formulation of the mean field model by comparing it with simulation results of the motivating discrete model. In fact, the motivating discrete model we consider is not a single, but rather a sequence of systems indexed and scaled by N, the market size. The scaling procedure involves increasing the number of agents, as well as the arrival rate, linearly with N, yet keeping constant the average number of available drivers in a passengers pickup region. Of much interest in the study of fluid approximations for queueing models is proving an interchange-of-limits argument. By that we mean proving three properties: (a) that for any fixed time t, observing the state of the N-th system at time t, as $N \to \infty$, this sequence of states converges to the state of the mean field system at time t (given an appropriate sequence of corresponding initial states); (b) that for each N, the N-th system converges to a steady-state at $t \to \infty$; and (c) that when both $N \to \infty$ and $t \to \infty$, one obtains convergence to the steady state of the fluid system, irrespective of the order according to which the limits are taken. In general, proving the interchange of limits is often involved, and supposedly very much so in the setup described here.

Because the purpose of this paper is to demonstrate the advantage of income-goal-aware matching in ride sharing platforms, we leave the task of proving the interchange of limits for future research. However we justify our mean field formulation based on an extensive simulation study described in Sections 2 and 7. These simulations show that the mean field model provides a practical and useful approximation for the discrete model, even when the market size N is of moderate size. The simulations also give a strong empirical evidence that the mean field system is indeed the fluid limit of the discussed sequence of systems, and that it converges to a stationary state as time tends to infinity.

To keep insights sharp we focus on a stylized model, imposing simplifying assumptions about the dynamics of the system, the geometric structure, drivers characteristics, etc. For this setup we prove our main managerial insight: that a matching policy which is aware of drivers' opportunity cost in equilibrium performs generally better than a policy that ignores it, because taking into account drivers' opportunity cost in equilibrium promotes higher willingness to participate among drivers with high opportunity cost. The maximal improvement ratio is proved to be 2, namely twice as much drivers in the former equilibrium as in the latter, and this bound is tight. The simple explanation behind this result implies that it is quite robust. In particular, we believe that assuming a more elaborate opportunity cost structure (e.g., more than two, as well as a continuum of types), or modifying the dynamics, will not yield a greater improvement ratio. Nevertheless, generalizing the model to cope with more than two types, non-uniform, possibly dependent pickup/drop-off locations, and more complex geometries, are important extensions that may appeal for practitioners.

Lastly, an important issue which we completely disregard in this paper is the question of how the platform can obtain knowledge about the different drivers' opportunity costs. In the introduction, we briefly mention that incentives and bonus programs are widely used in practice, from which useful knowledge about drivers income goals can be extracted. The question of devising efficient bonuses for drivers is interesting in its own right, and studying it from a mechanism design perspective may reveal new interesting research directions. Quantifying the value of drivers' commitment for the platform, which is the main purpose of this paper, is an essential prerequisite for an efficient design of any incentive program, and we hope the framework presented here can lay the foundation for such research.

APPENDIX

APX-A. Proof of Lemma 1

We show first uniqueness of a solution for $Q(x;t)$. Observe that under the assumption $Q(0;t)=0$ for all $t\in[0,\infty)$, together with the definition of $Q'(x;t)$, we have that

$$Q(x;t)=\int_{s=0}^{x}Q'(s;t)ds, \qquad \forall x\in[0,1),$$

therefore showing there exists a unique solution for $Q'(x;t)$ will imply uniqueness of the solution for $Q(x;t)$. By taking the derivative with respect to x of both sides of (8) we get, using Leibniz' rule of integration, for all $x\in[0,1)$

$$\frac{\partial Q'(x;t)}{\partial t}=\left(1-\int_{s=0}^{1}Q'(s;t)ds\right)m-\frac{\lambda}{\theta}\left(1-e^{-Q'(x;t)\theta\delta}\right) \tag{17}$$

Let $\mathcal{G}$ be the space of real, positive and absolutely integrable functions over $[0,1)$. For any $f\in\mathcal{G}$, define:

$$\varphi f(x)=\frac{1}{m}\left(1-\int_{s=0}^{1}f(s)ds\right)-\frac{\lambda}{\theta}\left(1-e^{-\theta\delta f(x)}\right).$$

with initial value $f_{t_0}(x)=f_0(x)$ We continue the proof by showing first that φ is Lipschitz continuous (with respect to the uniform norm), and then showing convergence of the sequence of Picard iterations to a fixed point.

To show that φ is Lipschitz continuous, note firstly that for any two functions $f,g\in\mathcal{G}$,

$$\left|\int_{s=0}^{1}f(s)ds-\int_{s=0}^{1}g(s)ds\right|\le\int_{s=0}^{1}\left|f(s)-g(s)\right|ds\le\int_{s=0}^{1}\|f-g\|_{\infty}ds$$
$$=\|f-g\|_{\infty},$$

and in addition, assuming that f is a positive function, for any $\alpha>0$

$$\left\|e^{-\alpha f}-e^{-\alpha g}\right\|_{\infty}=\left\|e^{-\alpha f}\left(1-e^{\alpha f-\alpha g}\right)\right\|_{\infty}\le\left\|e^{-\alpha f}\cdot(1-(1+\alpha f-\alpha g))\right\|_{\infty}$$
$$=\left\|e^{-\alpha f}\cdot\alpha(g-f)\right\|_{\infty}\le\alpha\|f-g\|_{\infty},$$

where the first inequality follows from the fact that for any $z \in \mathbb{R}$, $e^z \geq 1+z$. It follows that φ is Lipschitz continuous as a sum of such functions, and we denote its Lipschitz constant by L.

Fix an initial time t_0 and value $f_0 \in \mathcal{G}$. Define $\mathcal{H}$ as the set of all continuous functions of the form $F(x;t):[0,1)\times\mathbb{R}\to\mathbb{R}$ such that for every fixed t, $F(\cdot;t)\in\mathcal{G}$, and equip it with the uniform norm $\|F\|_\infty = \sup_{x,\tau}|F(x,\tau)|$. Define the operator Γ by

$$(\Gamma F)(x;t) = f_0 + \int_{\tau=t_0}^{t} \varphi F(x;\tau)d\tau,$$

Note that for every regular point t, $(\Gamma F)(\cdot;t)$ defines a real function over $x\in[0,1)$. Moreover, since φ is L-Lipschitz continuous, $f_0 - Lt \leq \Gamma F(\cdot;t) \leq f_0 + Lt$, meaning $\Gamma F(\cdot;t)$ is absolutely integrable and hence resides in $\mathcal{G}$, thus Γ maps the space $\mathcal{H}$ into itself. We next show that if t is chosen in a $\frac{1}{L}$-neighborhood of t_0 then Γ is a contraction mapping.

