

Ridesharing with Driver Location Preferences*

Duncan Rheingans-Yoo¹, Scott Duke Kominers², Hongyao Ma³ and David C. Parkes³

¹Harvard College

²Harvard Business School and Department of Economics

³Harvard School of Engineering and Applied Sciences

{d_rheingansyoo@college, kominers@fas, hma@seas, parkes@eecs}.harvard.edu

Abstract

We study revenue-optimal pricing and driver compensation in ridesharing platforms when drivers have heterogeneous preferences over locations. If a platform ignores drivers' location preferences, it may make inefficient trip dispatches; moreover, drivers may strategize so as to route towards their preferred locations. In a model with stationary and continuous demand and supply, we present a mechanism that incentivizes drivers to both (i) report their location preferences truthfully and (ii) always provide service. In settings with unconstrained driver supply or symmetric demand patterns, our mechanism achieves the full-information, first-best revenue. Under supply constraints and unbalanced demand, we show via simulation that our mechanism improves over existing mechanisms and has performance close to the first-best.

1 Introduction

Uber connected its first rider to a driver in the summer of 2009,¹ and since then, ridesharing platforms have dramatically changed the way people get around in urban areas. Ridesharing platforms allow a wide array of people to become drivers and—in contrast to traditional taxi systems—use dynamic “surge pricing” at times when demand exceeds supply. Properly designed, dynamic pricing improves system efficiency [Castillo *et al.*, 2017], increases driver supply [Chen and Sheldon, 2015], and makes the system reliable for riders [Hall *et al.*, 2015].

A growing literature studies how to structure prices for riders and compensation for drivers so as to optimally account for variation in supply and demand [Banerjee *et al.*, 2015; Bimpikis *et al.*, 2016; Castillo *et al.*, 2017; Ma *et al.*, 2019]. However, existing models leave aside driver heterogeneity. In practice, some drivers may prefer to drive in the city and others in the suburbs, and many prefer to end their days close to home. A matching system that treats drivers as homogeneous makes inefficient dispatches, with drivers preferring to fulfill each other's dispatches instead of their own.

The problem goes beyond simple efficiency loss. A core feature of ridesharing platforms is that drivers retain the flexibility to choose when and where to provide service. Every ride the platform proposes needs to be accepted voluntarily, forming an optimal response for the driver [Ma *et al.*, 2019]. This incentive alignment simplifies participation for drivers and also makes behavior more predictable. Without accounting for heterogeneity, a platform cannot fully understand a driver's preferences or achieve full incentive alignment.

Indeed, platforms have experimented with methods to incorporate driver heterogeneity. As of Summer 2019, Uber allows drivers to indicate—twice a day—that they would like to take trips in the direction of a particular location.² However, mechanisms that account for driver preferences can also have unintended consequences if not designed properly. By saying “I want to drive South,” a driver biases her dispatches in a way that could in principle promote more profitable trips.³

In this paper, we introduce the study of driver location preference in a mechanism design framework. In Section 2, we adapt a model originally conceived by Bimpikis *et al.* [2016] to an economy where drivers prefer a particular location. In Section 3, we present the *Preference-Attentive Ridesharing Mechanism* (PARM), which elicits driver preferences and sets a revenue-optimal pricing policy. We show that PARM is incentive-compatible, and that it achieves the full-information, first-best revenue when supply is unconstrained or when demand is symmetric. In Section 4, we study settings with constrained supply and asymmetric demand, using simulations to compare the revenue and welfare performance of PARM to existing ridesharing mechanisms. We show that PARM achieves close to first-best revenue and typically outperforms even the best case for preference-oblivious pricing (where strategic behavior hurts efficiency). Proofs not presented in the text are deferred to Appendix A of the full paper.

1.1 Related Work

Existing research on pricing and dispatching in ridesharing platforms does not account for driver heterogeneity.

We build on work of Bimpikis *et al.* [2016], who show that under a continuum model, and with stationary demand and

*The full version of this paper is available at arXiv:1905.13191.

¹<https://www.uber.com/newsroom/history/>, visited 02/25/2019.

²<https://help.uber.com/partners/article/set-a-driver-destination?nodeId=f3df375b-5bd4-4460-a5e9-afd84ba439b9>, visited 2/25/19.

³<https://therideshareguy.com/uber-drops-destination-filters-back-to-2-trips-per-day/>, visited 2/25/19.

unlimited supply, a ridesharing platform’s revenue is maximized when the demand pattern across different locations is balanced. They show via simulation that in comparison to setting a uniform price for all locations, pricing trips differently depending on trip origins improves revenue. Relative to the Bimpikis *et al.* [2016] model, we allow limited driver supply; moreover, each driver in our model has a preferred location. We thus introduce a reporting phase, in which drivers report their preferred locations. We then modify the matching and pricing formulations in order to align incentives.

Ma *et al.* [2019] study the incentive alignment of drivers in the presence of spatial imbalance and temporal variation of supply and demand. Castillo *et al.* [2017] show that dynamic pricing mitigates inefficient “wild goose chase” phenomena for platforms that employ myopic dispatching strategies. Modeling a shared vehicle system as a continuous-time Markov chain, Banerjee *et al.* [2017] establish approximation guarantees for a static, state-independent pricing policy. Ostrovsky and Schwarz [2019] study the economy of self-driving cars, focusing on car-pooling and market equilibrium. Queuing-theoretic approaches have also been adopted: Banerjee *et al.* [2015] show the robustness of dynamic pricing; Afèche *et al.* [2018] study the impact of driver autonomy and platform control; and Besbes *et al.* [2018] analyze the relationship between capacity and performance.

There are various empirical studies, analyzing the impact of dynamic pricing [Hall *et al.*, 2015; Chen and Sheldon, 2015], the labor market for drivers [Hall and Krueger, 2016; Hall *et al.*, 2017], consumer surplus [Cohen *et al.*, 2016], the value of flexible work [Chen *et al.*, 2017], the gender earnings gap [Cook *et al.*, 2018], and the commission vs. medallion lease-based compensation models [Angrist *et al.*, 2017].

2 Model

We consider a discrete time, infinite horizon model of a ridesharing network with discrete locations, $L = \{1, \dots, n\}$. Following the baseline model of Bimpikis *et al.* [2016], we assume unit distances, i.e., it takes one period of time to travel in between any pair of locations. At the beginning of each time period, for each location $i \in L$, there is a continuous mass $\theta_i \geq 0$ of riders requesting trips from i . The fraction of riders at i with destination $j \in L$ is given by $\alpha_{ij} \in [0, 1]$ (thus $\sum_{j \in L} \alpha_{ij} = 1$). We assume that the components of rider demand $\theta = \{\theta_i\}_{i \in L}$ and $\alpha = \{\alpha_{ij}\}_{i,j \in L}$ are stationary and do not change over time. Riders’ willingness to pay for trips are i.i.d. random variables with CDF F . Thus, for any $i, j \in L$, the number of trips demanded from i to j at price $p_{ij} \geq 0$ would be $\theta_i \alpha_{ij} (1 - F(p_{ij}))$. (Riders who are unwilling to pay the stated prices for their rides leave the market.)

Each driver has a preferred location $\tau \in L$. Drivers receive $I \geq 0$ additional utility whenever they start a period in their preferred locations (irrespective of whether they have a rider); this preferred location is private information and represents a driver’s *type*. For each location $\tau \in L$, the total mass of available drivers of type τ is given by $s^{(\tau)} \geq 0$. Drivers have a discount factor of $\delta \in (0, 1)$, and an outside option that delivers utility $w \geq 0$.⁴ We assume $\sum_{t=0}^{\infty} \delta^t I = I/(1 - \delta) < w$,

⁴Throughout the paper, we consider δ to be very close to 1—this

meaning that the utility from being in one’s favorite location at all times does not outweigh the outside opportunity.

A *ridesharing mechanism* elicits drivers’ preferred locations, matches drivers and riders to trips, sets riders’ trip prices and drivers’ compensation, and (potentially) imposes drivers’ penalties for strategic behavior. Before the beginning of the first time period, the mechanism elicits the preferred locations from potential drivers.

At the beginning of each period, a driver whose previous trip ended at location i chooses whether or not to provide service at location i . If a driver provides service, the mechanism may dispatch that driver to (i) pick up some rider with trip origin i , (ii) relocate to some location, or (iii) stay in the same location. If a rider going from i to j is picked up by some driver, then the rider pays the platform the trip price $p_{ij} \geq 0$. If a driver of reported type τ is dispatched from i to j , the platform pays them $c_{ij}^{(\tau)} \geq 0$, regardless of if her dispatch was to pick up a rider or relocate.⁵ Drivers who choose not to provide service in a period can relocate to any location j in the network, are not compensated by the mechanism in this period, and may be charged a penalty P_j .⁶ Denote $p \triangleq \{p_{ij}\}_{i,j \in L}$, $c \triangleq \{c_{ij}^{(\tau)}\}_{i,j,\tau \in L}$ and $P \triangleq \{P_j\}_{j \in L}$.

Based on rider demand (θ, α) and the reported supply of drivers of each type, a mechanism determines rider and driver flow, trip prices p , driver compensation c , and driver penalties P . Drivers decide whether or not to participate, considering pricing, penalties, and their outside options. Given entry decisions by drivers, and decisions made by drivers since entry, the platform then dispatches drivers to trips and processes payments and penalties accordingly in each period.

2.1 Steady-State Equilibrium

In this section we will analyze a steady-state equilibrium while ignoring penalties. This will be helpful because it establishes that under truthful reporting of types, drivers will always follow the proposed dispatches. We will separately handle incentives to report truthful types, considering the effect of penalties on aligning these incentives. It will turn out that drivers are only charged penalties for the first time they deviate, not for future deviations.

At the beginning of each period, let $x_i^{(\tau)}$ be the number of drivers of reported type τ at location i , and let $x^{(\tau)} \triangleq \sum_{i \in L} x_i^{(\tau)}$ denote the total number of drivers of reported type τ on the platform. Let the *trip flow* be $f \triangleq \{f_{ij}^{(\tau)}\}_{i,j,\tau \in L}$, where $f_{ij}^{(\tau)} \geq 0$ is the number of riders from i to j assigned to drivers of type τ . Let $y_{ij}^{(\tau)} \geq 0$ be the mass of drivers

is natural, since an annual interest rate of 4% implies an exponential discount factor of 0.9999992 over the course of ten minutes.

⁵It bears mentioning that $c_{ij}^{(\tau)}$ need not be a fixed proportion of p_{ij} . In fact, Bimpikis *et al.* [2016] find that for certain types of networks, making driver compensation a fixed proportion of trip price drastically reduces platform revenue.