Fix t and let $F,G\in\mathcal{H}$. Let $t^*\in[t_0,t]$ be the point satisfying

$$\|\Gamma F - \Gamma G\|_\infty = \|\Gamma F(\cdot;t^*) - \Gamma G(\cdot;t^*)\|_\infty = \sup_{x\in[0,1)} |\Gamma F(x;t^*) - \Gamma G(x;t^*)|.$$

Note that in the first term the maximum is taken jointly on both arguments while in the second only on the first (x) argument. Then,

$$\begin{aligned}\|\Gamma F - \Gamma G\|_\infty &= \sup_{x\in[0,1)} \left| \int_{\tau=t_0}^{t^*} \varphi F(x;\tau) - \varphi G(x;\tau) d\tau \right| \leq \int_{\tau=t_0}^{t^*} \|\varphi F(x;\tau) - \varphi G(x;\tau)\|_\infty d\tau \\ &\leq L \int_{\tau=t_0}^{t^*} \|F(x;\tau) - G(x;\tau)\|_\infty d\tau \leq L(t^*-t_0)\|F-G\|_\infty \leq L(t-t_0)\|F-G\|_\infty,\end{aligned}$$

and choosing t such that $t - t_0 < \frac{1}{L}$ implies that Γ is a contraction. It follows from Banach fixed point theorem that the operator has a unique fixed point. In particular, there is a unique function F satisfying $\Gamma F = F$, and so, we established the uniqueness of a local solution to the initial value problem.

In order to show global uniqueness we utilize Gronwall's inequality: Let $F,G\in\mathcal{H}$ be two different solutions and denote for all t, $F(x;t) = f_t(x)$ and $G(x;t) = g_t(x)$ with identical initial values $f_0 = g_0$. Consider the two following functions:

$$z(t) = \int_{x=0}^{1} (F(x;t) - G(x;t))^2 dx = \int_{x=0}^{1} (f_t(x) - g_t(x))^2 dx,$$

and

$$y(t) = \left(\int_{x=0}^{1} (F(x;t) - G(x;t))dx \right)^2 = \left(\int_{x=0}^{1} (f_t(x) - g_t(x))dx \right)^2.$$

Note by the assumption $f_0 = g_0$ we have that at t_0,

$$z(t_0) = y(t_0) = 0.$$

We shall show that $z \equiv 0$, by relying on that $y \equiv 0$, which we show first.

For every t, due to φ being L-Lipschitz we have

$$\begin{aligned}
\frac{d}{dt} y(t) &= 2\left(\int_{x=0}^{1} (f_t(x) - g_t(x))dx \right) \cdot \frac{\partial}{\partial t} \left(\int_{x=0}^{1} (F(x;t) - G(x;t))dx \right) \\
&= 2\left(\int_{x=0}^{1} (f_t(x) - g_t(x))dx \right) \cdot \left(\int_{x=0}^{1} (\varphi f_t(x) - \varphi g_t(x))dx \right) \\
&\leq 2\left(\int_{x=0}^{1} (f_t(x) - g_t(x))dx \right) \cdot \left(\int_{x=0}^{1} \|\varphi f_t - \varphi g_t\|_\infty dx \right) \\
&\leq 2\left(\int_{x=0}^{1} (f_t(x) - g_t(x))dx \right) \cdot L\|f_t - g_t\|_\infty = y(t)\beta(t)
\end{aligned}$$

where

$$\beta(t) = \frac{2L\|f_t - g_t\|_\infty}{\int_{x=0}^{1} (f_t(x) - g_t(x))dx}.$$

By Gronwall's inequality,

$$y(t) \leq y(t_0) e^{\int_{t_0}^{t} \beta(u)du} = 0,$$

and therefore $y \equiv 0$, implying that for all t,

$$\int_{x=0}^{1} f_t(x)dx = \int_{x=0}^{1} g_t(x)dx.$$

We move on to showing $z \equiv 0$. Using the above relation, together with the definition of φ, we get, for every t and $x \in [0,1)$,

$$\begin{aligned}
|\varphi f_t(x) - \varphi g_t(x)| &= \frac{\lambda}{\theta} \left| e^{-\theta\delta f_t(x)} - e^{-\theta\delta g_t(x)} \right| = \frac{\lambda}{\theta} \left| e^{-\theta\delta f_t(x)} (1 - e^{\theta\delta(f_t(x) - g_t(x))}) \right| \\
&\leq \frac{\lambda}{\theta} e^{-\theta\delta f_t(x)} \left| 1 - (1 + \theta\delta(f_t(x) - g_t(x))) \right| \leq \lambda\delta \left| f_t(x) - g_t(x) \right|,
\end{aligned}$$

and therefore,

$$\begin{aligned}
\frac{d}{dt}z(t) &= \int_{x=0}^{1} \frac{\partial}{\partial t}(F(x;t)-G(x;t))^2 dx \\
&= \int_{x=0}^{1} 2(F(x;t)-G(x;t)) \left(\frac{\partial F(x;t)}{\partial t} - \frac{\partial G(x;t)}{\partial t} \right) dx \\
&\leq \int_{x=0}^{1} 2|f_t(x)-g_t(x)| \cdot |\varphi f_t(x) - \varphi g_t(x)|\, dx \\
&\leq \int_{x=0}^{1} 2\lambda\delta (f_t(x)-g_t(x))^2 dx = 2\lambda\delta \cdot z(t)
\end{aligned}$$

Utilizing Gronwall's inequality once again we get

$$z(t) \leq z(t_0)e^{2\lambda\delta t} = 0,$$

and we conclude that $(F-G)^2 \equiv 0$, hence, proving global uniqueness.

Next we shall show that $Q_i'(x;t)$ exists uniquely.