⁶Drivers only choose whether to provide service at a location and cannot decline dispatches based on the trip destination. This is consistent with current ridesharing platforms, which hide trip destinations because of concern that drivers will cherry pick rides.

of type τ at i who are dispatched to relocate to j without a rider, and set $x \triangleq \{x_i^{(\tau)}\}_{i,\tau \in L}$ and $y \triangleq \{y_{ij}^{(\tau)}\}_{i,j,\tau \in L}$. No driver or rider can be matched multiple times in the same period, so assuming drivers always provide service, we have $\sum_{j \in L} f_{ij}^{(\tau)} + y_{ij}^{(\tau)} \leq x_i^{(\tau)}$ for all $i \in L$ and all $\tau \in L$, and $\sum_{\tau \in L} f_{ij}^{(\tau)} \leq \theta_i \alpha_{ij} (1 - F(p_{ij}))$ and for all $i, j \in L$.

For a trip with origin i and destination j , if the total rider demand exceeds driver supply (i.e., if $\sum_{\tau \in L} f_{ij}^{(\tau)} < \theta_i \alpha_{ij} (1 - F(p_{ij}))$), the mechanism may increase the trip price p_{ij} and achieve higher revenue. Therefore for revenue optimization, we can assume without loss that $\sum_{\tau \in L} f_{ij}^{(\tau)} = \theta_i \alpha_{ij} (1 - F(p_{ij}))$. When $x_i^{(\tau)} > 0$, meaning that some drivers with reported type τ are at location i , the probability that a given driver of reported type τ is dispatched to destination j is $(f_{ij}^{(\tau)} + y_{ij}^{(\tau)})/x_i^{(\tau)}$. Assuming a driver of type τ was truthfully reported her type and will provide service in all periods, her lifetime expected utility for starting from location i is of the form

$$\pi_i^{(\tau)} = \sum_{j \in L} (c_{ij}^{(\tau)} + \delta \pi_j^{(\tau)}) \frac{f_{ij}^{(\tau)} + y_{ij}^{(\tau)}}{x_i^{(\tau)}} + I \cdot \mathbb{1}\{i = \tau\}, \quad (1)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function. The first term in (1) is the expected compensation and future utility a driver gets when dispatched to one of the n possible destinations. The second term corresponds to the idiosyncratic utility drivers get from starting trips in their favorite locations.

Definition 1 (Steady-State Equilibrium). A *steady-state equilibrium* under pricing policy (p, c) is a tuple (f, x, y) s.t.:

- (C1) (Driver best-response) Drivers providing service always maximizes their payoff, i.e. $\forall i, \tau \in L, x_i^{(\tau)} > 0 \Rightarrow \forall k \in L, \pi_i^{(\tau)} \geq I \cdot \mathbb{1}\{i = \tau\} + \delta \pi_k^{(\tau)}$.
- (C2) (Flow balance) For all locations $i \in L$ and driver types $\tau \in L, x_i^{(\tau)} = \sum_{j \in L} f_{ji}^{(\tau)} + y_{ji}^{(\tau)}$.
- (C3) (Market-clearing) $\sum_{\tau \in L} f_{ij}^{(\tau)} = \theta_i \alpha_{ij} (1 - F(p_{ij}))$.
- (C4) (Individually rational driver entry) Participating drivers get at least their outside option w ; with excess supply of drivers with type τ , all participating type- τ drivers get exactly their outside option w .
- (C5) (Feasibility) Rider and driver flows are non-negative, i.e., $\forall i, j, \tau \in L, f_{ij}^{(\tau)}, y_{ij}^{(\tau)}, x_i^{(\tau)} \geq 0$; the supply constraints are satisfied, i.e., $\forall \tau \in L, \sum_{i \in L} x_i^{(\tau)} \leq s^{(\tau)}$.

The *full information first best revenue* (FB) is the highest revenue a mechanism can achieve in stationary-state equilibrium, if the mechanism has full knowledge of driver types (therefore does not need to determine dispatching and compensation in order to incentivize truthful reporting of types):

$$\begin{aligned} \max_{p,c} \sum_{i \in L} \sum_{j \in L} \sum_{\tau \in L} p_{ij} \cdot f_{ij}^{(\tau)} - c_{ij}^{(\tau)} (f_{ij}^{(\tau)} + y_{ij}^{(\tau)}) \quad (2) \\ \text{s.t. } (f, x, y) \text{ is a steady-state equilibrium under } (p, c). \end{aligned}$$

The design problem is to compute rider prices p , driver compensation c , and driver penalties P to optimize platform revenue in the steady state equilibrium, in a way that drivers will truthfully report their location preferences and will choose to always provide service.

3 The Preference-Attentive Ridesharing Mechanism (PARM)

We now introduce our *Preference-Attentive Ridesharing Mechanism (PARM)* and show that this mechanism (i) truthfully elicits drivers' location preferences, (ii) incentivizes drivers to provide service, and (iii) achieves first-best revenue when supply is unconstrained or when demand is symmetric.