Proof outline:

Set $i \in \{L, H\}$. We treat based on the above explanations, given the initial states, $\theta Q'(x;t)$ exists uniquely and therefore we can treat it as an exogenous input. Define

$$\begin{aligned}
w(t) := \frac{t\theta_i}{m\lambda r} \cdot (R_j(t) - R_i(t)) &= \frac{\theta_i}{m\lambda} \cdot \int_{u=0}^{t} (Q_i(1;u) - Q_j(1;u)) du \\
&= \frac{\theta_i}{m\lambda} \cdot \int_{u=0}^{t} \int_{s=0}^{1} (Q_i'(s;u) - Q_j'(s;u)) ds du
\end{aligned}$$

from which it is clear that $R_i(t) > R_j(t) \Leftrightarrow w(t) < 0$ and $R_i(t) = R_j(t) \Leftrightarrow w(t) = 0$. Note that $Q_j'(s;u)$ can be written as

$$Q_j'(s;u) = \frac{\theta Q'(s;u) - \theta_i Q_i'(s;u)}{\theta_j}.$$

Define functions $\phi : \mathcal{G} \to \mathbb{R}$, $\psi_i, \psi_j : \mathcal{G} \to \mathcal{G}$ and $\eta : \mathcal{G} \times \mathbb{R} \to \mathcal{G}$, such that given $y \in \mathcal{G}$,

$$\phi(y) = \frac{\theta_i}{m\lambda} \left(1 - \int_{s=0}^{1} y(s) ds \right) \tag{18}$$

$$\psi_i(y) = 1 - e^{-\delta\theta_i y} \text{ and } \psi_j(y) = 1 - e^{-\delta\theta_j y} \tag{19}$$

$$\eta(y,t) = \frac{\theta Q'(\cdot;t) - \theta_i y}{\theta_j} \tag{20}$$

and note that

$$w(t) = \int_{u=0}^{t} \left(\phi(\eta(Q_i'(\cdot;u),u)) - \phi(Q_i'(\cdot;u))\right) du$$

and that $\psi_i(0) = \psi_j(0) \equiv 0$. Define the set-valued positivity indicator function $\mathbf{I} : \mathbb{R} \to 2^{\mathbb{R}}$:

$$\mathbf{I}(s) := \operatorname{conv}\big(\mathbf{1}(s > 0)\big) = \begin{cases} \{0\} & \text{if } s < 0 \\ [0,1] & \text{if } s = 0 \\ \{1\} & \text{if } s > 0. \end{cases}$$

Note that $-\mathbf{I}(s)$ is OSL, i.e., for every $s, \tilde{s} \in \mathbb{R}$ and $z \in -\mathbf{I}(s)$, $\tilde{z} \in -\mathbf{I}(\tilde{s})$,

$$(s - \tilde{s})(z - \tilde{z}) \le 0 \le (s - \tilde{s})^2.$$

Let $\mathcal{D} := \mathcal{Q} \times \mathbb{R}^2$. We redefine our DE as the following differential inclusion problem, with the (set-valued) drift function φ:

$$\frac{\partial}{\partial t} \begin{pmatrix} y \\ \omega \\ \tau \end{pmatrix} \in \varphi(y, \omega, \tau) \times \{\phi(\eta(y,\tau)) - \phi(y)\} \times \{1\} \quad \text{a.e.,} \ (y, \omega, \tau) \in \mathcal{D} \tag{21}$$

where

$$\varphi(y, v, \tau) = \frac{\lambda}{\theta_i} \cdot \left\{ \phi(y) - \Big(1 - \psi_j\big(\eta(y,\tau)\big)\Big)\psi_i(y) - z \cdot \psi_j\big(\eta(y,\tau)\big)\psi_i(y) \ : \ z \in \mathbf{I}(\omega) \right\}. \tag{22}$$

For $\mathbf{y} = (y, \omega, \tau) \in \mathcal{D}$, we define the total (set-valued) drift term $\mathbf{F} : \mathcal{D} \to 2^{\mathcal{D}}$:

$$\mathbf{F}(\mathbf{y}) = \mathbf{F}(y, \omega, \tau) = \varphi(y, \omega, \tau) \times \{\phi(\eta(y,\tau)) - \phi(y)\} \times \{1\}$$

All values of $\mathbf{F}$ are convex and compact and $\mathbf{F}$ has a closed graph and therefore it is USC (upper semi-continuous). We endow the space $\mathcal{G}$ with the ℓ^2 norm and its corresponding inner product, i.e., for $y, \tilde{y} \in \mathcal{G}$

$$\langle y, \tilde{y} \rangle = \int_{s=0}^{1} y(s)\tilde{y}(s) ds, \ \text{ and } \ \|y\| = \left(\int_{s=0}^{1} (y(s))^2 ds \right)^{1/2}.$$

We naturally extend this to $\mathcal{D}$ by defining for all $\mathbf{y},\tilde{\mathbf{y}}\in\mathcal{D}$ with $\mathbf{y}=(y,\omega,\tau)$ and $\tilde{\mathbf{y}}=(\tilde{y},\tilde{\omega},\tilde{\tau})$:

$$\langle\mathbf{y},\tilde{\mathbf{y}}\rangle=\langle y,\tilde{y}\rangle+\omega\cdot\tilde{\omega}+\tau\cdot\tilde{\tau},\ \text{ and }\ \|\mathbf{y}\|=\left(\|y\|^2+\omega^2+\tau^2\right)^{1/2}$$

thus, defining and inner product and its induced norm over $\mathcal{D}$.

Our goal is to show that $\mathbf{F}$ satisfies the OSL (one-sided Lipschitz) condition, namely, for all $\mathbf{y},\tilde{\mathbf{y}}\in\mathcal{D}$, with $\mathbf{z}\in\mathbf{F}(\mathbf{y})$ and $\tilde{\mathbf{z}}\in\mathbf{F}(\tilde{\mathbf{y}})$,

$$\langle\mathbf{y}-\tilde{\mathbf{y}},\mathbf{z}-\tilde{\mathbf{z}}\rangle\le c_1\cdot\|\mathbf{y}-\tilde{\mathbf{y}}\|^2,$$

for some constant $c_1\in\mathbb{R}$.

For this aim we note that ϕ, ψ_i, ψ_j and η (w.r.t to its first argument) are all Lipschits continuous, that is, there exists some constant L s.t. for every $y,\tilde{y}\in\mathcal{G}$,

$$|\phi(y)-\phi(\tilde{y})|^2\le L\|y-\tilde{y}\|^2,\quad \|\psi_i(y)-\psi_i(\tilde{y})\|^2\le L\|y-\tilde{y}\|^2,\quad \|\psi_j(y)-\psi_j(\tilde{y})\|^2\le L\|y-\tilde{y}\|^2,\ \text{and}$$

$$\|\eta(y,t)-\eta(\tilde{y},t)\|^2\le L\|y-\tilde{y}\|^2\ \text{for all } t.$$

It follows that $\phi_j\circ\eta$ is Lipschitz for every t. In addition, ϕ_i and ϕ_j have bounded values, thus $\phi_i\cdot\phi_j$ is Lipschitz, and we conclude that the following function:

$$\phi(y)-\Big(1-\psi_j\big(\eta(y,\tau)\big)\Big)\psi_i(y)$$

is Lipschitz for every τ, and so is

$$\phi(\eta(y,\tau))-\phi(y).$$

Thus, the following set-valued map

$$\mathbf{F}_0(y,\omega,\tau):=\left\{\phi(y)-\Big(1-\psi_j\big(\eta(y,\tau)\big)\Big)\psi_i(y)\right\}\times\{\phi(\eta(y,\tau))-\phi(y)\}\times\{1\}$$

which is in fact a singleton for every (y,ω,τ), is Lipschitz, and as so satisfies the OSL condition.