3.1 Alternate Form of the Optimization

The optimization problem (2) need not be convex, and moreover, even when an optimal solution can be found, it may not incentivize drivers to report their types truthfully. Denoting $W \triangleq w(1 - \delta)$, we present an alternative problem (3), which guarantees that any optimal solution can be converted into an optimal solution for (2) using compensation scheme (4)—while preserving the objective. Specifically, we consider:

$$\begin{aligned} \max_{p,f,x,y} \left(\sum_{i,j,\tau \in L} f_{ij}^{(\tau)} p_{ij} \right) - W \sum_{i,\tau \in L} x_i^{(\tau)} + I \sum_{\tau \in L} x_\tau^{(\tau)} \quad (3) \\ \text{s.t. } x_i^{(\tau)} = \sum_{j \in L} f_{ji}^{(\tau)} + \sum_{j \in L} y_{ji}^{(\tau)}, \forall i \in L, \forall \tau \in L \\ \sum_{\tau \in L} f_{ij}^{(\tau)} = \theta_i \alpha_{ij} (1 - F(p_{ij})), \forall i, j \in L \\ \sum_{i \in L} x_i^{(\tau)} \leq s^{(\tau)}, \forall \tau \in L \\ \sum_{j \in L} y_{ij}^{(\tau)} = x_i^{(\tau)} - \sum_{j \in L} f_{ij}^{(\tau)}, \forall i \in L, \forall \tau \in L \\ f_{ij}^{(\tau)}, y_{ij}^{(\tau)}, x_i^{(\tau)} \geq 0, \forall i, j \in L, \forall \tau \in L. \end{aligned}$$

Our approach is analogous to a similar move by Bimpikis et al. [2016]—assuming that F is distributed $U[0, 1]$, the solution space is convex, and the optimization problem is quadratic. We also go a step further by accounting for driver heterogeneity and the possibility of zero demand at a location, the latter by paying drivers for relocation dispatches.

Consider the following compensation scheme:

$$c_{ij}^{(\tau)} = W - I \cdot \mathbb{1}\{i = \tau\}, \forall i, j, \tau \in L. \quad (4)$$

Lemma 1. Consider an optimal solution (p, f, x, y) to problem (3), and let c be the compensation scheme (4). Then:

- (i) (p, c) is an optimal solution to optimization problem (2) with steady-state equilibrium (f, x, y) ; and
- (ii) The expected lifetime utility (payment and location value) of a truthful driver is exactly w starting from every location, i.e. $\pi_i^{(\tau)} = w$ for all $i, \tau \in L$.

Briefly, feasible solutions to (3) satisfy conditions (C2), (C3) and (C5). Moreover, with $W > I$, the compensation c as in (4) is non-negative. Given (4), drivers receive utility

W in expectation per period (so w over their lifetimes); this implies (C1) and (C4). Furthermore, the solution is optimal, since no compensation scheme can lower the total payment to drivers while fulfilling the same rider trip flow f .

3.2 Constructing PARM

Definition 2. Given rider demand (θ, α) , the *Preference-Attentive Ridesharing Mechanism* (PARM):

1. Elicits the location preferences from drivers.
2. Solves (3) with an additional constraint

$$x_\tau^{(\tau)} \geq x_i^{(\tau)}, \forall i \in L, \forall \tau \in L \quad (5)$$

for dispatching and pricing, and determine driver compensation by (4).

3. If a driver with reported type- τ did not provide service and relocated to location $i \neq \tau \in L$, the platform treats her as a type- i driver from then on. If this is the first deviation for this driver, the driver pays penalty $P_\tau \triangleq \max\{\max_{k \in L}\{P^{k \rightarrow \tau}\}, 0\}$ for $P^{k \rightarrow \tau}$ as solved for in the following linear system:

$$\begin{aligned} \pi_i^{k \rightarrow \tau} &= \mathbb{1}\{i \neq \tau\} \left(W + \delta \sum_j \frac{f_{ij}^{(\tau)} + y_{ij}^{(\tau)}}{x_i^{(\tau)}} \pi_j^{k \rightarrow \tau} \right) + \\ &\quad \mathbb{1}\{i = \tau\}(\delta w - P^{k \rightarrow \tau}) + \mathbb{1}\{i = k\}I, \forall i, k \in L; \\ w &= \sum_i \pi_i^{k \rightarrow \tau} x_i^{(\tau)} / x^{(\tau)}, \forall k \in L. \end{aligned} \quad (6)$$

The system (6) has $n^2 + n$ linear equations and $n^2 + n$ unknowns (n^2 of the $\pi_i^{k \rightarrow \tau}$ and n of the $P^{k \rightarrow \tau}$). Intuitively, $\pi_i^{k \rightarrow \tau}$ is the expected utility of a driver of type k pretending to be of type τ and providing service everywhere except τ , where she instead relocates to k . By construction, $P^{k \rightarrow \tau}$ is the minimum penalty needed to equalize driver earnings between this deviation and truth telling plus always providing service. We take the maximum over such penalties so no driver can benefit from pretending to be of type τ and employ this strategy.⁷

If a driver declines to provide service but relocates to her reported preferred location, she is charged no penalty. A driver might have a legitimate (idiosyncratic) reason for not being able to provide service in a period, but if she relocates to a location she did not report as preferred, that is taken as an indication that her original report was not truthful.

We now prove, under the assumption that drivers always provide service and as a result are never charged any penalty, that imposing (5) is sufficient to guarantee truthful reporting.