For the discontinues part: for every y and τ, the function $\psi_i(y)\psi_j\big(\eta(y,\tau)\big)$ is positive and (uniformly) bounded by 1. Therefore $-\mathbf{I}(\omega)\cdot\psi_i(y)\psi_j\big(\eta(y,\tau)\big):\mathcal{D}\to 2^{\mathcal{G}}$ is OSL. Defining

$$\mathbf{F}_1(y,\omega,\tau):=\big\{-z\cdot\psi_i(y)\psi_j\big(\eta(y,\tau)\big)\ :\ \mathbf{I}(\omega)\big\}\times\{0\}\times\{0\}$$

we have that $\mathbf{F} = \mathbf{F}_0 + \mathbf{F}_1$ is a sum of OSL functions and thus it is OSL. For $\mathbf{y} = (y, \omega, \tau)$ define $F_1 : \mathcal{D} \to \mathcal{D}$ as $F_1(\mathbf{y}) = \big(\psi_i(y)\psi_j\big(\eta(y,\tau)\big), 0, 0\big)$, which, as explained, is Lipschitz. We have for all $\mathbf{y} \in \mathcal{D}$,

$$\begin{aligned}\sup\{\|\mathbf{z}\| : \mathbf{z} \in \mathbf{F}(\mathbf{y})\} &\leq \|\mathbf{F}_0(\mathbf{y})\| + \sup\{\|\mathbf{z}\| : \mathbf{z} \in \mathbf{F}_1(\mathbf{y})\} \\ &= \|\mathbf{F}_0(\mathbf{y})\| + \|F_1(\mathbf{y})\| \\ &= \|\mathbf{F}_0(\mathbf{y}) - \mathbf{F}_0(\mathbf{0}) + \mathbf{F}_0(\mathbf{0})\| + \|F_1(\mathbf{y}) - F_1(\mathbf{0})\| \\ &\leq c_2\|\mathbf{y} - \mathbf{0}\| + \|\mathbf{F}_0(\mathbf{0})\| + c_3\|\mathbf{y} - \mathbf{0}\| \\ &\leq c_0(1 + \|\mathbf{y}\|)\end{aligned}$$

for some constants $c_2, c_3 > 0$ and $c_0 \geq \max\{c_2 + c_3, \|\mathbf{F}_0(\mathbf{0})\|\}$ (note that $\mathbf{F}_0(\mathbf{y})$ for all $\mathbf{y}$ and $\mathbf{F}_1(\mathbf{0})$ are singletons in $\mathcal{D}$, thus slightly abusing notation). The steps follow because $\mathbf{F}_0$ and F_1 are Lipschitz (as single-valued functions) in $\mathcal{D}$, with $F_1(\mathbf{0}) = \mathbf{0}$. It follows (see book by Kunze, 2000) that the differential inclusion attains a unique solution.

APX-B. Proof of Lemma 2

First, we show that a unique solution to Equation (12) exists. By rearranging Equation (12) we get

$$(1-q)\theta = \lambda m \cdot \left(1 - e^{-q\theta\delta}\right). \tag{23}$$

A solution inside $[0,1]$ must exist because both sides of Equation (23) are continuous, positive functions of q and attain the value 0 at their minimum point over $q \in [0,1]$. Uniqueness follows from the fact that the LHS of (23) is (strongly) monotone decreasing and the RHS is monotone increasing in q.

Showing that the function qx satisfies Equation (11) is done by simple substitution: Assuming $Q^*(x) = Q^*_L(x) = Q^*_H(x) = qx$ for every $x \in [0,1)$ and $S^*_i(x) = s_i$ as defined in (13), basic algebraic manipulations yield $D_i(x) = (1 - e^{-q\theta\delta})\theta_i/\theta$ for each $i \in \{L, H\}$, therefore Equation (11) for every $x \in (0,1)$ and $i \in \{L, H\}$ takes the form

$$x \cdot (1-q)\frac{\theta_i}{m} - \lambda x \cdot (1 - e^{-q\theta\delta})\frac{\theta_i}{\theta} = 0$$

which is equivalent to (12) and therefore satisfies the equality.

APX-C. Proof of Lemma 3

Assuming $\theta > 0$, the equality in (12) can be rewritten as

$$q + \lambda \frac{1 - e^{-q\delta\theta}}{m\theta} = 1 \tag{24}$$

where λm^{-1} is a positive constant. The LHS is an increasing function in q. Differentiating it w.r.t. θ we arrive at

$$\frac{d}{d\theta}\left(q + \lambda \frac{1 - e^{-q\delta\theta}}{m\theta}\right) = \lambda \cdot \frac{e^{-q\delta\theta}}{m\theta^2}\left(1 + q\delta\theta - e^{q\delta\theta}\right)$$

and we note that $e^{q\delta\theta} > 1 + q\delta\theta$, thus, the LHS of Equation (24) is decreasing in θ. Therefore, if q_1 is the solution for (24) with $\theta = \theta_1$, and q_2 is the solution for (24) with $\theta = \theta_2 > \theta_1$, then it must hold that $q_2 > q_1$.

APX-D. Proof of Proposition 2

By Lemma 3, the revenue rate, $R^*(\theta) = (1 - q(\theta))r$ is decreasing with θ, and assuming $\kappa_L < \overline{R} = R^*(0)$, we have that an equilibrium $\theta^e > 0$ exists. We focus on showing uniqueness by dividing into three cases based on the value of $R^*(\Theta_L)$:

If $R^*(\Theta_L) < \kappa_L$, then, by monotonicity, there exists a unique $\theta^e \in (0, \Theta_L)$ satisfying $R^*(\theta^e) = \kappa_L$ and for every $\theta > \theta^e$, $R^*(\theta^e) < \kappa_L < \kappa_L$, therefore θ^e is the unique equilibrium, which corresponds to item (1).