Theorem 1. *Assuming all drivers always provide service, it is a dominant strategy for drivers to report their location preferences truthfully under PARM.*

Proof. Observe that by being truthful, each driver gets utility $W = w(1 - \delta)$ per period—getting paid $W - I$ at preferred

⁷ The penalty P_τ is set to zero if $P^{k \rightarrow \tau} < 0$ for all $k \in L$. This case arises if this deviation is itself bad for drivers of all types, in which case the only way to make the deviating drivers' utility equal to w is to pay those drivers.

locations, and W at every other location. As a result, $\pi_i^{(\tau)} = w$ for all $i, \tau \in L$. Suppose an infinitesimal driver of type $k \in L$ reports she is of type $\tau \neq k$. At all $i \neq k, \tau$ she gains utility W per period. At k , she makes $W + I$ because the platform, treating her as a type- τ driver, is still paying her W . At τ , she is paid $W - I$ and does not get the extra utility I .

With $\delta \rightarrow 1$, misreporting τ in place of k leads to an increase in the expected payoff in static steady-state equilibrium if and only if in equilibrium, the driver with reported type τ spends more time in location k than in location τ . Considering the location of a driver with reported type τ as a Markov chain, then $\{x_i^{(\tau)} / x^{(\tau)}\}_{i \in L}$ is the stationary distribution. (5) then guarantees that a driver with reported type τ spends a plurality of her time at location τ , therefore no driver benefits from misreporting her type if all drivers always provide service. \square

We now consider drivers who may strategically decline to provide service and show such deviations are not useful under PARM, which updates its belief about a driver's type after deviations and imposes a penalty on the first such deviation.

Theorem 2. *Under PARM, it is an ex post Nash equilibrium for drivers to report their types truthfully and to always provide service.*

Briefly, Theorem 1 and the following Lemma 2 imply that (i) a profitable misreport must be paired with post-reporting deviation(s), and (ii) the most profitable deviation must be the driver providing service everywhere except her reported preferred location. The penalties ensure the driver does not get a utility higher than w from this deviation (or any other), so there does not exist a profitable deviation.

Drivers are never charged any penalty under the equilibrium outcome, but the threat of a penalty is necessary to ensure truthful reporting. In certain special economies, a misreporting driver might spend many periods at her true preferred location before being sent to her reported preferred location. Without penalties, she may simply decline service and relocate back to her actual preferred location, thereby sacrificing one period of income for the possibility of many periods of extra idiosyncratic utility. See Appendix B of the full paper for an example and discussions.

Lemma 2. *Consider a driver of true type $k \in L$ and reported type $\tau \in L$, and assume that the rest of the drivers always provide service. If $\tau = k$ (truthful), always providing service is a best response. If $\tau \neq k$, one of the following is a best-response: (i) always providing service, or (ii) providing service at every location except τ , where the driver instead drives to k .*

We now outline the proof of Lemma 2. We first show that a truthful driver earns W at every location, and it is always optimal for her to provide service. An untruthful driver gets the least utility when at her reported preferred location τ , so relocating to τ is worse than providing service; in any period, she should either provide service or relocate to a location $i \neq \tau$. In fact, because she is charged the same penalty for relocating to any location $i \neq \tau$, her optimal relocation is her true preferred location k . Intuitively, if she relocates to i , she will be

at i and treated as type i from then on, which is sub-optimal for her unless $i = k$. Given that the optimal relocation is then her true preferred location—after which she will make W every period—the only location where she might profitably not provide service is τ , her reported preferred location and the only place she currently makes less than W in-period.

3.3 Cases with First-Best Revenue and No Penalty

Although the IC constraint (5) may reduce revenue, we can characterize some settings where imposing IC does not lead to a revenue loss: when supply is unconstrained, or when rider demand is symmetric, PARM achieves the full-information first-best revenue. Furthermore, no penalty is necessary to ensure incentive compatibility.

Theorem 3. *Suppose $s^{(\tau)} = \infty$ for all $\tau \in L$. Then PARM achieves full-information first-best revenue, and no penalty is necessary to ensure incentive compatibility.*

Briefly, the IC constraint (5) does not bind because drivers cost less to the platform when at their preferred locations. If there are more drivers with reported type τ at location i than at τ , the platform can improve its revenue by replacing the type- τ drivers with type- i drivers (there are unlimited type- i drivers). Incentive compatibility holds without penalties because each driver visits her reported preferred location before visiting any other location too many times. Thus, she cannot profitably use a misreport-plus-deviation to sacrifice one period of income for many periods of idiosyncratic utility (as described following Theorem 2). Note that the preceding argument makes no assumption on the demand pattern and requires only the availability of supply.

Definition 3. Rider demand (θ, α) is *symmetric* if $\forall i, j, k, l \in L$ we have $\theta_i = \theta_k$ and $\alpha_{ij} = \alpha_{kl}$.

Theorem 4. *Suppose that rider demand is symmetric. Then we can construct a solution to optimization problem (3) with incentive compatibility constraint (5) such that PARM achieves full-information first-best revenue, and no penalty is necessary to ensure incentive compatibility.*

To understand Theorem 4, we prove two additional lemmas.

Lemma 3. *With symmetric demand, any optimal solution to (3) satisfies $f_{ii}^{(\tau)} + y_{ii}^{(\tau)} \leq f_{\tau\tau}^{(\tau)} + y_{\tau\tau}^{(\tau)}$ for all $i, \tau \in L$.*

Intuitively, type- τ drivers cost less when at location τ , so it is optimal for the marginal ride they give at location τ to have a lower price than at other locations. If demand is symmetric, this means drivers with reported type- τ provide more rides at location τ than at any other location.