If $R^*(\Theta_L) \in [\kappa_L, \kappa_H]$, then, for every $\theta < \Theta_L$, $R^*(\theta) > \kappa_L$, thus θ violates the equilibrium condition 9 for type L. In addition, every $\theta > \Theta_L$ is such that $R^*(\theta) < \kappa_H$, and therefore violates the equilibrium condition for type H. Hence, $\theta^e = \Theta_L$ is the unique equilibrium, which corresponds to item (2).

Finally, if $R^*(\Theta_L) > \kappa_H$, then either there exists a unique $\theta^e \in (\Theta_L, \Theta_H]$ satisfying $R^*(\theta^e) = \kappa_H$, or $R^*(\Theta_H) > \kappa_H$ in which case $\Theta = \Theta_L + \Theta_H$ is the unique equilibrium, which corresponds to item (3).

APX-E. Proof of Lemma 4

Similar to the proof of Lemma 1.

APX-F. Proof of Lemma 5

First we show that there exists a unique triple $(\hat{q}_L, \hat{q}_H, \hat{q}) \in [0,1]^3$ solving System (15). Note that $\hat{q}$ solves Equation (12) and therefore by Lemma 2 exists uniquely in the interval $[0,1]$. Given the solution $\hat{q}$, the pair $(\hat{q}_L, \hat{q}_H)$ either solves

$$\begin{cases} \theta_L \hat{q}_L + \theta_H \hat{q}_H = \theta \hat{q}, \\ 1 - \hat{q}_H = \lambda m \cdot \left(1 - e^{-\hat{q}_H \theta_H \delta}\right)/\theta_H. \end{cases} \tag{25}$$

or solves

$$\begin{cases} \theta_L \hat{q}_L + \theta_H \hat{q}_H = \theta \hat{q}, \\ 1 - \hat{q}_H = (1 - \hat{q}_L)\kappa_H/\kappa_L. \end{cases} \tag{26}$$

Assume that $\hat{q}$ is fixed, and consider System (25). We show that this system has a unique solution in $[0,1]^2$: By Lemma 2, the second equation of (25) possesses a unique solution in the interval $[0,1]$. Denote this solution by a_H, and define $a_L = \frac{\theta \hat{q} - \theta_H a_H}{\theta_L}$, then (a_L, a_H) uniquely solve (25). By substituting $\hat{q} = 1 - \lambda m \cdot \left(1 - e^{-\hat{q}\theta\delta}\right)/\theta$ and $a_H = 1 - \lambda m \cdot \left(1 - e^{-a_H \theta_H \delta}\right)/\theta_H$ in the definition of a_L we get, after rearranging,

$$a_L = 1 - \frac{\lambda m}{\theta_L}\left(e^{-a_H \theta_H \delta} - e^{-\hat{q}\theta\delta}\right).$$

By Lemma 3 and the fact that $\theta_H < \theta$, it must hold that $a_H < \hat{q}$, and therefore $e^{-a_H \theta_H \delta} > e^{-\hat{q}\theta\delta}$, thus, $a_L < 1$. Moreover, $\hat{q}$ is a convex combination of a_L and a_H, thus $a_H < \hat{q} < a_L$, therefore $a_L \in [0,1]$.

Consider now System (26) given $\hat{q}$. System (26) is linear and non-singular, thus the existence of a unique solution (b_L, b_H) trivially follows. Since $\kappa_H/\kappa_L > 1$ and since $\hat{q}$ is a convex combination of b_L and b_H, this solution satisfies $b_H < \hat{q} < b_L$. Thus, $b_H < 1$, and so $b_L = 1 - (1 - b_H)\kappa_L/\kappa_H < 1$, meaning that $b_L \in [0,1]$.

From System (15), a solution $(\hat{q}_L, \hat{q}_H, \hat{q})$ must satisfy $\hat{q}_H = \max\{a_H, b_H\}$ which is positive due to the fact that $a_H > 0$, and less than unity because both $a_H < 1$ and $b_H < 1$. As both (a_L, a_H) and (b_L, b_H) satisfy the first equation in (15) with the same value $\hat{q}$, it follows that $a_H \geq b_H$ iff $a_L \leq b_L$.

Therefore $\hat{q}_L = \min\{a_L, b_L\} \in [0,1]$. We conclude that a unique solution $(\hat{q}_L, \hat{q}_H, \hat{q})$ to System (15) exists, and lies in $[0,1]^3$.

We now confirm that the two functions $\hat{Q}_i^*(x) = \hat{q}_i x, i \in \{L, H\}$ satisfy the mean field steady-state equation (14): First we note that if $a_H > b_H$, then $a_L < b_L$, therefore $\hat{q}_i = a_i, i \in \{L, H\}$. In this case we have that

$$\frac{\hat{R}_H}{\kappa_H} = \frac{1-\hat{q}_H}{\kappa_H} r < \frac{1-b_H}{\kappa_H} r = \frac{1-b_L}{\kappa_L} r < \frac{1-\hat{q}_L}{\kappa_L} r = \frac{\hat{R}_L}{\kappa_L},$$

meaning that $\hat{s}_H = 1$, and therefore $\hat{D}_H(x) = 1 - e^{-\hat{q}_H \theta_H \delta}$. Equation (14) with $i = H$ then becomes

$$x \cdot (1-\hat{q}_H)\frac{\theta_H}{m} - \lambda x \cdot (1 - e^{-\hat{q}_H \theta_H \delta}) = 0 \tag{27}$$

which, when substituting $\hat{q}_H = a_H$, is clearly satisfied by definition of a_H. For $i = L$, we have that $\hat{s}_L = 0$, thus $\hat{D}_L(x) = (1 - e^{-\hat{q}_L \theta_L \delta}) e^{-\hat{q}_H \theta_H \delta}$. Note from the definition of $\hat{q}$ we have, for every $x \in (0,1]$,

$$x \cdot (1-\hat{q})\frac{\theta}{m} - \lambda x \cdot (1 - e^{-\hat{q}\theta\delta}) = 0. \tag{28}$$

Subtracting (27) from the equation above we arrive at

$$\frac{x}{m} \cdot (\theta - \theta_H - \theta\hat{q} + \theta_H \hat{q}_H) - \lambda x \cdot \left(1 - e^{-\hat{q}\theta\delta} - 1 + e^{-\hat{q}_H \theta_H \delta}\right) = 0,$$

and with $\theta_L = \theta - \theta_H$ and $\theta_L \hat{q}_L = \theta\hat{q} - \theta_H \hat{q}_H$ this becomes equivalent to

$$x \cdot (1-\hat{q}_L)\frac{\theta_L}{m} - \lambda x \cdot (1 - e^{-\hat{q}_L \theta_L \delta}) e^{-\hat{q}_H \theta_H \delta} = 0,$$

which coincides with Equation (14) for $i = L$, therefore proving that $(\hat{Q}_L^*(x), \hat{Q}_H^*(x))$ is the unique steady-state solution under the assumption that $a_H > b_H$.