Lemma 4. *If the demand pattern is symmetric, we can construct an optimal solution to (3) such that for all $i, j \in L$ and all $\tau \in L$, $f_{ij}^{(\tau)} = f_{ji}^{(\tau)}$ and $y_{ij}^{(\tau)} = y_{ji}^{(\tau)} = 0$.*

Intuitively, there is no need for drivers to relocate when demand is fully symmetric. Moreover, given any optimal solution to (3), we can construct an alternative optimal solution, where the flow of drivers of each type can be decomposed as cycles with length 2, i.e., $f_{ij}^{(\tau)} = f_{ji}^{(\tau)}$.

We can now sketch the proof of Theorem 4. With symmetric demand, Lemma 3 implies driver flow for within-location

trips satisfies the IC constraint (5). For all inter-location trips, Lemma 4 lets us focus only on bilateral driver flow between pairs of locations i and j . Type- τ drivers cost less at τ , so they will naturally fill more rides between τ and j than between i and j —and this holds for all j , so type- τ drivers fill more rides in and out of τ than i . Combining the two cases, drivers of type τ do not spend more time at another location $i \neq \tau$ than they do at τ , so imposing the IC constraint (5) is without loss of revenue. As in Theorem 3, incentive compatibility holds without penalties because each driver will visit her reported preferred location before visiting any other too many times. Thus, she cannot profitably use a misreport-plus-deviation to sacrifice one period of income for many periods of idiosyncratic utility (as described following Theorem 2).

4 Simulation Results

In this section, we use simulations to analyze the revenue and social welfare under PARM for settings outside the cases covered by Theorems 3 and 4—i.e., settings with limited supply and unbalanced demand.


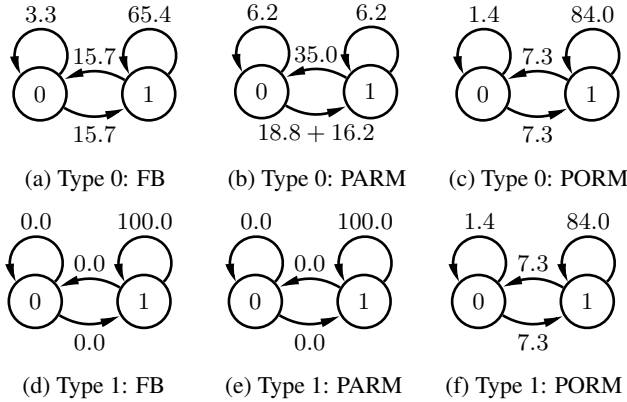

Social welfare is defined as the total rider value plus drivers’ utilities from being in their preferred locations, minus the total opportunity costs incurred by drivers. We compare PARM with the full-information *first-best*, and also a *Preference-Oblivious Ridesharing Mechanism (PORM)* which sets prices as in Bimpikis *et al.* [2016] without considering drivers’ location preferences, while assuming that drivers always follow dispatches. In Section 4.2, we also study the equilibrium outcome under PORM, allowing driver autonomy. For ease of illustration, we consider two locations $L = \{0, 1\}$ throughout the analysis.

4.1 Varying Demand Patterns

Suppose that there are an equal number of drivers favoring each location: $s^{(0)} = s^{(1)} = 100$. Drivers have outside option $w = 40$, discount factor $\delta = 0.99$, and gain utility $I = 0.2W = 0.2w(1 - \delta)$ per period from being in their preferred locations. Each rider has value independently drawn $\sim U[0, 1]$.

Varying Total Demand. We first assume an unbalanced trip flow $\alpha_{00} = \alpha_{10} = 0.25$ and $\alpha_{01} = \alpha_{11} = 0.75$ (i.e., three quarters of riders from each location would like to go to location 1). Fixing the total demand at location 1 at $\theta_1 = 1000$, and varying θ_0 from 0 to 1000, the revenue and welfare under PARM and benchmarks are as in Figure 1. Although PARM only necessarily achieves first-best revenue when $\theta_0 = 1000$ (symmetric demand), we see that PARM achieves the first-best and outperforms PORM unless θ_0 is very small, such that demand from the two locations is highly asymmetric.

When $\theta_0 \ll \theta_1$, almost all rides originate and terminate at location 1, thus the first-best and PORM dispatch most drivers of both types to provide service at location 1. Figure 2 illustrates the rider trip flows fulfilled by drivers of each type under different mechanisms, when $\theta_0 = 50$. To satisfy PARM’s incentive compatibility (IC) constraint, however, drivers of type 0 must spend a plurality of their time at location 0. Therefore, PARM completes fewer trips at location 1, dispatches more type 0 drivers to fulfill (the less profitable) between-location trips, and asks many type 0 drivers to




 Figure 1: Revenue and welfare varying demand θ_0 at location 0.

 Figure 2: Total rider trips fulfilled by drivers of each type, with $(\alpha_{i0}, \alpha_{i1}) = (0.25, 0.75)$, $\theta = (50, 1000)$, and $s = (100, 100)$.

 Figure 3: Revenue and welfare varying α_{i0} for $i = 0, 1$.

relocate back to 1 once they arrive at location 0 (the numbers after the “+” sign represent driver relocation flow), resulting in lower revenue and social welfare.

Varying Imbalance in Demand. Fixing $\theta_0 = \theta_1 = 1000$ and varying α_{i0} for $i = 0, 1$ (i.e., changing the proportion of rides with destination 0), the revenue and welfare achieved by different mechanisms are shown in Figure 3. Similar to Figure 1, PARM achieves first-best revenue and outperforms PORM for a wide range of α_{i0} (although demand is only symmetric when $\alpha_{i0} = 0.5$). For similar reasons as in the above scenario, we see a decline of revenue and welfare under PARM when demand becomes highly unbalanced—in this case, when α_{i0} approaches 0 or 1 and almost all riders have the same destination.

4.2 PORM in Equilibrium

In this section, we analyze a scenario for which we are able to compute the equilibrium outcome given the pricing under PORM, and under the setting where drivers are given the flexibility to decide how to drive. Consider two locations $L =$


 Figure 4: Equilibrium revenue and welfare varying I/W .

 Figure 5: Rider trips fulfilled by type-1 drivers, with $s = (0, 200)$, $\theta = (1000, 1000)$, $\alpha_{00} = \alpha_{11} = 1$, $\alpha_{01} = \alpha_{10} = 0$, and $I = 0.2W$.

$\{0, 1\}$ and drivers of type 1 only: $s^{(0)} = 0$, $s^{(1)} = 200$. All trips start and end in the same location, i.e., $\alpha_{00} = \alpha_{11} = 1$. Being oblivious to drivers’ preferences, PORM sets the same trip price for the two locations and expects the spatial distribution of drivers to be proportional to the distribution of demand. In equilibrium, however, more drivers decide to drive in location 1 (the preferred location), such that in each period drivers in 1 are dispatched with probability less than 1 and achieve the same expected utility as drivers in 0.

Varying Location Preference I . In Figure 4, we fix demand $\theta_0 = \theta_1 = 1000$ and plot revenue and welfare as I , the idiosyncratic driver utility, varies from 0 to W . As I increases, welfare and revenue under PARM coincide with the first-best and increase as expected. However, revenue under PORM (assuming driver compliance) remains constant since the mechanism is oblivious to drivers’ preferences. We also see a decrease in welfare and revenue achieved in equilibrium under PORM, since more drivers decide to supply in location 1, instead of in location 0 as dispatched, resulting in unfulfilled rides in 0 and idle drivers in 1. Beyond $I = 0.5W$, revenue and welfare remain constant, since all drivers are already supplying location 1.

Figure 5 illustrates rider trip flow fulfilled by the type 1 drivers when $I = 0.2W$. PARM assigns more drivers to location 1 than location 0, but PORM does not. However, in equilibrium more drivers end up at location 1 anyway, leading to 25 units of drivers idling at location 1.

Varying Demand Ratio θ_0/θ_1 . In Figure 6, we fix $\theta_1 = 1000$, $I/W = 0.2$, and vary θ_0 from 0 to 2000. We see that PARM revenue coincides with the first-best and significantly exceeds the revenue of PORM. The revenue and welfare of the equilibrium outcome under PORM is much lower, however, because drivers over-supply the preferred location 1, leaving rider trips in 0 unfulfilled. It is curious that with highly unbalanced demand, an increase in θ_0 initially leads to reduced equilibrium revenue and welfare—this is because with higher demand at location 0, PORM sets a higher price at location 1 and accepts fewer location 1 trips in order to com-




Figure 6: Equilibrium revenue and welfare varying θ_1/θ_0

plete more trips in 0. The drivers, however, are only willing to drive in 0 when θ_0 is high enough that the low probability of getting a ride in 1 offsets the extra utility I .

Varying Demand Ratio θ_1/θ_0 . In Figure 7, we set $I/W = 0.2$, $\theta_0 = 1000$ and vary θ_1 from 0 to 2000. We see a similar trend as in Figure 1, with PARM doing worse than even equilibrium PORM for very small values of θ_1 . Figure 8 illustrates driver flow for $\theta_1 = 100 = 0.1\theta_0$ — PARM employs drivers to idle at location 1 in order to satisfy IC, and this is very costly. It is worth noting that as θ_1 increases, the social welfare achieved by the equilibrium outcome under PORM in fact does not increase, due to the increased amount of idle drivers at location 1.





Figure 7: Equilibrium revenue and welfare varying θ_1/θ_0 .



(a) Type 1: PORM (b) Type 1: PORM Eq. (c) Type 1: PARM

Figure 8: Rider trip and idle driver flows, with $\theta = (1000, 100)$, $\alpha = [(1.0, 0.0), (0.0, 1.0)]$, $s = (0, 200)$, and $I = 0.2W$.

Varying Driver Supply s_1 . We now examine the effect of varying the supply of type 1 drivers (while still keeping the supply of type 0 drivers at 0). In Figure 9, we set $I = 0.2W$, $\theta_0 = \theta_1 = 1000$, and vary s_1 from 0 to 1000. Revenue and welfare under PARM coincide with first-best and outperform PORM. All the mechanisms improve in profit and welfare as supply increases, but PARM is better able to use the additional drivers. Under PORM in equilibrium, drivers again over-supply the preferred location 1, causing rides at location 0 to get dropped. Eventually, there are so many drivers that they can fill all the demand, even with drivers idling at location 1. At this point, equilibrium PORM revenue coincides with PORM revenue, though the welfare is still lower.