On the other hand, if $a_H \le b_H$, then $\hat{q}_i = b_i, i \in \{L, H\}$. Under this assumption,

$$\frac{\hat{R}_H}{\kappa_H} = \frac{1-\hat{q}_H}{\kappa_H} r = \frac{1-\hat{q}_L}{\kappa_L} r = \frac{\hat{R}_L}{\kappa_L},$$

and therefore, for $i, j \in \{L, H\}$ s.t. $i \neq j$, $\hat{s}_i$ can be written as

$$\hat{s}_i = \frac{\beta_i}{1 - e^{-\hat{q}_i \theta_i \delta}} + \frac{\beta_j e^{-\hat{q}_j \theta_j \delta}}{1 - e^{-\hat{q}_j \theta_j \delta}}$$

where $\beta_i = \theta_i\kappa_i/(\theta_L\kappa_L + \theta_H\kappa_H)$, noting that $\beta_j = 1-\beta_i$. Basic algebraic manipulations then yield $D_i(x) = (1-e^{-\hat{q}\theta\delta})\beta_i$, and so, the drift term at each point $x \in [0,1)$ is given by

$$\begin{aligned}
\frac{d\hat{Q}_i^*}{dt} &= x\cdot(1-\hat{q}_i)\frac{\theta_i}{m} - \lambda x\cdot(1-e^{-\hat{q}\theta\delta})\beta_i \\
&= \frac{x}{m}\cdot((1-\hat{q}_i)\theta_i - (1-\hat{q})\theta\beta_i) \\
&= x\theta_i\cdot\frac{(1-\hat{q}_i)(\theta_i\kappa_i+\theta_j\kappa_j) - (1-\hat{q}_i)\theta_i\kappa_i - (1-\hat{q}_j)\theta_j\kappa_i}{m\cdot(\theta_i\kappa_i+\theta_j\kappa_j)} \\
&= x\theta_i\cdot\frac{(1-\hat{q}_i)(\theta_i\kappa_i+\theta_j\kappa_j) - (1-\hat{q}_i)\theta_i\kappa_i - (1-\hat{q}_i)\theta_j\kappa_j}{m\cdot(\theta_i\kappa_i+\theta_j\kappa_j)} = 0
\end{aligned}$$

for each $i,j \in \{L,H\}, i \neq j$. The second equality follows from (28), the third from $(1-\hat{q})\theta = (1-\hat{q}_L)\theta_L + (1-\hat{q}_H)\theta_H$, and the forth from $(1-\hat{q}_L)\kappa_H = (1-\hat{q}_H)\kappa_L$ which is an implication of $\hat{q}_i = b_i, i \in \{L,H\}$. Hence Equation (14) is satisfied.

APX-G. Proof of Lemma 6

Our characterization of the participation intensities implies that for any $\theta \in (0, \Theta_L+\Theta_H]$, the type dependent participation intensities are given by $\theta_L = \min\{\theta, \Theta_L\}$ and $\theta_H = \{\theta-\Theta_L\}^+$. Suppose $\theta \in (0,\Theta_L]$, thus, $\theta_L = \theta$ and $\theta_H = 0$. In this case the two policies are equivalent and therefore it follows by Lemma 3 that the proportion of available drivers is increasing in θ. We therefore focus on $\theta > \Theta_L$, in which case $\theta_L = \Theta_L$ is constant w.r.t to θ, and $\theta_H = \theta - \Theta_L$.

First, we note from Lemma 3, that $\hat{q}(\theta)$ is increasing. Given $\theta \in (\Theta_L, \Theta_L+\Theta_H]$, let $(a_L(\theta), a_H(\theta))$ be the solution of

$$\begin{cases} \Theta_L a_L + (\theta-\Theta_L)a_H = \theta\hat{q}(\theta), \\ 1-a_H = \lambda m\cdot\left(1-e^{-a_H\cdot(\theta-\Theta_L)\delta}\right)/(\theta-\Theta_L), \end{cases} \tag{29}$$

and let $(b_L(\theta), b_H(\theta))$ be the solution of

$$\begin{cases} \Theta_L b_L + (\theta-\Theta_L)b_H = \theta\hat{q}(\theta), \\ 1-b_H = (1-b_L)\kappa_H/\kappa_L. \end{cases} \tag{30}$$

It is shown in APX-F that each solution pair $(a_L(\theta), a_H(\theta))$ and $(b_L(\theta), b_H(\theta))$ exists uniquely and that $\hat{q}_L(\theta) = \min\{a_L(\theta), b_L(\theta)\}$ and $\hat{q}_H(\theta) = \max\{a_H(\theta), b_H(\theta)\}$. In addition, for every θ, $a_H(\theta) <$

$\hat{q}(\theta) < a_L(\theta)$, and similarly, $b_H(\theta) < \hat{q}(\theta) < b_L(\theta)$. We show that $a_L(\theta), a_H(\theta), b_L(\theta)$ and $b_H(\theta)$ are all increasing in θ.

We start with the monotonicity of $a_L(\theta)$ and $a_H(\theta)$. Consider the second equation of (29). From Lemma 3 we have that $a_H(\theta)$ is increasing. In order to show that $a_L(\theta)$ is increasing we show first that the function $g(\theta) = \theta\hat{q}(\theta)$ is increasing and convex. That $g(\theta)$ is increasing trivially follows from θ and $\hat{q}(\theta)$ being positive increasing (see Lemma 3). By Equation (12) it holds that

$$g(\theta) = \theta - \frac{1}{\lambda m}(1 - e^{-\delta g(\theta)}),$$

and simple algebraic manipulations yield

$$\frac{d}{d\theta} g(\theta) = \left(1 + \frac{\delta}{\lambda m} e^{-\delta g(\theta)}\right)^{-1}.$$

Note that $e^{-\delta g(\theta)}$ is decreasing and positive and therefore the RHS of the latter equation is increasing, meaning that $g(\theta)$ is convex. Now note from (29) that $a_H(\theta) = \hat{q}(\theta - \Theta_L)$ and that

$$a_L(\theta) = \frac{\theta\hat{q}(\theta) - (\theta - \Theta_L)a_H(\theta)}{\Theta_L} = \frac{g(\theta) - g(\theta - \Theta_L)}{\Theta_L}.$$

Because $g(\theta)$ is convex, $a_L(\theta)$ is increasing.