Figure 9: Equilibrium revenue and welfare varying s_1

5 Discussion

We have proposed the Preference-Attentive Ridesharing Mechanism (PARM) for pricing and dispatch in the presence of driver location preferences. It is an equilibrium under PARM for drivers to report their preferred locations truthfully and always provide service. PARM achieves first-best revenue in settings with unconstrained driver supply or symmetric rider demand, and we show via simulations that even outside those scenarios, PARM achieves close to first-best welfare and revenue and outperforms a mechanism that is oblivious to location preferences.

Our analysis suggests that incorporating drivers' location preferences is compatible with other aspects of ridesharing pricing and marketplace design—even though drivers could in principle game the system by expressing preferences for locations associated with more highly compensated rides. There are two key elements to our approach that both seem likely to provide practical insight beyond the specific framework and mechanism considered here: First, we recognize that respecting drivers' location preferences creates value, which can at least partially substitute for cash compensation. Then, we incentivize truthful location preference revelation through a variation on a revealed preference approach. PARM uses drivers' deviations from proposed dispatches to learn about their preference types—a driver who chooses to drive to i instead of her assigned location is inferred to prefer location i and subsequently faces the compensation profile of other drivers with that preference. For this approach to work, it is important that preferences do not change frequently over the course of the day. Otherwise, it would be much harder to enforce incentive compatibility by tracking endogenous responses to dispatch assignments.

Acknowledgments

We appreciate the helpful comments of Peter Frazier, Jonathan Hall, Hamid Nazerzadeh, Michael Stewart, and the Lab for Economic Design. Rheingans-Yoo gratefully acknowledges the support of a Harvard Center of Mathematical Sciences and Applications Economic Design Fellowship, a Harvard College Program for Research and Science and Engineering Fellowship, and the Harvard College Research Program. Rheingans-Yoo also appreciates the academic inspiration and personal support of Ross Rheingans-Yoo. Kominers gratefully acknowledges the support of National Science Foundation grant SES-1459912, as well as the CMSA's Ng Fund and Mathematics in Economics Research Fund. Ma gratefully acknowledges the support of a Siebel scholarship.

References

- [Afèche *et al.*, 2018] Philipp Afèche, Zhe Liu, and Costis Maglaras. Ride-hailing networks with strategic drivers: The impact of platform control capabilities on performance. Working Paper, 2018.
- [Angrist *et al.*, 2017] Joshua D. Angrist, Sydney Caldwell, and Jonathan V. Hall. Uber vs. taxi: A driver’s eye view. *NBER Working Paper No. 23891*, 2017.
- [Banerjee *et al.*, 2015] Siddhartha Banerjee, Ramesh Johari, and Carlos Riquelme. Pricing in ride-sharing platforms: A queueing-theoretic approach. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 639–639. ACM, 2015.
- [Banerjee *et al.*, 2017] Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. Pricing and optimization in shared vehicle systems: An approximation framework. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 517–517. ACM, 2017.
- [Besbes *et al.*, 2018] Omar Besbes, Francisco Castro, and Ilan Lobel. Spatial capacity planning. Working Paper, 2018.
- [Bimpikis *et al.*, 2016] Kostas Bimpikis, Ozan Candogan, and Daniela Saban. Spatial pricing in ride-sharing networks. Working Paper, 2016.
- [Castillo *et al.*, 2017] Juan Camilo Castillo, Dan Knoepfle, and E. Glen Weyl. Surge pricing solves the wild goose chase. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 241–242. ACM, 2017.
- [Chen and Sheldon, 2015] M. Keith Chen and Michael Sheldon. Dynamic pricing in a labor market: Surge pricing and the supply of Uber driver-partners. Working Paper, 2015.
- [Chen *et al.*, 2017] M. Keith Chen, Judith A. Chevalier, Peter E. Rossi, and Emily Oehlsen. The value of flexible work: Evidence from Uber drivers. *NBER Working Paper No. 23296*, 2017.
- [Cohen *et al.*, 2016] Peter Cohen, Robert Hahn, Jonathan Hall, Steven Levitt, and Robert Metcalfe. Using big data to estimate consumer surplus: The case of Uber. *NBER Working Paper No. 22627*, 2016.
- [Cook *et al.*, 2018] Cody Cook, Rebecca Diamond, Jonathan Hall, John List, and Paul Oyer. The gender earnings gap in the gig economy: Evidence from over a million rideshare drivers. Working Paper, 2018.
- [Hall and Krueger, 2016] Jonathan V. Hall and Alan B. Krueger. An analysis of the labor market for Uber’s driver-partners in the United States. *NBER Working Paper No. 22843*, 2016.
- [Hall *et al.*, 2015] Jonathan V. Hall, Cory Kendrick, and Chris Nosko. The effects of Uber’s surge pricing: A case study. Working Paper, 2015.
- [Hall *et al.*, 2017] Jonathan V. Hall, John J. Horton, and Daniel T. Knoepfle. Labor market equilibration: Evidence from Uber. Working Paper, 2017.
- [Ma *et al.*, 2019] Hongyao Ma, Fei Fang, and David C. Parkes. Spatio-temporal pricing for ridesharing platforms. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. ACM, 2019.
- [Ostrovsky and Schwarz, 2019] Michael Ostrovsky and Michael Schwarz. Carpooling and the economics of self-driving cars. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. ACM, 2019.