We now turn our attention to the pair $(b_L(\theta), b_H(\theta))$. Denote $\alpha = \kappa_H/\kappa_L$, thus from the second equation in (30) we have $b_H(\theta) = 1 - \alpha + \alpha b_L(\theta)$, by which $b_H(\theta)$ is increasing iff $b_L(\theta)$ is increasing. Substituting this in the first equation of (30) we get, by rearranging,

$$(\Theta_L + \alpha(\theta - \Theta_L))b_L(\theta) = \theta\hat{q}(\theta) + (\alpha - 1)(\theta - \Theta_L).$$

Taking derivative w.r.t θ, after algebra we obtain

$$\frac{d}{d\theta} b_L(\theta) = \frac{\theta \frac{d}{d\theta}\hat{q}(\theta) + \alpha(1 - b_L(\theta)) - (1 - \hat{q}(\theta))}{\Theta_L + \alpha(\theta - \Theta_L)}.$$

Note first that $b_H(\theta) < \hat{q}(\theta)$, which implies $1 - \hat{q}(\theta) < 1 - b_H(\theta) = \alpha(1 - b_L(\theta))$, and secondly, $(d/d\theta)\hat{q}(\theta) > 0$ which is from Lemma 3, hence $(d/d\theta)b_L(\theta) > 0$. Thus, $b_L(\theta)$ is increasing, and so is $b_H(\theta)$.

Finally, because $a_L(\theta)$ and $b_L(\theta)$ are increasing, $\hat{q}_L(\theta) = \min\{a_L(\theta), b_L(\theta)\}$ is increasing, and because $a_H(\theta)$ and $b_H(\theta)$ are increasing, $\hat{q}_H(\theta) = \max\{a_H(\theta), b_H(\theta)\}$ is increasing.

APX-H. Proof of Proposition 3

From Lemma 6, $\hat{R}_L^*(\theta)$ and $\hat{R}_H^*(\theta)$ are decreasing in θ, and we further have that $\hat{R}_L^*(\theta) < \hat{R}_H^*(\theta)$ for all θ. Then exactly one of the following two cases must hold:

- If $\hat{R}_H^*(\Theta_L) < \kappa_H$, then for all $\theta > \Theta_L$, $\hat{R}_H^*(\theta) \le \kappa_H$, therefore type-$H$ drivers do not participate in equilibrium. In addition,

  —if $\hat{R}_L^*(\Theta_L) > \kappa_L$, then Θ_L is the unique equilibrium participation intensity, and

  —if $\hat{R}_L^*(\Theta_L) \le \kappa_L$, then there exists a unique $\theta \in (0, \Theta_L]$ such that $\hat{R}_L^*(\theta) = \kappa_L$, and this value θ is a unique equilibrium participation intensity.

- If $\hat{R}_H^*(\Theta_L) > \kappa_H$, then either

  —$\hat{R}_H^*(\Theta_L + \Theta_H) \ge \kappa_H$, and $\Theta_L + \Theta_H = \Theta$ is the unique equilibrium participation intensity, or

  —if $\hat{R}_H^*(\Theta_L + \Theta_H) < \kappa_H$, and therefore there exists a unique equilibrium participation intensity $\theta \in (\Theta_L, \Theta_L + \Theta_H]$, which satisfies $\hat{R}_H^*(\theta) = \kappa_H$.

APX-I. Proof of Lemma 7

As explained, under both policies, participation of any type-H drivers implies that all type-L drivers participate. Therefore it suffices to show that in equilibrium under MinWeightRev, more type-H drivers participate than under MinRev. To show this, we shall prove that for every θ, $\hat{R}_H^*(\theta) > R^*(\theta)$. From the proof of Lemma 5, we have that $\hat{q}_H(\theta) < \hat{q}(\theta)$, thus,

$$\hat{R}_H^*(\theta) = (1 - \hat{q}_H(\theta))r > (1 - \hat{q}(\theta))r = (1 - q(\theta))r = R^*(\theta)$$

From our equilibrium condition (9), it follows that $\hat{\theta}^e \ge \theta^e$.

APX-J. Proof of Proposition 4

To prove the first item, first note that $\Phi \ge 1$ is a direct implication of Lemma 7. Assume by way of contradiction that $\hat{\theta}^e > 2\theta^e$. Because the equilibrium participation of type-L drivers is the same under both policies, it must be that the surge in participation is due to increased participation of type-H drivers, i.e.,

$$\hat{\theta}_H^e = \hat{\theta}^e - \Theta_L > 2\theta^e - \Theta_L = 2\theta_H^e + 2\Theta_L - \Theta_L = 2\theta_H^e + \Theta_L \ge \theta^e.$$

It also implies that under MinRev, not all type-H drivers participate in equilibrium, namely $R^*(\theta^e) \leq \kappa_H$.

Recall the definition of $(a_L(\theta), a_H(\theta))$ and $(b_L(\theta), b_H(\theta))$ from APX-G, which satisfy $\hat{q}_H(\theta) = \max\{a_H(\theta), b_H(\theta)\}$, for all θ. It follows that

$$\hat{q}_H(\hat{\theta}^e) \geq a_H(\hat{\theta}^e) = q(\hat{\theta}^e - \Theta_L) > q(\theta^e)$$

where $q(\theta)$ is defined in (12).

$$\hat{R}^*_H(\hat{\theta}^e) = (1 - \hat{q}_H(\hat{\theta}^e))r < (1 - q(\theta^e))r = R^*(\theta^e) \leq \kappa_H$$

which is a contradiction to the assumption that $\hat{\theta}^e$ induces equilibrium under MinWeightRev.

To show the second item we first note that

$$\lambda^*(\theta) = \lambda \cdot (1 - e^{-\delta\theta q(\theta)}) = \frac{\theta - \theta q(\theta)}{m}$$

which, given the participation intensity θ, is independent of the policy. The second equality follows from Equation (12). It has been proven in APX-G that the function $g(\theta) = \theta q(\theta)$ is positive increasing and (strongly) convex in θ, therefore $\lambda^*(\theta)$ is concave. In addition, $\lambda^*(\theta)$ is increasing as it is a composition of increasing functions, $f(x) = 1 - e^{-\delta x}$ over $g(\theta)$. It follows that for any $\phi \in (1,2]$ and for any θ, $\lambda^*(\phi^{-1}\theta) > \phi^{-1}\lambda^*(\theta)$, and therefore

$$1 \leq \frac{\lambda^*(\theta)}{\lambda^*(\phi^{-1}\theta)} < \frac{\lambda^*(\theta)}{\phi^{-1}\lambda^*(\theta)} = \phi.$$

Thus, whenever $\Phi > 1$ we have that $\Psi \in [1, \Phi) \subset [1, 2)$, and because for every $\phi \in [1,2]$, $\lim_{\theta \to 0} \lambda^*(\theta)/(\lambda^*(\phi^{-1}\theta)) = \phi$, we have that $\lim_{\Theta \to 0} \Psi = 2$ iff $\lim_{\Theta \to 0} \Phi = 2$.

APX-K. Different formulation for mean field equations under MinRev

A different formulation of the mean field equations can be regarded for the MinRev that neither requires any knowledge of driver's revenues or employing a prioritization scheme:

For all $x \in [0,1)$, $t \in [0, \infty)$, and $i \in \{L, H\}$,

$$\frac{\partial Q_i(x;t)}{\partial t} = \frac{x}{m}(1 - Q_i(1;t)) - \frac{\lambda}{\theta_i} \int_{s=0}^{x} \left(1 - e^{-Q'(s;t)\theta\delta}\right) D_i(s;t) ds, \tag{31}$$

where

$$D_i(s;t) = \begin{cases} \frac{\theta_i Q_i'(s;t)}{\theta Q'(s;t)} & \text{if } Q'(s;t) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

The interpretation is that when all drivers are symmetric in the eyes of the platform, the type of a matched drivers is determined by a random draw of a driver among all candidates. The probability of the number of the total number candidates being non-zero corresponds to $1 - e^{-Q'(s;t)\theta\delta}$, and conditioned on this event, the probability of the type of the matched driver being i is given by $\frac{\theta_i Q_i'(s;t)}{\theta Q'(s;t)}$. More intuition is generated by Lemma 8 below. In this formulations the drift is Lipschitz continuous.

LEMMA 8. *Consider an urn with a random non-negative integer number B of blue balls and a random non-negative integer number R of red balls. Assume that $B \sim \text{Pois}(b)$ and $R \sim \text{Pois}(r)$ for some $r, b > 0$. Conditioned on the urn being non empty, we pick a ball uniformly at random. Then the probability of picking a blue ball is $\frac{b}{b+r}$.*

Before proving the result algebraically, we provide and intuitive explanation: Assume for simplicity that r and b are rational numbers, so that there exists some $\lambda > 0$ such that λb and λr are both integers. Then B can be written as a sum $B = B_1 + \cdots + B_{\lambda b}$ and similarly $R = R_1 + \cdots + R_{\lambda r}$, where $\{B_i\}, \{R_j\}$ are all i.i.d Poisson random variables with parameter λ^{-1}. Imagine now that the urn contains a random number B_i of blue balls with the number i on them, $i = 1, \ldots, \lambda b$, and R_j red balls with the number j on them, $j = 1, \ldots, \lambda r$, and given the urn is not empty we draw one uniformly at random. It is clear from symmetry that the chances of drawing any possible combination of color and number is the same for all combination, and is equal $\frac{1}{\lambda \cdot (b+r)}$. Since exactly λb of the combinations contain the color blue, it follows that the probability of drawing a blue ball is $\frac{\lambda b}{\lambda \cdot (b+r)} = \frac{b}{b+r}$.

A straightforward algebraic proof can be provided: conditioning on the urn having exactly one ball in it, from the branching property of Poisson, the probability of its color being blue is $\frac{b}{b+r}$.

When the number of balls is $n > 1$, the colors of two different balls are independent, hence the probability of a randomly chosen ball having the color blue is invariant of n.

Acknowledgments

SRIBD presidential post-doctoral fellowship

References

Afeche P, Liu Z, Maglaras C (2018) Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. *Columbia Business School Research Paper* (18-19):18–19.

Bai J, So KC, Tang CS, Chen X, Wang H (2019) Coordinating supply and demand on an on-demand service platform with impatient customers. *Manufacturing & Service Operations Management* 21(3):556–570.

Banerjee S, Riquelme C, Johari R (2015) Pricing in ride-share platforms: A queueing-theoretic approach. *Available at SSRN 2568258* .

Benjaafar S, Hu M (2020) Operations management in the age of the sharing economy: what is old and what is new? *Manufacturing & Service Operations Management* 22(1):93–101.

Besbes O, Castro F, Lobel I (2020) Surge pricing and its spatial supply response. *Management Science* .

Bimpikis K, Candogan O, Saban D (2019) Spatial pricing in ride-sharing networks. *Operations Research* 67(3):744–769.

Braverman A, Dai JG, Liu X, Ying L (2019) Empty-car routing in ridesharing systems. *Operations Research* .

Cachon GP, Daniels KM, Lobel R (2017) The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management* 19(3):368–384.

Department of Transportation N (2019) Improving efficiency and managing growth in new york's for-hire vehicle sector. *NYC Taxi and Limousine Commission and Department of Transportation, Final Report* .

Gast N, Gaujal B (2012) Markov chains with discontinuous drifts have differential inclusion limits. *Performance Evaluation* 69(12):623–642.

Gurvich I, Lariviere M, Moreno A (2019) Operations in the on-demand economy: Staffing services with self-scheduling capacity. *Sharing economy*, 249–278 (Springer).

Hassin R (2016) *Rational queueing* (CRC press).

Hassin R, Haviv M (2003) *To queue or not to queue: Equilibrium behavior in queueing systems*, volume 59 (Springer Science & Business Media).

Hu M (2019) *Sharing economy: making supply meet demand* (Springer).

Hu M (2020) From the classics to new tunes: A neoclassical view on sharing economy and innovative marketplaces. *Production and Operations Management* .

Iglesias R, Rossi F, Zhang R, Pavone M (2019) A bcmp network approach to modeling and controlling autonomous mobility-on-demand systems. *The International Journal of Robotics Research* 38(2-3):357–374.

Kunze M (2000) *Non-smooth dynamical systems*, volume 1744 (Springer Science & Business Media).

Özkan E (2020) Joint pricing and matching in ride-sharing systems. *European Journal of Operational Research* 287(3):1149–1160.

Ozkan E, Ward A (2017) Dynamic matching for real-time ridesharing. *Available at SSRN 2844451* .

Parrott JA, Reich M (2018) An earnings standard for new york city's app-based drivers. *New York: The New School: Center for New York City Affairs* .

Taylor TA (2018) On-demand service platforms. *Manufacturing & Service Operations Management* 20(4):704–720.

Tsitsiklis JN, Xu K (2012) On the power of (even a little) resource pooling. *Stochastic Systems* 2(1):1–66.

Xu J, Hajek B (2013) The supermarket game. *Stochastic Systems* 3(2):405–441.